Cheap Mass:
Learning, Scale, and the Technology of Conflict

Martin Ljunge*

August 2026

Abstract

Military power depends both on a state's fiscal resources and on the industrial technology that turns spending into effective force. Cheap, mass-produced, attritable drones sharply raise the productivity of military spending. I model the shift in two layers: an *industrial* layer mapping money into delivered military effect, and a *contest* layer mapping effect into the probability of prevailing.

Three results follow. First, equilibrium win probability is a side's share of the productivity of spending, so symmetric cheapening raises destructiveness per dollar without raising equilibrium budgets. Second, learning-by-doing makes force production exhibit increasing returns, which turns the *decisiveness* of the contest, a primitive in contest theory, into a derived function of a measurable learning rate; past a threshold rate no stable interior arming equilibrium exists, and a side able to commit production first deters the other outright. Third, a rapid and anticipated shift in industrial power can create a commitment problem, and under a bounded-transfer bargaining protocol I derive the exact condition under which the declining side prefers preventive war.

Affordability and fixed-budget constraints preserve a role for wealth; among sides able to arm at their unconstrained optimum, relative military power is governed by the productivity rather than the level of spending. The model is calibrated to Ukraine-era and US procurement costs, and at the benchmark learning rate the arming game exhibits increasing returns but sits below the decisive-race threshold. Data from the Russian long-range UAV campaign reject the fixed delivered-effect target that the simplest replacement-learning prediction requires. The broader implication is that an arms race may appear in production, iteration and capability without appearing in defense budgets.

Keywords: economics of conflict; contest success functions; learning-by-doing; defense economics; drones.
JEL: D74, H56, O33, L64.

1 Introduction

For three decades the dominant model of military effectiveness in advanced states was *precision*: a small number of highly accurate, expensive, slow-to-produce platforms and munitions, each capable of destroying a target with near certainty. The war in Ukraine has revealed an alternative that earlier doctrine treated as a contradiction. First-person-view (FPV) quadcopters costing a few hundred dollars, long-range one-way-attack drones costing tens of thousands, uncrewed surface vessels, and ground robots have substituted for, and in places displaced, exquisite precision

*Research Institute of Industrial Economics (IFN), Stockholm. Email: martin.ljunge@ifn.se. Working paper; comments welcome.

inventories.[1] Plichta and Rossiter (2024) call the combination "affordable mass precision." The economic content of the shift is a collapse in the *cost of delivering accurate destructive effect at distance*, together with the fact that this capability is now *scalable and attritable* rather than scarce and precious: a $400 FPV drone can destroy a tank valued at several million dollars, and a $50,000 one-way-attack drone can force the expenditure of a $3–4 million interceptor.

This paper asks how that shift changes the economics of conflict: the mapping from resources to military effectiveness, equilibrium arming, the offense–defense balance, the industrial organization of defense, and, by extension, the likelihood and character of war. The economics of conflict has rich tools for these questions (contest success functions, arms-race dynamics, the guns-versus-butter tradeoff, and bargaining models of war) but it has not formally modeled the cheap-mass/attritable-precision transition. The best existing analysis sits in strategic studies and international relations. This paper brings the transition inside an economic model.

Approach. I build the model in two layers. An *industrial layer* maps military spending into effective force through a cost-minimizing choice of weapon mix, summarized by a single object: the unit cost of effective force P (equivalently its inverse, the *productivity of spending* $\psi = 1/P$). A *contest layer*, a Tullock contest success function (Tullock, 1980; Hirshleifer, 1989; Skaperdas, 1996), maps forces into the probability of prevailing and, with guns-versus-butter, into equilibrium arming. The separation yields the paper's organizing identity: a side's probability of prevailing depends on its military budget only through the effective force that budget buys. The entire cheap-mass paradigm is therefore a statement about the industrial layer, a fall in the unit cost of force achieved by many cheap, attritable, learning-driven units. The baseline arsenal is a single weapon fielded in chosen number and quality (the quantity–quality trade-off of Becker and Lewis, 1973); I develop its comparative statics in a tractable two-good form that connects to the cost-exchange data, feeding successively richer industrial layers (exogenous cost, then learning, then an arms race, then a full two-stage production game) through the same contest.

Main results. (1) The equilibrium win probability equals a side's *share* of the productivity of spending, the effective force each dollar buys (Lemma 2). (2) A fall in the price of mass favors an asymmetric adopter and re-tools its arsenal toward mass (Proposition 2); but when both sides cheapen, equilibrium *budgets are unchanged* while destructiveness per dollar rises without bound (Proposition 3), an arms race in capability, not spending. (3) Learning-by-doing gives the cheap-mass technology increasing returns: because each unit built lowers the cost of the next, doubling the mass budget more than doubles the force it fields (Proposition 4). Increasing returns sharpen the contest, and beyond a threshold rate of learning the race turns *decisive*: no stable interior arming equilibrium exists, and a first mover able to commit production wins outright (Propositions 5 and 6). (4) Whether cheap mass is durable is a race between offensive and defensive learning curves (Proposition 7); and that race runs only where weapons are *consumed*: experience accrues only where an operational requirement forces reconstitution, so a weapon held in inventory does not descend its own learning curve whatever its learning rate; and how much any given source of depletion lowers cost depends on how large a share of accumulated experience it generated (Lemma 3, Proposition 8). The distinction is operational rather than technological, and which regime a campaign occupies is an empirical question the paper takes to data (Appendix D). (5) In

[1] Two terms recur throughout and are used in their defense-community sense. *Exquisite* describes small numbers of high-cost, high-capability systems bought on long timelines. *Attritable* describes systems cheap enough that losing one is an accepted cost of the mission rather than the erosion of a scarce stock. The economic content of the paper is that the second category has become capable enough to substitute for the first.

the full two-stage production game the shape of cheap-mass adoption turns on how strong learning is relative to how substitutable mass and precision are. When learning is weak the cost-minimizing arsenal is a mix at every scale, with the mass share rising continuously as the force grows; when learning is strong adoption is instead *discontinuous*, pure precision below a threshold scale and a jump to a mass-heavy mix above it (Proposition 9). Higher stakes tip the arsenal toward cheap mass, discretely in the strong-learning case (Proposition 10), and the plane of stakes against learning rates partitions into precision, mass, and decisive-race regions (Proposition 11). (6) In a bargaining model of war, cheap mass does not mechanically cause war: a productivity advantage re-centers the peace range rather than closing it (Proposition 14), and with war costs derived from the model itself, symmetric cheap mass *widens* the static peace range, because arming is unchanged while destruction per dollar rises (Proposition 15). The war risk is in the *dynamics*: fast, learning-driven, potentially decisive shifts in relative power are exactly the conditions the bargaining literature identifies as most war-prone (Section 11).

Contribution. The paper integrates three literatures that have not been combined for this problem: contest theory in the economics of conflict (Hirshleifer, 1991a; Skaperdas, 1996; Garfinkel and Skaperdas, 2007); the endogenous-cost economics of learning-by-doing (Wright, 1936; Arrow, 1962; Thompson, 2012); and the formalization of the offense–defense balance through contest functions (Garfinkel and Dafoe, 2019). It converts the "affordable mass precision" observation of Plichta and Rossiter (2024) into a tractable economic framework with comparative statics and testable predictions, and it connects to the diffusion of military technology (Horowitz, 2010; Fuhrmann and Horowitz, 2017; Gilli and Gilli, 2016) and to general-purpose-technology accounts of military change (Bresnahan and Trajtenberg, 1995; Ding and Dafoe, 2023). Section 2 sets out the stylized facts; Section 3 the related literature; Sections 4–8 the static model; Section 9 adds an explicit experience state and solves a two-period learning race; Section 10 calibrates the model; Section 11 embeds it in the bargaining theory of war; Section 12 discusses and Section 13 concludes. Proofs are in Appendix A; Appendix B rederives every result natively on the multiplicative base; the global analysis of the arsenal problem, including the adoption threshold and the dropout boundary, is in Appendix C; and Appendix D asks which learning regime the largest one-way-attack campaign on record occupies.

2 The drone paradigm: stylized facts

The following facts (Ukraine, 2022–2026) motivate and later calibrate the model. Wartime operational figures are self-reported, contested, and fast-moving; they are orders of magnitude, not point estimates.

Cheap mass at the front. FPV drones cost roughly $300–600 per unit. Ukrainian output scaled from a few thousand units in 2022 to a 2025 flow of roughly three million, with capacity claims near eight million by early 2026; monthly output rose from about 20,000 (early 2024) to about 200,000 (2025). The three-million figure is a *delivery* count: Ukrainian forces received about three million FPV drones during 2025, some 2.5 times the 2024 total, of which the Defence Procurement Agency supplied 2.4 million directly (Ministry of Defence of Ukraine, 2025). Deliveries and production need not coincide, and the distinction matters for any learning curve fitted to these series. Long-range one-way-attack drones cost as little as $55,000, with ranges of 1,000–2,000 km; naval drones sank warships and downed aircraft; and nearly 15,000 ground robotic systems were delivered to units in 2025 (Ministry of Defence of Ukraine, 2025), taking over a growing share of frontline logistics and casualty evacuation. Cleaned cost points span $220–775 for FPVs (median $400), about $6,000 for

an AI-guided one-way-attack drone, $35,000–$80,000 for a Shahed/Geran (median $59,000), and $1–3.5 million for the exquisite interceptors fired against them (Center for Strategic and International Studies, 2025).[2]

The use–produce–develop loop. Ukraine's Brave1 defense-tech cluster, run by the Ministry of Digital Transformation, is the organizational core of the drone effort. It coordinates a broad ecosystem of competing developers (over 1,000 participating companies and more than 1,600 registered technological solutions within the platform's first year of operation, and over 750 grants awarded by mid-2026) and, through its "Test in Ukraine" program and an "Amazon-for-the-military" marketplace, runs many designs in parallel and feeds combat results back to developers within days (KSE Institute, 2024; Brave1 Defense Tech Cluster, 2026). Cluster headcounts are self-reported and move quickly; the figures above are the dated, separately corroborated ones, and the argument below needs only the order of magnitude and the parallel structure. The innovation cycle compresses from years to weeks, so effectiveness is a moving target set by *iteration speed* and by the *breadth* of experimentation, not by a fixed platform specification. This decentralized, many-supplier architecture, contrasted with an adversary that concentrates production in fewer variants, returns in Section 12 as an endogenous determinant of the learning rate.

A live countermeasure race. Electronic warfare reportedly downs a majority of drones, provoking fiber-optic (unjammable) links and, increasingly, onboard autonomy, so the value of any given drone depends on the adversary's evolving defenses. Two things travel under the latter name and only one of them enters the model. *Onboard* autonomy lets an individual round finish its mission when the link to its operator is jammed; it is fielded, and Ukraine's mid-range strike campaign is credited in part to it (O'Brien, 2026b). *Theatre-level* artificial intelligence, integrating sensing and targeting across a campaign, is a larger and more contested claim: the 2026 US air campaign against Iran was conducted under an explicit "AI-first" doctrine and is judged by some observers to have targeted poorly (O'Brien, 2026a). Only the first is modeled below, where it raises the effect a cheap unit delivers (Remark 1).

Cost-exchange inversion. Figure 1 summarizes the assembled cost panel (63 source rows; Center for Strategic and International Studies, 2025; Cancian and Park, 2026 and sources therein, plus the FY2027 US procurement justification books). The key economic object is the cost to attack relative to the cost to defend or to the value of the target destroyed. The figure keeps those comparisons separate. Three features stand out. First, unit costs span nearly five orders of magnitude: buying a given amount of destructive effect through cheap mass costs far less than buying it through exquisite precision (panel a). That gap between the two *prices of delivered effect* is the primitive the model turns on, and Section 4.3 gives it notation. Second, the interception exchange for a *given* threat, the ratio of what defeating it costs the defender to what fielding it costs the attacker, depends on the defense it meets: a Shahed one-way-attack drone is offense-favored against exquisite interceptors (defeating it costs 17–55× its own price in raw cost ratios, before adjusting for the interceptor's kill probability, which raises the defender's cost per kill further) but *defense*-favored against a cheap interceptor drone (whose panel-median price is about 1/28 of the Shahed), foreshadowing the arms-race logic of Section 7. Third, target value per attacker dollar is a distinct exposure measure whose expected-damage interpretation depends on the probability of destroying the target. Table 1 records the resulting first calibration of the price primitives.

Beyond Ukraine. The same economics ran in 2026 in a second theatre, against the world's most exquisite defensive inventory. In the war between the United States and Iran, Shahed-class one-way-attack drones costing tens of thousands of dollars were met by PAC-3 MSE interceptors

[2]Conflict-event data record on the order of 77,000 and 32,000 air/drone-strike events in Ukraine and Russia, respectively, over the war (Armed Conflict Location & Event Data Project , ACLED).

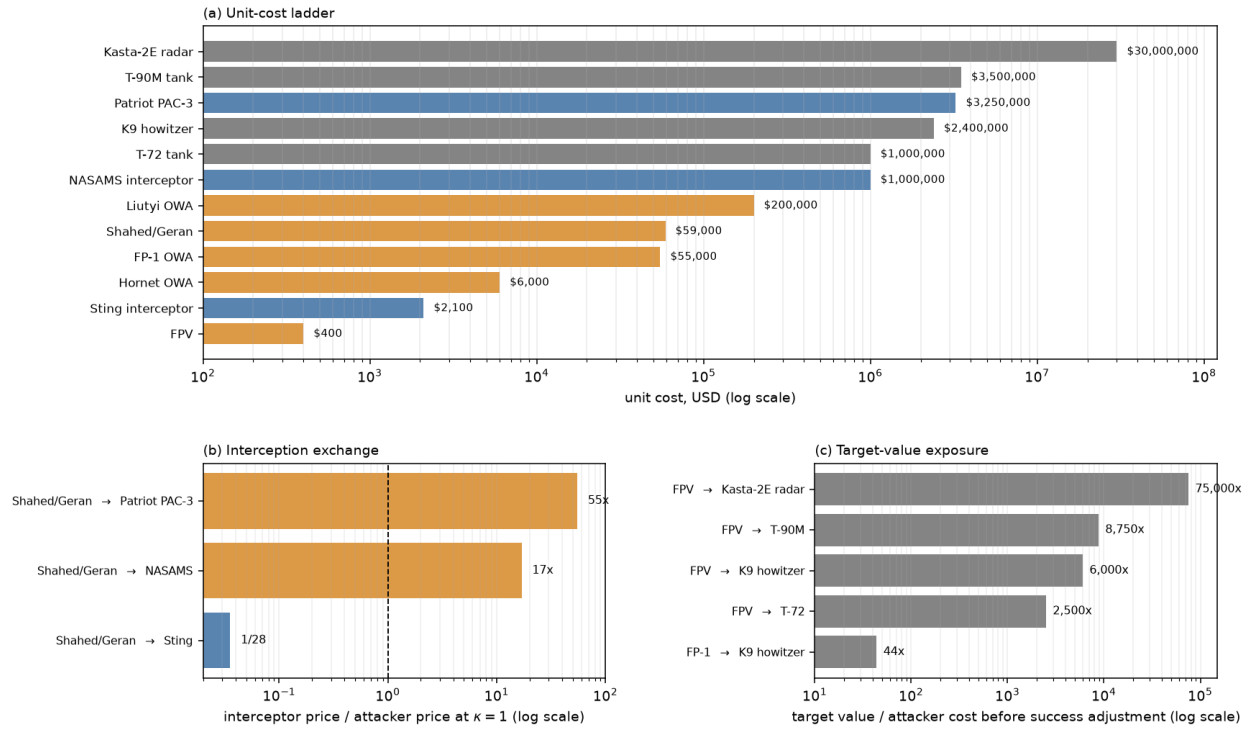


Figure 1: Three empirical cost comparisons from the assembled drone cost panel (Ukraine 2022–2026). (a) The unit-cost ladder spans nearly five orders of magnitude, from a $400 FPV to a $30 m radar. (b) Interception exchange, defender interceptor price divided by attacker price at $\kappa = 1$. These bars are lower bounds on defender cost per kill when interceptor kill probability is below one. Against a Shahed threat the exchange favors offense for Patriot and NASAMS (17–55×) and defense for the Sting interceptor drone (about 1/28 in raw prices and 1/19 at its reported 0.68 kill probability). (c) Target value per attacker cost before adjustment for the probability of destroying the target. These exposure multiples are upper bounds on expected target damage per attacker dollar when that probability is below one. The plotted T-90M benchmark is $3.5 m.

at roughly $4 m and THAAD interceptors at several times that. At the April 2026 ceasefire, CSIS estimated that seven key US munitions had been drawn down to levels requiring one to four years to rebuild, against manufacturing lead times of roughly 52 months from appropriation to delivered round (Cancian and Park, 2026). Reported depletion deepened after the ceasefire lapsed, and the figures are open-source estimates that the Department of Defense has disputed; the episode is used here for its qualitative pattern, not its point values. That pattern is the one this paper models: a cheap and rapidly produced threat imposing cost on a defense whose replacement rate is set by a procurement pipeline measured in years. It also motivates a distinction developed in Section 7, between holding a magazine and consuming one.[3]

[3]Unit costs for the Shahed family are themselves contested. The $20,000–$50,000 range usually quoted prices the Russian-built variant; Batmanghelidj (2026) argues that Iranian factor costs imply a production cost closer to $7,000. Either figure leaves the cost-exchange conclusion intact, and a lower one strengthens it.

Table 1: First calibration of the price primitives to Ukraine-era costs (panel medians; ranges in the text). Plain-language labels keep the stylized-facts section self-contained; Sections 4.3 and 7 introduce the formal notation.

Economic input	Data proxy	Value
Cheap-mass unit cost	FPV (median)	\$400
One-way-attack unit cost	Shahed/Geran (median)	\$59 k
Cheap-defense unit cost	Sting interceptor drone (median)	\$2,100
Precision unit cost	Patriot PAC-3	\$3.25 m
Mid-tier interceptor cost	NASAMS/AIM-9X	\$1 m
Interceptor kill probability	Sting success rate	0.68 (0.60–0.90)
Mass/precision price gap	cost ladder	$\approx 8{,}000\times$

3 Related literature

Economics of conflict. The organizing abstraction is Hirshleifer's (1991b) "technology of conflict", a mapping from resources devoted to fighting into relative success, and the *paradox of power* (Hirshleifer, 1991a), that the shape of this technology governs whether the weaker side can contest effectively. The workhorse is the contest success function (CSF): Skaperdas (1996) axiomatizes it, Hirshleifer (1989) distinguishes ratio and difference forms, and Tullock (1980) supplies the decisiveness parameter. Garfinkel and Skaperdas (2007) survey the field and Jia et al. (2013) treat estimation. The re-emergence of mass also has a classical antecedent: Lanchester's (1916) square law makes numbers count disproportionately in aimed-fire attrition (see Kress, 2020 for modern variants); the contest layer here delivers the economic counterpart, with the decisiveness of numbers emerging from learning-driven cost rather than being assumed in the combat dynamics. Closest in spirit, Hwang (2012) studies how the technology of conflict shifts guns, butter, and war in general equilibrium; the present paper opens the technology itself, deriving it from arsenal design and production experience. I place the paradigm shift inside the CSF as a change in the productivity of spending and, through learning, in the effective decisiveness of the contest.

Defense economics and arms races. Within the defense-economics literature (Hartley and Sandler, 1995), Brito and Intriligator (1995) formalize arms races as coupled investment paths; the drone–counter-drone race is their modern instance. Two strands are especially close to the industrial layer here. Anton and Yao (1987) put the experience curve at the center of defense procurement, studying price competition between learning suppliers; this paper puts the same curve on the battlefield side, where the buyer's rival, not a second source, disciplines the race down the curve. Blanken and Lepore (2011) model strategic timing in military-technology competition, a complement to the industrial-lead result of Section 6.2. Hall and Coyne (2014) give a political-economy account of the U.S. drone industry's procurement trajectory. The drone era also contrasts with the slow, capital-intensive procurement that defense economics studies, suggesting that procurement structure is itself an economic variable the paradigm changes.

Offense–defense and bargaining. Jervis (1978) founds offense–defense theory and Glaser and Kaufmann (1998) make the balance measurable; most useful here, Garfinkel and Dafoe (2019) formalize the balance with contest functions and show it *scales* with investment. Fearon (1995) and Powell (2006) supply the bargaining channel through which the changed cost and speed of conflict

affect the probability of war; Chang and Luo (2017) integrate arming, endogenous destruction, and settlement in a guns-versus-butter bargaining model, the closest formal antecedent of Section 11, where the mechanism differs: destruction there is a choice along a fixed technology, while here it rises through the productivity of the industrial layer at unchanged budgets.

Learning and general-purpose technology. Wright (1936) and Arrow (1962) give the learning curve; Thompson (2012) reappraises the evidence. Bresnahan and Trajtenberg (1995) and Ding and Dafoe (2023) frame autonomy/AI as a possible general-purpose (military) technology, cautioning that current cost-exchange ratios need not be permanent.

Quantity versus quality, and "good enough." The trade-off at the center of this paper, many cheap units versus few exquisite ones, is the defense counterpart of the quantity–quality problem that Becker and Lewis (1973) placed at the heart of the economics of fertility and human capital (also Becker, 1960; Becker and Tomes, 1976). Because per-unit quality is paid on *every* unit, the shadow price of quality rises with fleet size exactly as in their multiplicative budget, so a learning-driven fall in the unit cost of mass amplifies the substitution toward quantity (Section 4.2). Two features distinguish the military problem: the payoff to (n, ζ) is a *contest* against a strategic adversary rather than household utility, and the price of quantity is *endogenous* through learning. The viability of cheap mass rests on accuracy being "good enough" rather than maximal, the satisficing logic of Simon (1955), and on the disruptive dynamic of Christensen (1997), in which cheaper, good-enough products displace over-engineered incumbents from below: Ukraine's drones and cruise-missile substitutes are "good enough" at a small fraction of the cost of exquisite US/NATO equivalents. The precision paradigm being displaced is instead close to Kremer's (1993) O-ring production, where quality is complementary and one high-quality munition suffices; the drone shift is a move from an O-ring regime to a Becker quantity–quality regime.

Diffusion and the phenomenon. Horowitz (2010) shows adoption depends on organizational capital as well as cost; Fuhrmann and Horowitz (2017) and Gilli and Gilli (2016) study drone diffusion. Plichta and Rossiter (2024) and Plichta (2025) describe the one-way-attack revolution. Economics is not silent on adjacent questions: conflict-technology shifts (Hwang, 2012), endogenous destruction in bargaining (Chang and Luo, 2017), experience curves in procurement (Anton and Yao, 1987), and the political economy of drone acquisition (Hall and Coyne, 2014). What none of these supplies, and what this paper contributes, is a force-production technology in which the quantity–quality margin and production experience jointly determine the productivity and the effective decisiveness of the contest, the mapping the drone era has changed.

4 The model

4.1 Environment

Two players $i \in \{A, B\}$ contest a stake of value $V > 0$. Each is endowed with resources W_i and divides them between military spending ("guns") $B_i \geq 0$ and civilian consumption ("butter"), which takes the residual, subject to the *resource constraint*

$$B_i \leq W_i. \tag{1}$$

Payoff is linear in civilian consumption, $U_i = p_i V + (W_i - B_i)$, which gives the reduced form

$$U_i = p_i\, V - B_i \quad (+\, W_i), \tag{2}$$

where p_i is the probability i prevails and the constant W_i is immaterial to the choice. Arming is therefore *not* capped at a fixed budget: each dollar of guns has a unit opportunity cost in foregone civilian goods, the guns-versus-butter margin, and the equilibrium B_i is pinned down by the first-order conditions below. (Endowments are taken large enough that the interior B_i is affordable, $B_i \le W_i$.) A hard defense-budget cap $B_i \le \bar{B}_i$ is an alternative closure, closer to appropriations practice; where it binds it turns the model into a fixed-budget contest and, with the multiplicative arsenal, delivers the Becker fixed-income amplification (Remark 2). The two stages are industrial choice and outcome realization: in Stage 1 players choose production and technology, fixing the unit cost of force; in Stage 2 the contest resolves given the fielded forces.

An arsenal turns spending into *effective force* X_i through a Stage-1 technology with cost $B_i = C_i(X_i)$. The *baseline* technology is the quantity–quality arsenal at the heart of the paper: a side fields n_i units each of quality $\zeta_i \ge 0$, with combat power

$$X_i = n_i\, \zeta_i^{\,b}, \qquad b \in (0,1), \tag{3}$$

where $b < 1$ encodes "good-enough" accuracy, per-unit quality has diminishing returns, so beyond a point extra precision adds little to the outcome.[4] "Mass" means many low-quality units, "precision" few high-quality ones. Following the quantity–quality logic of Becker and Lewis (1973), per-unit quality is paid on *every* unit, so the arsenal-composition constraint is *multiplicative*; with learning-by-doing on the base unit cost (Wright, 1936; Arrow, 1962),

$$B_i = \underbrace{c_{m0}\, n_i^{\,1-\beta}}_{\text{base cost, Wright learning } (\beta)} + \underbrace{w\, n_i\, \zeta_i}_{\text{quality, paid per unit}}. \tag{4}$$

For calibration it is convenient to work with a two-good version of the same trade-off (Section 4.3) that replaces the continuous quality dial by two graded goods: a cheap drone m (low effectiveness e_m, low price c_m) and an exquisite munition q (high e_q, high c_q), combined in a CES aggregator with substitution $\sigma > 1$ and $\rho = (\sigma - 1)/\sigma$ under an *additive* arsenal-composition constraint:

$$X_i = \left[\, (e_m m_i)^\rho + (e_q q_i)^\rho \,\right]^{1/\rho}, \qquad c_m\, m_i + c_q\, q_i = B_i. \tag{5}$$

Here "mass vs. precision" is the mix of two priced goods: a low-quality unit of the baseline corresponds to good m, a high-quality unit to good q. This version is convenient because its two prices per unit of effectiveness, $\pi_m = c_m/e_m$ and $\pi_q = c_q/e_q$, map onto the distinct systems the cost-exchange data report (Figure 1); Appendix B shows the two versions deliver the same results, and identifies the two objects that generalize.

For a specified common mission, let e_j denote expected delivered effect per procured unit and define the *precision-effectiveness advantage*

$$\mathcal{E} \;\equiv\; e_q/e_m \;>\; 1, \qquad \text{so that} \qquad \pi_q/\pi_m = (c_q/c_m)/\mathcal{E}. \tag{6}$$

Survival, acquisition, and conditional success enter only insofar as they are incorporated in e_j; $\mathcal{E}$ is therefore the entire wedge between a system's sticker price and the effect it delivers, and it is the object the calibration of Section 10 must assume rather than estimate.

Two ratios are now in play and it is worth being explicit about how they meet, since one counts effect and the other money. Mass is the cheaper way to buy delivered effect when $\pi_m < \pi_q$, that

[4] A microfoundation: if a unit's contribution is its kill probability $P_k(\zeta)$, bounded by one and concave in quality inputs, returns to per-unit quality diminish and $b < 1$ obtains. The O-ring pole instead makes quality components *complementary*, so a single low-quality component destroys the unit's contribution (Kremer, 1993).

is when $c_m/e_m < c_q/e_q$, which rearranges to $\mathcal{E} < c_q/c_m$: precision earns its price only if it is more times more *effective* than it is more times more *expensive*. A precision round that costs thirty times a drone must deliver more than thirty times the effect per round, or the drone is the better buy. Everything the calibration does is a comparison of those two multiples.

Remark 1 (Autonomy lowers $\mathcal{E}$). The microfoundation of $b < 1$ prices one prediction worth stating, because the parameter it moves is the one the calibration cannot pin. If a unit's contribution is a kill probability bounded by one and concave in quality inputs, a technology that improves guidance on *every* unit raises delivered effect by more where the unit starts further from the bound. Onboard autonomy is such a technology: it substitutes for the operator link that electronic warfare severs, and it does so on the cheap units that had the least accuracy to begin with. Mass therefore gains proportionally more than precision, which is already near the ceiling, and $\mathcal{E} = e_q/e_m$ *falls*. The gain is also unusually cheap to spread. Quality in the baseline is paid on *every* unit — the $w\,n\,\zeta$ term of (4), and the reason the shadow price of quality rises with fleet size — but autonomy is in large part software, developed once and copied across the fleet. It is not free at the unit level — it needs sensors, compute, power and integration, and each of those is paid per airframe — but a large share of its development cost is not. The condition is therefore exactly

$$d\ln e_m \;>\; d\ln c_m,$$

effectiveness rising proportionally faster than unit cost, which is what makes $\pi_m = c_m/e_m$ fall. Because the arsenal tips to mass exactly when $\mathcal{E} < c_q/c_m$, autonomy can carry a mission across that boundary without the *price ratio* c_q/c_m moving at all: the crossing comes from the effectiveness side, and needs no change in stakes or budgets either. It works in the same direction as learning, which lowers c_m directly, so the two compound on the two arguments of π_m.

Nor need the improvement arrive from outside. If effectiveness itself improves with cumulative production — the effectiveness elasticity α_e of Remark 6, read as learning rather than saturation — then the side producing more mass is also the side improving it faster, and by the concavity above it gains more per increment than the precision pole can. On that reading $\mathcal{E}$ falls *endogenously* as mass is built, and the organization that learns (Section 6.4) buys not only a lower c_m but a higher e_m — the second worth more at the cheap pole than at the exquisite one.

4.2 The baseline: quantity, quality, and the cost of force

Cost-minimizing arsenal design under (3)–(4) sets both how many units to field and how good each should be; the contest that turns the resulting force into a win probability is stated in Section 4.4. Because quality is paid on every unit, the shadow prices display the Becker–Lewis interaction: $\pi_n \equiv \partial B/\partial n = (1-\beta)c_{m0}n^{-\beta} + w\zeta$ rises with quality, and $\pi_\zeta \equiv \partial B/\partial \zeta = w\,n$ rises with quantity.

Proposition 1 (Baseline cost of force and quantity–quality amplification)**.** *Cost-minimizing design in* (3)–(4) *gives quality per unit, fleet size, and unit cost of force*

$$\zeta^*(X) \;\propto\; \left(\frac{c_{m0}}{w}\right)^{1/(1-b\beta)} X^{-\beta/(1-b\beta)}, \qquad n^*(X) \;\propto\; \left(\frac{w}{c_{m0}}\right)^{b/(1-b\beta)} X^{\,1/(1-b\beta)},$$

$$C(X) \;\propto\; c_{m0}^{(1-b)/(1-b\beta)} w^{b(1-\beta)/(1-b\beta)} X^{(1-\beta)/(1-b\beta)}.$$

Hence: (i) without learning ($\beta = 0$) quality per unit is scale-invariant*; (ii) with learning ($\beta > 0$) the cost of force has* increasing returns *(the cost-of-force exponent $(1-\beta)/(1-b\beta) < 1$), with effective contest decisiveness equal to its reciprocal $(1-b\beta)/(1-\beta) > 1$, and a larger required force is*

met with more, lower-quality *units (ζ^* falls, n^* rises more than one-for-one): the quantity–quality substitution is amplified by scale. The limits in b span the poles: as $b \to 0$ effective decisiveness is $1/(1-\beta)$ (pure mass), and as $b \to 1$ it is 1: quality contributes to force one-for-one, so the composition margin, and with it the increasing returns, fades.*

Which exponent, and what "decisiveness" means. The increasing returns live in the *cost-of-force* exponent. Because $b < 1$ makes $1 - b\beta > 1 - \beta$, the exponent $(1-\beta)/(1-b\beta)$ on X in $C(X)$ is below one: a larger force costs *less* than proportionally to field. The other two exponents point the opposite way, the fielded number n^* rises *more* than proportionally $(1/(1-b\beta) > 1)$ while quality per unit ζ^* falls $(-\beta/(1-b\beta) < 0)$, so the extra spending buys more units, not better ones. The *effective contest decisiveness* is the reciprocal of that cost exponent, $(1-b\beta)/(1-\beta)$. In a Tullock contest, decisiveness is the exponent with which relative resources map into the probability of winning (the R of the CSF, §4.4): because learning makes force scale super-linearly in spending, $X \propto B^{(1-b\beta)/(1-\beta)}$, win probability responds to *budgets* as if the contest were sharper than its raw R. The higher this value, the more a small spending lead tips the outcome; once it exceeds 2 no interior arming equilibrium survives (Proposition 5).

The paradigm shift. A fall in the base unit cost c_{m0}, cheaper drones through commercial supply chains and learning, lowers the unit cost of force and, because quality is paid per unit, lowers optimal quality per unit ($\zeta^* \propto c_{m0}^{1/(1-b\beta)}$) while raising the number fielded: the arsenal tips toward cheap mass. The tip is amplified by scale and stakes, higher V raises equilibrium force and, through $\pi_\zeta = wn$, the marginal cost of quality. The "good-enough" property $b < 1$ makes the substitution viable; as the returns to per-unit quality approach constant ($b \to 1$) it vanishes, and under O-ring complementarity across quality components it is removed entirely. Unlike the fixed-income amplification of Becker and Lewis (1973), here spending is endogenous (guns-versus-butter), so the amplification runs through learning-driven scale rather than a binding budget line, a distinction the hard-cap Remark 2 makes sharp.

Remark 2 (Hard budget cap and the fixed-income amplification). Replace the guns-versus-butter margin (1) with a hard defense budget $B_i \le \bar{B}_i$. Where the cap binds (below the unconstrained optimum) the side spends $\bar{B}_i$ in full, and the contest is fixed-budget: the two results established below — that the win probability is the productivity share (Lemma 2) and that cheaper mass raises destructiveness rather than budgets (Proposition 3) — hold with budgets now *literally* fixed. Under the multiplicative arsenal this is exactly the Becker–Lewis fixed-income problem, maximize $X = n\zeta^b$ subject to $c_{m0}n^{1-\beta} + wn\zeta = \bar{B}$, whose solution is the dual of Proposition 1:

$$X^* \propto \bar{B}^{\,(1-b\beta)/(1-\beta)}, \qquad \zeta^* \propto \bar{B}^{-\beta/(1-\beta)}, \qquad n^* \propto \bar{B}^{\,1/(1-\beta)}.$$

So force rises *more than proportionally* with the budget: the exponent $(1-b\beta)/(1-\beta)$ exceeds one whenever $b < 1$, being the reciprocal of the cost-of-force exponent that Appendix B shows is below one. And *a larger budget optimally tips the arsenal toward cheap mass* (ζ^* falls, n^* rises more than proportionally): a well-funded military facing learning should buy *more, cheaper* units, not fewer exquisite ones.

4.3 The additive two-good formalization (calibration)

I develop the paradigm's comparative statics in the two-good version of the environment, where they take their sharpest closed forms and connect directly to the cost-exchange data; Appendix B confirms that each result carries over to the quantity–quality baseline. Minimizing $c_m m + c_q q$ for a given effective force (5) yields a quality-adjusted price index. With constant unit prices (no

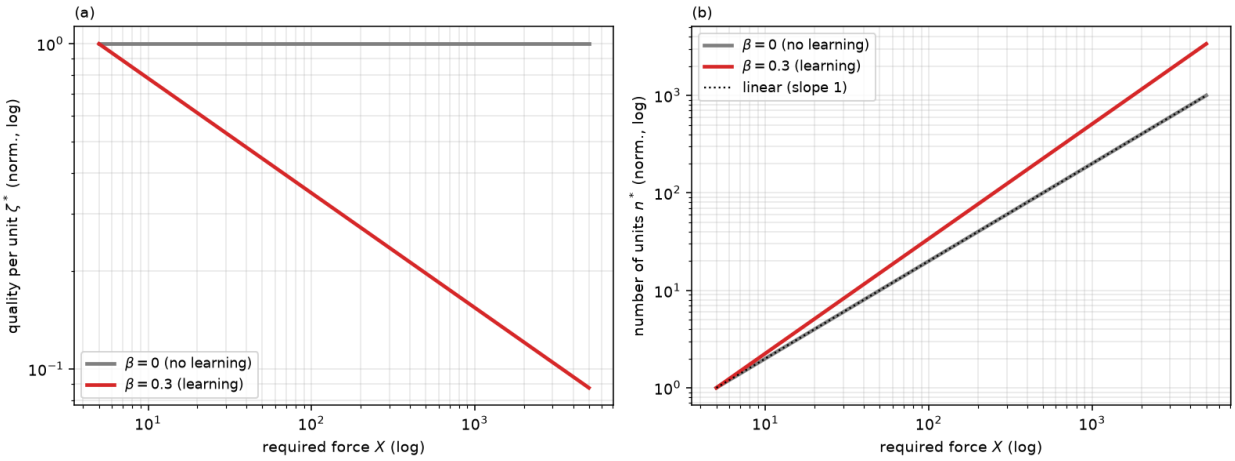


Figure 2: The baseline quantity–quality arsenal. (a) Without learning quality per unit is scale-invariant; with learning it falls as the required force grows. (b) Fleet size grows super-linearly with force (above the slope-1 line): larger arsenals tip to more, lower-quality units.

learning) this cost is linear, which makes the win probability a clean "productivity share" and anchors the cost-exchange calibration; learning is reintroduced below on the mass input.

Lemma 1 (Unit cost of force)**.** *The minimum cost of producing one unit of effective force* (5) *is*

$$P_i = \left(\pi_m^{1-\sigma} + \pi_q^{1-\sigma} \right)^{1/(1-\sigma)}, \qquad \psi_i \equiv 1/P_i, \tag{7}$$

and the cost share of mass is $s_m = \pi_m^{1-\sigma}/(\pi_m^{1-\sigma} + \pi_q^{1-\sigma})$*. A budget* B_i *buys* $X_i = B_i/P_i = \psi_i B_i$ *of effective force.*

Thus ψ_i is the "productivity of spending", the effective force a dollar buys. The cost-exchange ratio of Figure 1 is, in this notation, π_m relative to the value of the target: the object the paradigm collapses.

4.4 Contest layer

The contest layer is common to both versions of the arsenal. Given forces $X_i = \psi_i B_i$, the probability i prevails is a Tullock CSF with decisiveness $R > 0$:

$$p_i = \frac{X_i^R}{X_i^R + X_j^R} = \frac{(\psi_i B_i)^R}{(\psi_i B_i)^R + (\psi_j B_j)^R}. \tag{8}$$

Lemma 2 (Contest equilibrium)**.** *For* R *in the range admitting an interior pure-strategy equilibrium* ($R \leq 2$ *in the symmetric case; see Cornes and Hartley, 2005 for general conditions):*

(i) equilibrium effective efforts satisfy $X_A/X_B = \psi_A/\psi_B$*, so the win probability is the* productivity share

$$p_A = \frac{\psi_A^R}{\psi_A^R + \psi_B^R}; \tag{9}$$

(ii) equilibrium spending is equal across sides, $B_A = B_B$*; for the lottery CSF* $R = 1$*,* $B_A = B_B = P_A P_B V/(P_A + P_B)^2$*;*

(iii) in the symmetric case ($\psi_A = \psi_B$), $B^* = RV/4$, independent of the productivity level.

Lemma 2 is the workhorse: everything about the paradigm runs through ψ_i, i.e. through the industrial layer P_i. The sections that follow feed increasingly rich industrial layers through (9), all in the two-good form; the same productivity-share result holds in the multiplicative baseline, where productivity is again the inverse of a cost parameter, now that of the quantity–quality cost function (Appendix B states it).

What is standard and what is new. Lemma 2 is deliberately standard: with linear force costs it restates the Tullock contest with heterogeneous marginal costs. The existence bound invoked in Proposition 5 is classical too — but what that literature treats as a primitive, the decisiveness of the contest, is here *derived*, as $\gamma = 1/(1-\beta)$, from a learning rate that is in principle measurable. Borrowing the bound is not the same as borrowing the result. That is the contribution: an industrial layer mapping measurable primitives (unit costs, the learning elasticity β, the quality elasticity b, cost-exchange ratios) into the productivity ψ and into the contest's *effective* decisiveness. Results with no counterpart in the contest literature are the arsenal-design exponents of Proposition 1; the threshold shift $\beta = 1/(2-b)$ (Appendix B); the deterrence value of an industrial lead (Proposition 6); the learning-race characterization of the offense–defense balance (Proposition 7); the operational conditions under which production experience accrues at all (Proposition 8); the scale- and stake-dependence of the arsenal mix (Propositions 9–11); and the two-period race with an explicit experience state (Propositions 12–13), whose antecedents lie in the industrial-organization literature on learning-by-doing pricing (Spence, 1981; Cabral and Riordan, 1994) rather than in contest theory.

5 The paradigm shift with exogenous cost

Hold the industrial layer exogenous and ask what a fall in the price of mass does.

Proposition 2 (The paradigm shift favors the adopter)**.** *A fall in π_m (cheaper drones $c_m\downarrow$ and/or better accuracy $e_m\uparrow$) lowers P_i, raises ψ_i, and raises p_i. If $\sigma > 1$ it also raises the mass cost share s_m (the side re-tools toward mass); if $\sigma < 1$, s_m moves the other way.*

An asymmetric adopter (one side reaches the cheap-mass frontier first, the other constrained by organizational capital, Horowitz, 2010; Gilli and Gilli, 2016) gains win probability and shifts toward mass (Figure 3a). *Early adoption and absorptive capacity.* Because adoption is gated by organizational capital rather than by price (Horowitz, 2010), the binding policy lever is doctrine, training, and procurement agility, not the drone bill; the early adopter's edge is a first-mover advantage of the kind Proposition 6 makes precise, which raises the return to moving early. The effect on *budgets*, however, is subtle, and the symmetric case carries the paradigm's macro-signature.

Proposition 3 (An arms race in capability, not spending)**.** *If the price of mass falls for both sides, then $P_A = P_B$ fall together: the win probability stays at $1/2$ and equilibrium budgets are unchanged at $B^* = RV/4$, while effective force per dollar $\psi = 1/P$ rises without bound.*

Proposition 3 cautions against reading the drone era off defense budgets. Cheap mass need not raise military spending; it raises the destruction per unit of spending (Figure 3b). The transformation in Ukraine is measured in units produced, not in a proportional budget explosion.

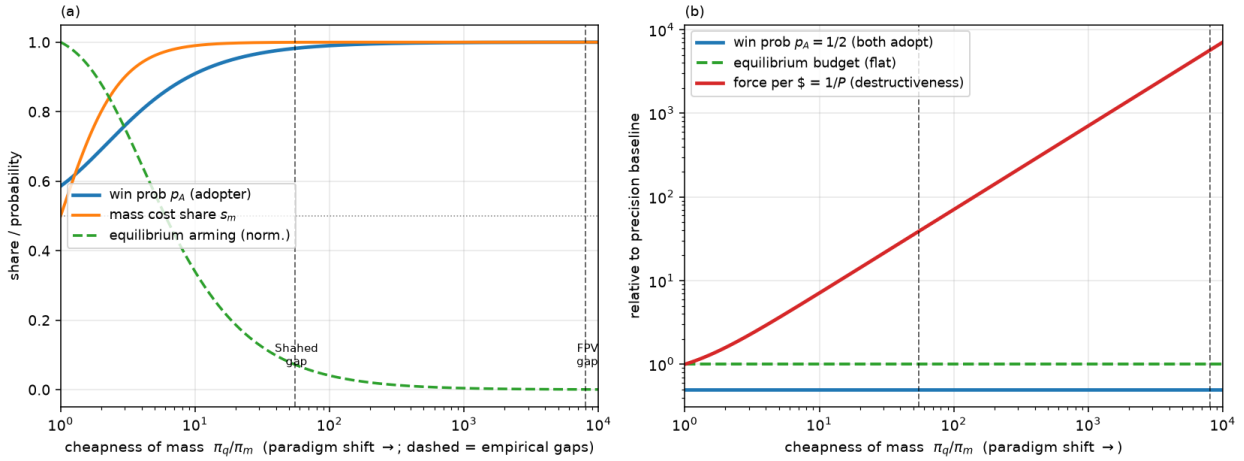


Figure 3: The paradigm shift as a fall in the price of effective force. (a) Asymmetric adoption: the adopter's win probability and mass share rise; equilibrium arming falls as asymmetry grows. (b) Symmetric adoption: win probability and budgets are flat while destructiveness per dollar ($1/P$) soars. Dashed vertical lines mark the calibrated empirical cost gaps π_q/π_m from Table 1 (Shahed vs. precision ≈ 55; FPV vs. precision $\approx 8{,}000$): at observed gaps asymmetric adoption is essentially decisive, and symmetric adoption raises destructiveness per dollar by one-to-four orders of magnitude at unchanged budgets.

6 Learning and increasing returns

The paradigm shift of Section 5 treated costs as fixed. I now let the cheap-mass good *learn*: in the two-good form, the unit cost of drones falls with cumulative production (Wright, 1936; Arrow, 1962; Thompson, 2012),

$$c_m(m) = c_{m0}\, m^{-\beta}, \qquad \beta \in (0,1), \tag{10}$$

so spending B_m on mass buys $m = (B_m/c_{m0})^{1/(1-\beta)}$ drones.[5] (The multiplicative baseline of Section 4.2 builds the same Wright learning into its cost from the outset; the two formulations coincide in the limits of Appendix B.)

Proposition 4 (Learning implies increasing returns to mass)**.** *Effective force from mass spending is $e_m m = \theta\, B_m^{\gamma}$, with elasticity*

$$\gamma = \frac{1}{1-\beta} > 1, \qquad \theta = e_m\, c_{m0}^{-1/(1-\beta)}. \tag{11}$$

It is strictly convex in B_m: doubling the mass budget more than doubles mass force. The cheap-mass technology exhibits increasing returns; the constant-cost precision technology does not.

Whether a force actually travels down this curve is a separate question, since cumulative output must grow to move along it, and Section 7 shows that growth is governed by conditions that are operational rather than technological (Proposition 8). Until then, treat β as a technological constant.

[5] Here $c_m(m)$ is the *average* unit cost of a run of m, so the run costs $c_{m0}m^{1-\beta}$. Under the marginal-cost reading the run costs $c_{m0}m^{1-\beta}/(1-\beta)$, identical up to a constant absorbed into c_{m0}; the convention matters only for calibrating the level of c_{m0}. With inherited experience $Q_0 > 0$ the run cost is $c_{m0}[(Q_0+m)^{1-\beta} - Q_0^{1-\beta}]$, which removes the singularity at zero and converges to the text's form for $m \gg Q_0$, the relevant regime for wartime mass production.

Remark 3 (Static learning and spillovers). Two caveats discipline the learning specification. First, (10) collapses a dynamic learning curve into a static cost function, the reduced form introduced by Spence (1981) (also Fudenberg and Tirole, 1983): it is appropriate when the production run maps to a single campaign, so that cumulative output and fielded quantity coincide; Section 9 supplies the explicit state variable and shows which dynamic claims survive it. Second, (10) makes learning proprietary to own output. Let instead each side's experience pool be $m_i + \lambda m_j$, with $\lambda \in [0,1]$ the share of the rival's output that spills over (copied designs, common commercial supply chains), so a run costs $B_i = c_{m0}\, m_i\, (m_i + \lambda m_j)^{-\beta}$. At a symmetric point the own-spending force elasticity is then $(1+\lambda)/(1+\lambda-\beta)$, and the decisive-race threshold (for $R = 1$) rises from $\beta = 1/2$ to $\beta = (1+\lambda)/2$: with full spillovers ($\lambda = 1$) no learning rate below one produces a decisive race.Spillovers shrink the decisive-race region, so *closing* an innovation ecosystem (secrecy, export controls, proprietary feedback loops) is an investment in lowering λ, a testable attenuation of every learning result below.

Increasing returns turn production scale into a strategic weapon. Restricting attention to the mass technology (the regime the paradigm pushes toward) and feeding Proposition 4 into the contest gives the paper's sharpest result.

Proposition 5 (Effective decisiveness and winner-take-all)**.** *With both sides on the cheap-mass technology and $R = 1$, the arming game is a Tullock contest with effective decisiveness $\gamma = 1/(1-\beta)$,*

$$p_i = \frac{B_i^{\gamma}}{B_i^{\gamma} + B_j^{\gamma}}.$$

Hence: (i) for $\beta \le 1/2$ (so $\gamma \le 2$) a symmetric interior equilibrium exists, with spending $B^* = \gamma V/4$ rising in the learning rate; (ii) for $\beta > 1/2$ (so $\gamma > 2$) no pure-strategy equilibrium exists: the arming game enters a decisive-race *region in which simultaneous play is randomized escalation (Section 6.1), production leads buy more-than-proportional force advantages, and the literal winner-take-all outcome obtains under sequential production commitment (Proposition 6).* (*For general R the effective decisiveness is $R\gamma$ and the threshold is $R\gamma = 2$.*)

Proposition 5 is the formal content of the idea that when the cheaper side can also out-produce the other, static cost advantages compound dynamically.[6] Learning raises the effective decisiveness of the contest; strong enough learning ($\beta > 1/2$) makes mass production a race one side wins outright (Figure 4). Iteration speed and the use–produce–develop loop acquire strategic value precisely because they are investments in β and θ. The two-player Tullock existence threshold at decisiveness 2 is standard (Pérez-Castrillo and Verdier, 1992; Baye et al., 1994). The same force operates in the multiplicative baseline, where the threshold generalizes to $\beta = 1/(2-b)$: the "good-enough" parameter raises the learning needed for a decisive race (Appendix B).

6.1 Robustness of the winner-take-all threshold

The mechanism extends beyond the Tullock ratio form in three respects. The threshold's precise *location* is form-specific; the increasing-returns mechanism is general.

The driver is increasing returns, not the ratio form. Write the arming payoff in budgets, $U_i(B_i) = p_i(X_i, X_j)\, V - B_i$ with $X_i = \theta B_i^{\gamma}$. Learning makes force convex in the budget ($\gamma > 1$), and in any contest whose success probability is a function of relative impact, the composite impact function inherits the exponent $R\gamma$: the learning elasticity multiplies whatever decisiveness the contest technology supplies. Existence results for general contest technologies tie the pure-strategy

[6] O'Brien (2024) reaches the same conclusion informally, arguing that a small edge in autonomous-system effectiveness, particularly in drone survivability, would compound into a large one.

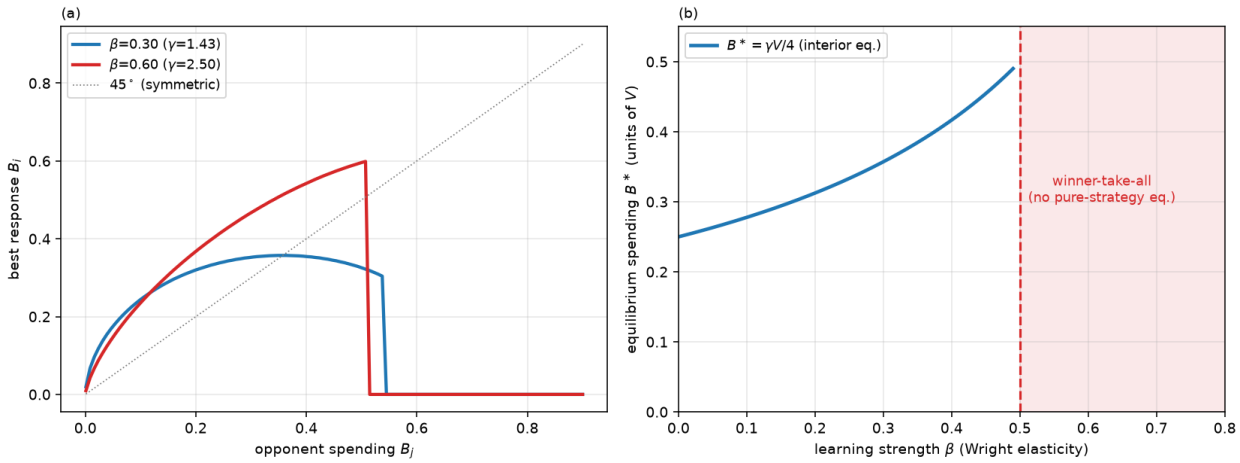


Figure 4: Learning turns mass production strategic. (a) Best responses: for $\beta < 1/2$ an interior symmetric equilibrium exists; for $\beta > 1/2$ the best response jumps (winner-take-all). (b) Equilibrium arming $B^* = \gamma V/4$ rises with the learning rate up to $\beta = 1/2$, beyond which no pure-strategy equilibrium exists.

equilibrium to exactly this curvature (Cornes and Hartley, 2005), so a large enough γ defeats existence for any given contest form; a sufficiently flat contest (low R) postpones the threshold but cannot remove the mechanism, since the bound $R\gamma = 2$ is reached at $\beta = 1 - R/2$ for every $R < 2$. The decisive-race region is therefore a property of *learning*, read off, conveniently, through the Tullock bound.

The value is form- and n-specific. The number 2 is the two-player Tullock existence bound; the n-player symmetric bound is $R \leq n/(n-1)$, so more contestants make the decisive region easier to reach (three sides, $R \leq 3/2$; many sides, $R \leq 1$). A difference-form CSF (Hirshleifer, 1989) relocates rather than removes the effect and forfeits the productivity-share identity (9) that organizes the paper; the all-pay auction (Baye et al., 1996) is the $R \to \infty$ limit of the same force; and effort or capacity caps restore pure-strategy existence at a boundary (Che and Gale, 1998), economically natural here, since production capacity is bounded. The content of Proposition 5 is thus "a threshold exists," not the specific value $\beta = 1/2$.

Above the threshold the equilibrium is mixed. The symmetric mixed-strategy equilibrium features randomized escalation with near-complete rent dissipation (Baye et al., 1994). That is the economically appropriate description of the region, an arms race with no stable interior rest point, in which small production leads are worth pre-empting, and it reinforces rather than weakens the commitment channel of Section 11. Throughout, "winner-take-all" names the threshold; strictly, it marks where stable interior arming ends and the value of moving first becomes discontinuous. The sequential-move (Stackelberg) case, in which one side commits its production first, restores a pure-strategy equilibrium through first-mover advantage, and there the label is literal; it is developed next. The dynamic version, with an explicit experience stock and both sides learning over time, is developed in Section 9, which solves the two-period race; the infinite-horizon version, and the joint evolution of offensive and defensive experience, remain open (Section 12).

6.2 Sequential moves and the value of an industrial lead

The winner-take-all region need not be read as indeterminacy. It is the region in which the *timing* of production becomes decisive: a side able to commit its industrial base first turns the instability

to its advantage.

Proposition 6 (Sequential moves and industrial lock-in)**.** *Let the two sides move sequentially in the pure-mass arming game of Proposition 5 (effective decisiveness $\gamma = 1/(1-\beta)$): a* leader *can concede, ending the game at payoff zero, or commit production spending $B_L \in [\underline{B}, V]$, where $0 < \underline{B} < V/2$ is an arbitrarily small minimum positive commitment. The* follower *observes B_L before choosing $B_F \in [0, V]$. Then* (i) *a subgame-perfect equilibrium in pure strategies exists for every $\beta \in (0,1)$, including the winner-take-all region $\beta > 1/2$ where the simultaneous-move game has none; and* (ii) *in that region the equilibrium is* deterrence*: the leader commits*

$$\bar{B} = \frac{V}{\gamma}(\gamma - 1)^{(\gamma-1)/\gamma},$$

the follower drops out, and the leader wins V with certainty for a net payoff $V - \bar{B} > 0$, against zero for both sides in the simultaneous-move mixed-strategy equilibrium. The first-mover rent $V - \bar{B}$ is strictly positive and declines toward zero as $\beta \to 1$ (the all-pay limit). (At $\bar{B}$ the follower is indifferent between dropping out and an entry that also nets zero; the equilibrium adopts the standard leader-preferred resolution of this indifference. With a minimal production increment ε, any feasible overcommitment $\bar{B} + \varepsilon$ makes deterrence strict and converges to the stated payoff as $\varepsilon \downarrow 0$.)

The value of an industrial lead. Proposition 6 turns the winner-take-all region into a determinate and policy-relevant statement. Where the simultaneous race has no stable outcome, the ability to commit production *first* is decisive: the leader locks in the win by fielding a scale the follower cannot profitably contest. The strategic asset is the credibility of an early, hard-to-reverse production commitment, an *industrial lead*. This is the formal counterpart of the idea that in a fast-learning, winner-take-all technology, being first to scale is worth more than being efficient at the margin. It sharpens the commitment channel of Section 11: the side that can lock in an industrial lead has every incentive to move first, and the side that cannot may prefer preventive action to a settlement it expects to be foreclosed. And it reframes procurement and industrial policy (Section 12): the return to pre-positioned production capacity, and to the speed of the use–produce–develop loop that builds a lead, is a first-mover rent, largest where learning is strong enough to make the race winner-take-all but before it dissipates in the all-pay limit. The construction follows the sequential-contest literature (Dixit, 1987; Leininger, 1993; Morgan, 2003).

6.3 Asymmetric learning rates: a numerical illustration

The motivating conflict is asymmetric, and the paradigm's sharpest question is what a *learning-rate edge* is worth. Figure 5 computes the Nash equilibrium of the *pure-mass* arming game — the game of Proposition 5, in which both sides field only the cheap-mass good, so the arsenal-mix margin is switched off and learning operates on quantity alone — when side A's learning elasticity varies against a rival fixed at the benchmark $\beta_B = 0.30$, at low and high stakes (measured by V/c_{m0}, the stake relative to the initial unit cost). Three results stand out. First, the learning edge is *stake-dependent*: the same gap in β buys a far larger shift in win probability in a high-stakes conflict, because higher stakes push production deeper into the learning curve, the equilibrium counterpart of Proposition 10. Second, equilibrium arming is hump-shaped in the gap: spending peaks when the sides are matched and falls as the laggard is discouraged, so an emerging learning asymmetry can show up as *falling* budgets alongside a shifting balance, reinforcing the measurement caution of Proposition 3. Third, at high stakes the pure-strategy equilibrium *fails* at a modest gap ($\beta_A \gtrsim 0.39$ against 0.30, both well below the symmetric threshold 1/2): a global grid scan finds no mutual best

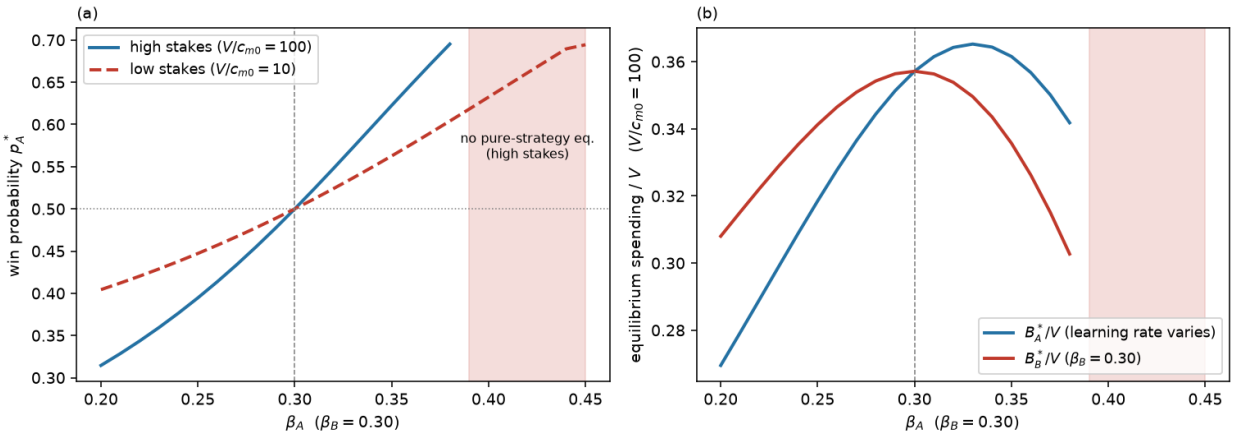


Figure 5: Asymmetric learning rates in the pure-mass arming game, in which both sides field only the cheap-mass good. The horizontal axis is side A's learning elasticity β_A, swept over $[0.20, 0.45]$; side B's is held fixed at the benchmark $\beta_B = 0.30$, so A is the slower learner to the left of 0.30 and the faster to its right ($R = 1$ throughout). (a) A's win probability: the same learning edge is worth more at high stakes; in the shaded region a global grid scan finds no pure-strategy equilibrium in the high-stakes game (the asymmetric decisive-race region, reached with both learning rates below 1/2). (b) Equilibrium spending (high stakes): arming peaks when the sides are matched and the faster learner outspends the laggard beyond parity.

response, because the laggard's payoff develops a second local optimum (drop out versus match) between which its best response jumps. Asymmetry brings the decisive-race region down to learning rates that symmetric analysis would call safe, and it does so exactly where the stakes are highest. The illustration is confined to the pure-mass game; the asymmetric two-stage equilibrium in both weapons and both layers is not solved here.

6.4 Where the learning rate comes from

The paper treats the learning rate β as a technological constant. The war suggests it is a *choice* about how learning is organized. Ukraine's Brave1 ecosystem fields many competing suppliers running parallel design variants with fast, real-world feedback (a use–produce–develop loop; Section 2, KSE Institute, 2024), while the adversary concentrates production in a few variants improved incrementally. In the language of the economics of innovation this is a breadth–depth trade-off: concentration drives each line far down its own experience curve (high cumulative output per variant) but explores little of the design space, whereas distributed experimentation accepts lower volume per variant in exchange for a better best-of-many frontier and recombination across designs (Nelson, 1961; Weitzman, 1979, 1998). It is the military instance of Sah and Stiglitz's (1986) contrast between *polyarchy* (many independent trials, so more good designs survive) and *hierarchy* (centralized selection), and of absorptive capacity as an investment (Cohen and Levinthal, 1990; Nelson and Winter, 1982). The formal state records cumulative production. Combat feedback, test cadence, supplier breadth, and design recombination are candidate determinants of the reduced-form learning parameters; empirically, they may shift β, the cost level c_{m0}, productivity θ, or effectiveness e_m. Read this way the deep strategic asset is the *organization that learns*: the side that builds the better experimentation ecosystem can raise the effective rate of improvement, and with it the increasing returns of Proposition 4, a path to the winner-take-all threshold of Proposition 5, and

durability in the learning race of Section 7 below. It reframes the offense–defense balance and the industrial lead of Proposition 6 as, at bottom, contests between learning organizations.

Remark 4 (Endogenous learning rate). Let each side choose its learning elasticity β_i in a prior stage at a convex organizational cost $\Phi(\beta_i)$ ($\Phi' > 0$, $\Phi'' > 0$), the cost of running a broader experimentation ecosystem. Let $\Pi_i^*(\beta_i, \beta_j; V)$ denote its equilibrium continuation payoff after arsenal choice and arming costs. An interior organizational choice satisfies

$$\Phi'(\beta_i^*) = \frac{\partial \Pi_i^*(\beta_i, \beta_j; V)}{\partial \beta_i}.$$

If the continuation game has increasing differences, $\partial^2 \Pi_i^* / (\partial \beta_i \partial V) > 0$, the chosen β_i^* rises with the stake: larger conflicts call for greater investment in the rate of learning, an arms race in the *organization* of learning rather than in the *stock* of weapons. Because β governs both the winner-take-all threshold (Proposition 5) and the durability of an offense advantage (Section 7), each becomes an organizational choice: a side may invest to cross the threshold or to out-learn its rival, and the offense–defense balance is the equilibrium of two chosen learning rates, not a fixed technological fact. The mapping from organizational breadth to the reduced-form parameters can be microfounded through a parallel-experimentation-and-selection model (Nelson, 1961; Sah and Stiglitz, 1986; Weitzman, 1979).

7 Offense–defense and the arms race

Countermeasures enter through effectiveness: the opponent's counter-drone effort (EW, interceptors) lowers each drone's effectiveness, $e_m = e_{m0}\,\varphi(d_j)$ with $\varphi' < 0$. Consider the cost-exchange condition and its dynamics. Suppose defeating one drone costs the defender c_d/κ (interceptor price c_d divided by kill probability κ), while fielding one costs the attacker c_m. Offense wins the marginal exchange (mass is the cost-effective way to impose cost) iff the *offense advantage*

$$A_{\text{off}} \equiv \frac{c_d/\kappa}{c_m} > 1. \tag{12}$$

The calibration of Table 1 makes this condition concrete and reveals its knife-edge dependence on the defense technology (Figure 1b).[7] Against an exquisite interceptor, offense is overwhelmingly favored: a \$59 k Shahed met by a \$1–3.25 m interceptor gives $A_{\text{off}} \geq 17$–55 at $\kappa = 1$. Figure 1c reports the separate target-exposure multiple T/c_m for FPVs against ground targets. The *same* Shahed met by a \$2,100 panel-median Sting interceptor drone inverts the interception exchange: at $\kappa = 1$ the ratio is about 1:28, and at the reported 0.68 kill probability $A_{\text{off}} = (c_d/\kappa)/c_m \approx 0.052$, i.e. about 1:19, *defense*-favored either way. The cheap-defense response, not exquisite air defense, is what restores favorable economics. But costs are endogenous, and both sides learn.

Remark 5 (Target value and the exchange-ratio fallacy). The offense advantage (12) compares flow costs and is silent about what the defense protects. With a protected asset of value T that a leaker destroys with probability π_T, firing the interceptor is statically rational iff $c_d/\kappa < \pi_T T$: defending a \$500 m asset with a \$3 m interceptor against a \$59 k drone passes this test easily. The two objects answer different questions. Target value governs the *level* decision, whether to intercept a given

[7] A reporting convention: kill probabilities are observed for only some of the defenses in the panel. Where κ is not separately measured I set $\kappa = 1$, so the reported figure is the raw cost ratio and therefore a *lower bound* on the defender's cost per kill — the true A_{off} is weakly larger, and more offense-favorable, than every number below and in Figure 1b. Only the cheap interceptor drone, for which a success rate is reported, is quoted both ways.

shot; the exchange ratio governs the *margin*, who can sustain the race, because the attacker chooses how many threats to present and, at drone prices, can saturate any finite budget of individually rational interceptions (Remark 6). Exquisite point defense of high-value assets and a losing cost-exchange position therefore coexist; (12) is the arms-race object, not an operational firing rule.[8]

Proposition 7 (Durability of cheap mass is a learning-rate race)**.** *Let offensive and defensive unit costs fall at Wright rates β_m, β_d with their own cumulative outputs Q_m, Q_d, and let the kill probability evolve with defensive experience, so*

$$A_{\text{off}} \;=\; A_0\,\frac{Q_m^{\beta_m}}{Q_d^{\beta_d}}\cdot\frac{1}{\kappa(Q_d)}, \qquad \frac{d\ln A_{\text{off}}}{dt} \;=\; \beta_m g_m - \beta_d g_d - \frac{d\ln\kappa}{dt},$$

with g_m, g_d the growth rates of cumulative output. Cheap-mass offense is durable (A_{off} grows) iff the attacker's learning-weighted production growth exceeds the defender's plus the improvement in kill rates, $\beta_m g_m > \beta_d g_d + d\ln\kappa/dt$. On the equal-growth, constant-κ path this reduces to $A_{\text{off}}(Q) \propto Q^{\beta_m-\beta_d}$: durability iff the attacker out-learns ($\beta_m > \beta_d$), self-limitation (defense catches up and precision-like economics are restored) iff the defender out-learns.

The offense–defense balance is the outcome of a race between two learning curves (Figure 6). This is the economic content behind the offense–defense *scaling* of Garfinkel and Dafoe (2019), specialized to cost learning, and it answers the central open question (is cheap mass durable?) with a condition rather than an assertion.

Two readings of the condition should be kept apart. The equal-growth reduction ($A_{\text{off}} \propto Q^{\beta_m-\beta_d}$, durability iff $\beta_m > \beta_d$) is analytically convenient. Empirically, durability compares both the learning elasticities and movement along the two experience curves. Published flows illustrate the contrast: PAC-3 MSE output stood at roughly 600 per year in 2026, while Ukrainian FPV output rose from about 20,000 to about 200,000 per month between early 2024 and 2025 (Cancian and Park, 2026) (Section 2). Translating those flows into $g_i = \dot{Q}_i/Q_i$ requires the corresponding cumulative-output bases. Over any observed horizon the sufficient statistic for realized cost decline is $\beta_i \ln(Q_{iT}/Q_{i0})$: the learning elasticity times the distance traveled down the curve. Remark 6 derives the production flows that move Q_i. A full differential-game treatment (Brito and Intriligator, 1995) is left to future work.

Counter-drone procurement. Proposition 7 is a procurement rule for defense as much as for offense: since the balance is a race between learning curves, the way to hold or restore favorable economics is to field the *faster-learning* and cheaper defense (interceptor drones, electronic warfare, onboard autonomy), not the more exquisite one. Meeting a \$50 k attack drone with a multi-million-dollar interceptor loses the cost-exchange even when it wins the engagement; the cheap-defense response wins it. The strategic variable is the defensive learning rate β_d, which argues for investment in iteration speed and mass-producible countermeasures over a small inventory of exquisite interceptors.

Autonomy can favor either side. On the attacking side it raises e_m and can make survival less sensitive to countermeasure intensity, lowering the elasticity ϵ of Proposition 8(v) and with it the

[8] The 2026 Iran war shows how that coexistence resolves. Air and missile defenses protected bases and infrastructure whose value made each interception individually rational — air-to-air missiles costing on the order of \$1 m were expended against drones costing tens of thousands — and largely succeeded engagement by engagement. What they could not sustain was the magazine: interceptor stocks were drawn down far enough that the defender's own analysts judged the remaining inventory a constraint on future operations (Cancian and Park, 2026). A high-value asset can thus become harder to defend without a single engagement having been lost. The exchange ratio bites through the stock rather than through any one intercept, which is why the response was to field a *cheaper* interceptor rather than a better one — the procurement corollary of Proposition 7 below.

subsidy that jamming pays to the attacker's learning. On the defending side it does the same work for the interceptor, raising the kill probability κ or lowering c_d, and the first fielded uses have included autonomous interception of Shahed-type targets. Its net effect depends on how strongly it moves attack effectiveness, defensive kill probability, and the two cost curves. Equation (12) governs the *interception exchange*, in which only the interceptor price, its kill probability and the attacker's unit cost appear; $\pi_m = c_m/e_m$ governs *force production*, where the attacker's effectiveness enters and (12) does not. Autonomy moves both, by different channels, and it is the durability condition of Proposition 7 that nets them against each other. The corollary is being acted on. Facing sustained one-way-attack salvos in 2026 and falling interceptor inventories, US forces turned toward Ukrainian-designed kinetic interceptor drones costing thousands rather than millions per round. These are the same class of system, and close to the same price, as the cheap interceptor that inverts A_{off} in Table 1. The move is the cheap-defense response of (12) rather than a more exquisite interceptor, and it illustrates the rule without testing it.

Remark 6 (Attrition and saturation). The two features that most define drones, expendability and the saturation of finite defenses, enter through the same effectiveness term e_m and can be treated statically, before any full dynamics. *Attrition* is a survival wedge on the price of force. Since e_m is delivered effect per procured unit (6), it already nets out losses; write $e_m = \bar{e}_m\, s$, separating the effect a unit achieves *conditional* on arriving, $\bar{e}_m$, from the share $s = \varphi(d_j)$ that arrives. The delivered-effect price is then $\pi_m = c_m/e_m = c_m/(\bar{e}_m s)$, and countermeasures raise it by lowering s with $\bar{e}_m$ unchanged. In a stationary environment the standing force is a stock–flow steady state (production rate over loss rate μ). High attrition therefore favors cheap, replaceable mass and penalizes exquisite precision (each loss is expensive), *widening* the cheap-mass advantage exactly in the drone-war environment. *Saturation* makes effectiveness increasing in own mass: with defender magazine depth $\bar{M}_d$ and leak-through $L(m/\bar{M}_d)$ rising in salvo size, delivered force is $\propto m\, L(m)$. Writing $e_m(m) \propto m^{\alpha_e}$ with saturation elasticity $\alpha_e > 0$ gives delivered force $\propto m^{1+\alpha_e}$, so with learning ($m \propto B^{1/(1-\beta)}$) the effective decisiveness is $(1+\alpha_e)/(1-\beta)$ and the winner-take-all threshold falls to $\beta = (1-\alpha_e)/2 < \frac{1}{2}$. Saturation and learning are thus complementary increasing-returns channels, so the threshold of Proposition 5 is over-determined rather than fragile. The derivation nowhere uses *why* effectiveness rises with own quantity. If it rises because cumulative production improves guidance rather than because a finite defense is saturated, then α_e is an effectiveness-learning elasticity and the same bound applies: learning taken as capability rather than as price lowers the winner-take-all threshold exactly as saturation does, and the two compound. The reading holds in the static reduced form used here, where cumulative output and fielded quantity coincide (Remark 3); in a dynamic treatment performance learning would need its own state, alongside the cost experience of Section 9. The arithmetic of $L(\cdot)$ is unforgiving for the defender: against salvos of forty airframes, an interception rate of 0.9 still admits four arrivals per wave, and a launch rate in the low hundreds per week admits twenty or more arrivals weekly from a defense performing well by any historical standard. Leakage is a share of a quantity the attacker chooses. What requires the full dynamic game is the *adjustment path*: when finite magazines are exhausted, how quickly they can be reconstituted, and how much should have been stockpiled before a war. The steady-state economics do not require it. The interceptor drawdowns of 2026 are of exactly this kind. What they show is a finite magazine emptying faster than industry can refill it: a question about timing, and about how large a stockpile should have been held before the war began. The model here answers a different question: who can sustain the exchange once both sides are producing, and at what cost per unit of effect delivered. The episode is therefore evidence about the mechanism this paper sets aside rather than the one it analyzes.

Attrition also enters a second time, dynamically, as the generator of experience. For a reusable

platform whose missions do not deplete the inventory, let F be the standing force and μ the non-mission loss rate. The inventory identity is then $x = \dot{F} + \mu F$: production covers net accumulation plus replacement. This is the force-stock special case of the more general accounting below. Since production is what accumulates experience, $\dot{Q} = x$ gives

$$g \;\equiv\; \frac{\dot{Q}}{Q} \;=\; \frac{\dot{F} + \mu F}{Q}, \qquad \text{so} \qquad \Delta \ln c \;=\; -\beta \ln\!\big(Q_T/Q_0\big),$$

and the growth rates entering Proposition 7 are not free parameters but the sum of a build-up term and a replacement term. Realized cost decline is a learning rate times a number of doublings, and the two terms are separate sources of doublings. Under a *stationary* force requirement $\dot{F} = 0$ the replacement term is all that remains, $g = \mu F/Q$, and durability reads $\beta_m \mu_m (F_m/Q_m) > \beta_d \mu_d (F_d/Q_d) + d\ln\kappa/dt$.

Two implications follow. A weapon that is stockpiled rather than consumed learns only while its inventory is being built: once the target stock is reached and $\mu \approx 0$, both terms vanish and no further experience accrues whatever its β. And learning decelerates mechanically as F/Q falls, so Wright's diminishing returns arrive through the stock rather than through the technology. For a consumable weapon, mission depletion is an additional source of production and must be accounted for separately rather than folded into μ.

Two conditions qualify this, and both are substantive. First, high attrition raises g only where losses are *replaced*, and replacement beats withdrawal only where the unit is cheap: the operative primitive is economic replaceability, not physical expendability. Second, and less obviously, the force level must be pinned by an operational requirement rather than chosen freely. Under free choice force is a depreciating asset, so a higher loss rate *lowers* the return to building it and production falls in μ; the replacement channel operates only where holding a line, sustaining a strike tempo or defending an airspace makes the force requirement binding. Proposition 8 states the conditional result exactly; the free-choice case is discussed after it.

7.1 When does experience actually accumulate?

The right primitive is the required *delivered effect*, not the required force. A consumable is not held because a standing stock is intrinsically valuable but because a flow of delivered effects must be generated; the stock is a buffer, not the target. Let $D_t = s_t(z_t)\, y_t$, where y_t is the flow of missions and $s_t(z_t) \in (0, 1]$ is the share that survives the defender's countermeasure intensity z_t. This index combines interception, electronic warfare, and other defenses; the interceptor kill probability κ in (12) is one component. Survival is decreasing in z_t, so a more capable defense raises the number of missions needed to deliver a given effect. Let $\xi \in [0, 1]$ be the inventory-depletion rate per mission: $\xi = 1$ for launches, rounds, and one-way sorties, and $\xi = 0$ for a reusable platform that returns from a sortie. Non-mission attrition is $\ell_t = \mu_t F_t$. Thus an intercepted one-way drone is counted once, through the survival term s_t, not again as ℓ_t.

Lemma 3 (Operational constraints induce reconstitution)**.** *Fix F_t, Q_t, an exogenous non-mission loss ℓ_t, and a strictly increasing production cost $C(Q_t, x_t)$. A side minimizes C subject to*

$$s_t(z_t) y_t \geq \bar{D}_t, \qquad F_{t+1} = F_t + x_t - \xi y_t - \ell_t \geq \bar{F}_{t+1}, \qquad x_t, y_t \geq 0.$$

For the cost-minimizing selection, $y_t^* = \bar{D}_t / s_t(z_t)$ and

$$x_t^* = \left[\bar{F}_{t+1} - F_t + \xi \frac{\bar{D}_t}{s_t(z_t)} + \ell_t\right]_+ .$$

In the positive-reconstitution region (the expression in brackets is positive), both operational constraints bind and the inventory path is pinned rather than chosen.

Proof. For $\xi > 0$, any $y_t > \bar{D}_t/s_t(z_t)$ weakly tightens the inventory constraint and strictly raises the least feasible production; for $\xi = 0$, choose the smallest y_t among cost-minimizing ties. Substitution of that minimum mission flow into the inventory constraint gives the smallest feasible x_t, which is the displayed positive part. □

Two decision environments. The distinction behind the proposition is easiest to see from a commander's sequence of decisions. In the first environment, the commander starts with a requirement: a given number of effects must be delivered tomorrow, a line must be held, or a specified number of platforms must be available. If the defense reduces survival, more missions are required; if missions or attrition deplete the force, production fills the resulting gap. In that case production is a residual of the requirement, not a margin on which the commander is optimizing, and every replacement adds to experience. In the second environment, the commander chooses the force or tempo freely. A higher loss rate then makes each unit less durable and can make it optimal to build or field *less*; replacement is not mechanically implied. The same observed attrition can therefore generate learning in the first environment and reduce production in the second. The proposition deliberately describes only the first, pinned-requirement case; it does not assign a sign to free choice.

Proposition 8 (Conditional replacement learning)**.** *Let inventory, experience and cost evolve as*

$$F_{t+1} = F_t + x_t - \xi y_t - \ell_t, \qquad \ell_t = \mu_t F_t, \qquad Q_{t+1} = Q_t + x_t, \qquad c_t = c_0 Q_t^{-\beta},$$

with $Q_0, c_0 > 0$, $\beta \in [0,1)$, $\xi \in [0,1]$, and $\mu_t \in [0,1]$. Suppose the economy is in the positive-reconstitution region of Lemma 3, so $F_t = \bar{F}_t$ and $x_t \geq 0$. Write $\Delta\bar{F}_t \equiv \bar{F}_{t+1} - \bar{F}_t$. Then

(i) production is determined rather than chosen*, $x_t = \Delta\bar{F}_t + \xi y_t + \mu_t \bar{F}_t$, and $g_t \equiv (Q_{t+1} - Q_t)/Q_t = x_t/Q_t$;*

(ii) cumulative experience has three sources,

$$Q_T = Q_0 + \underbrace{\left(\bar{F}_T - \bar{F}_0\right)}_{\text{net build-up}} + \underbrace{\xi \textstyle\sum_t y_t}_{\text{mission depletion}} + \underbrace{\textstyle\sum_t \mu_t \bar{F}_t}_{\text{non-mission attrition}};$$

(iii) $\partial Q_T/\partial y_t = \xi$ and $\partial Q_T/\partial \mu_t = \bar{F}_t$. Thus terminal cost is strictly decreasing in a source exactly when $\beta > 0$ and that source is positive ($\xi > 0$ for mission depletion and $\bar{F}_t > 0$ for attrition). The experience derivatives themselves do not depend on β;

(iv) define the source contribution ratios

$$a_T^y \equiv \frac{\xi \sum_t y_t}{Q_T}, \qquad a_T^\ell \equiv \frac{\sum_t \mu_t \bar{F}_t}{Q_T}.$$

For a common proportional scaling of a positive source,

$$\frac{d \ln c_T}{d \ln y} = -\beta a_T^y, \qquad \frac{d \ln c_T}{d \ln \mu} = -\beta a_T^\ell;$$

(v) if a delivered-effect target binds and a common proportional countermeasure shift has time-specific survival elasticities $\epsilon_t \equiv -\,d\ln s_t/d\ln z_t \geq 0$*, then*

$$\frac{d\ln c_T}{d\ln z} = -\beta \sum_t \frac{\xi y_t}{Q_T}\epsilon_t \;\leq\; 0.$$

The inequality is strict if $\beta > 0$*,* $\xi > 0$*, and at least one positive mission flow has* $\epsilon_t > 0$*. With common* $\epsilon_t = \epsilon$*, this is* $-\beta a_T^y \epsilon$*;*

(vi) if $\xi y_t = \mu_t = \Delta\bar{F}_t = 0$ *for* $t \geq t_0$*, then* $x_t = 0$ *and* $Q_t = Q_{t_0}$ *for all* $t \geq t_0$*: learning ceases entirely, whatever* β*.*

Proof. (i) Substituting $F_t = \bar{F}_t$, $F_{t+1} = \bar{F}_{t+1}$ and $\ell_t = \mu_t \bar{F}_t$ into the inventory law and solving for x_t gives $x_t = \Delta\bar{F}_t + \xi y_t + \mu_t \bar{F}_t$. The expression for g_t is the definition with $Q_{t+1} - Q_t = x_t$.

(ii) Telescoping $Q_T = Q_0 + \sum_t x_t$ and substituting (i); the $\Delta\bar{F}_t$ sum telescopes to $\bar{F}_T - \bar{F}_0$.

(iii) Differentiate (ii) holding $\{\bar{F}_t\}$ fixed. Multiplying either nonzero experience derivative by $\partial c_T/\partial Q_T = -\beta c_0 Q_T^{-\beta-1}$ gives the stated strict cost comparison.

(iv) Scale every mission flow by a common factor. By (ii), Q_T is affine in that factor with slope $\xi \sum_t y_t$; differentiating at unity gives the first elasticity. The attrition case is identical.

(v) A binding target gives $y_t = \bar{D}_t/s_t(z_t)$, hence $d\ln y_t/d\ln z_t = \epsilon_t$. Differentiating (ii) and then $\ln c_T$ gives the displayed weighted sum. The common-elasticity case factors out ϵ.

(vi) Part (i) gives $x_t = 0$, hence $Q_{t+1} = Q_t$. □

Parts (iv)–(v) are the operative statements. The elasticity of terminal marginal cost to a source of depletion is $-\beta$ times that source's contribution ratio. For a continuously consumed and replenished system, a_T^y can approach one; for a system held in inventory it approaches zero. These are contribution ratios rather than literal shares when net inventory build-up is negative.

The proposition is thus an accounting result about *reconstitution*, not a general claim that attrition always accelerates learning. It has three operational readings. For one-way weapons, $\xi = 1$: a binding delivered-effect target pins launches, so every additional mission both depletes the stock and calls forth production. This is consumable, tempo-driven warfare. For reusable platforms, $\xi = 0$: missions leave the stock intact, but a binding readiness or inventory requirement turns non-mission attrition into the source of replacement. Finally, when neither requirement binds, the commander is back in the free-choice environment. The model then makes no directional claim: losses may prompt replacement, retrenchment, or a different operational response. The distinction is operational rather than technological.

Part (v) isolates the countermeasure *learning channel* within the first of those environments. Defenses that reduce survival can lower the attacker's *future marginal unit cost* when they force reconstitution, but they simultaneously reduce current delivered effect and can raise total resource cost. The sign of cost per delivered effect, the offense–defense balance, and the optimal countermeasure intensity require the subsequent dynamic game.

The regimes are therefore not fixed attributes of a weapon, and a system can move between them. A magazine held for deterrence sits in the stockpile corollary, part (vi): with $\xi y_t = \mu_t = \Delta\bar{F}_t = 0$ nothing is produced and no experience accrues, however favorable the technology. The same magazine, once a campaign begins to consume it, enters the positive-reconstitution region and starts generating experience mechanically. US interceptor inventories made that transition in 2026. Before it, the exquisite pole was the stockpile case almost exactly: Tomahawk cumulative procurement moved so little across successive budget cycles that its implied doubling time runs to decades. After it, over a thousand rounds were expended, and the production response (multi-year contracts written explicitly as demand signals, with planned annual output several times the

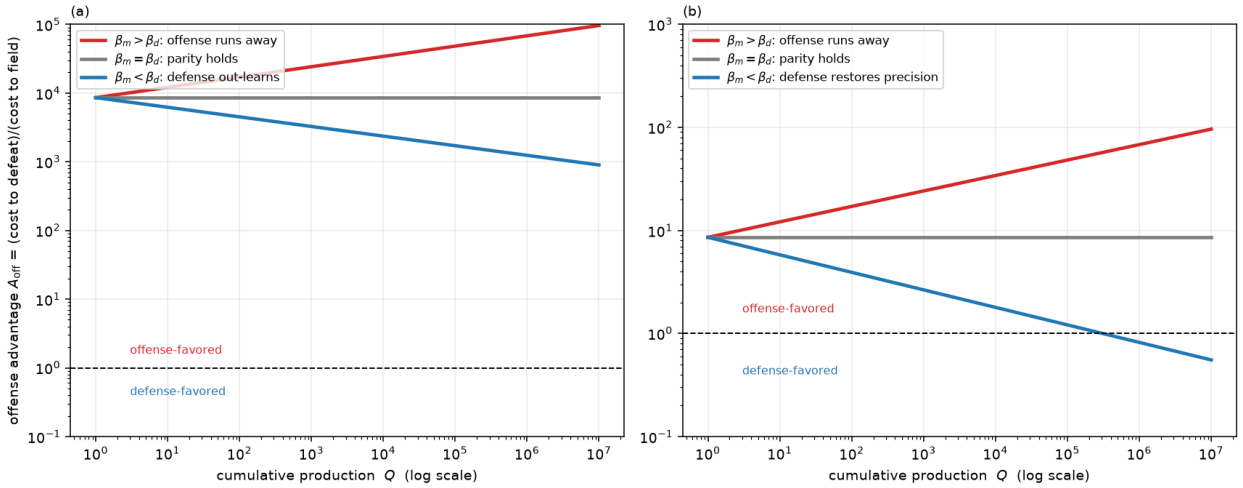


Figure 6: Arms-race stability. (a) With a large starting cost gap, offense stays favored across a wide range even when the defender out-learns. (b) With a modest gap (cheap interceptor drones), a faster-learning defender crosses back below parity, restoring precision-like economics.

prewar rate for both PAC-3 MSE and THAAD; Cancian and Park, 2026) is the reconstitution path of part (i) being financed after the fact rather than in advance. Two cautions attach. The transition was forced by consumption, not chosen, so it says nothing about whether pre-positioning that capacity would have been optimal; and whether the resulting experience converts into cost decline is a question about β, whereas whether any experience is generated at all was a question about the operational regime. The proposition speaks only to the second.

The same logic warns against reading the regimes off the paper's title dichotomy. Cost determines whether replacement is *feasible*, not whether it *happens*: an expensive interceptor consumed in an intense campaign enters the replacement regime, while a cheap drone held in a peacetime stockpile does not. What the model separates is consumed and replaced systems from held ones, which is an operational distinction that cuts across the cheap–exquisite one.

8 The full two-stage equilibrium

Sections 5–6 analyzed separable slices. The full model lets each side choose *both* weapons under learning, so the arsenal mix, the production scale, and the contest are determined jointly. I work in the additive two-good formalization of Section 4.3 (a cheap drone m and an exquisite munition q combined by the CES aggregator (5)), now with Wright learning on the mass good. Stage 1 each side produces (m_i, q_i) at cost $c_{m0} m_i^{1-\beta} + c_q q_i$, i.e. the additive budget with the constant mass price $c_m m$ replaced by the learning cost $c_{m0} m^{1-\beta}$; Stage 2 the contest (8) resolves.[9] The single-weapon multiplicative baseline of Section 4.2 delivers the same forces in continuous form (Appendix B).

Remark 7 (Interior mixing under learning). One might expect learning to drive *specialization*, an all-mass or all-precision corner, since the mass cost $m^{1-\beta}$ is concave; the isoquant's curvature resists it. Which force wins is exactly the condition in Proposition 9. Writing the arsenal by its mass

[9]The CSF realizes the outcome given fielded forces, so the strategic game is a production contest with cost-minimizing arsenal design nested in the arming first-order condition. The two-stage label marks the economic separation of industrial choices from combat. The sequential-move variant (Section 6.2) and the prior learning-rate choice (Remark 4) put genuine strategic content into the production stage.

contribution share $\omega \in [0,1]$, with $(e_m m)^\rho = \omega X^\rho$, conditional cost is

$$C_X(\omega) \;=\; c_{m0}\,(X/e_m)^{1-\beta}\,\omega^{(1-\beta)/\rho} \;+\; \pi_q X\,(1-\omega)^{1/\rho}, \tag{13}$$

and the cost of the first sliver of mass behaves like $\omega^{(1-\beta)/\rho-1}$: it vanishes when $\beta < 1/\sigma$ (mass always enters) and diverges when $\beta > 1/\sigma$ (the precision corner is locally stable, and cheap mass must arrive at scale or not at all). The marginal cost of mass, $MC_m(m) = c_{m0}(1-\beta)m^{-\beta}$, falls with volume, which is what makes the mix scale-dependent in both regimes.

Proposition 9 (Scale-dependent mix and the shape of adoption)**.** *Let $\sigma > 1$ and $\beta \in (0,1)$, with conditional cost* (13)*.*

(i) If $\beta < 1/\sigma$, the cost minimum is a unique interior mix at every scale, and the mass contribution share $\omega^*(X)$ rises continuously from 0 *to* 1 *as X grows: cheap-mass adoption is* smooth*.*

(ii) If $\beta > 1/\sigma$, the all-precision corner is the global optimum below a threshold scale $\hat{X} > 0$, and at $\hat{X}$ the optimum jumps *to an interior mass-heavy mix, with $\omega^*(X)$ rising thereafter: adoption is* discrete*, with minimum efficient scale $\hat{X}$. The threshold is bounded by the all-mass/all-precision average-cost crossover, $\hat{X} < X^c = e_m\,(c_{m0}/(\pi_q e_m))^{1/\beta}$. More sharply, if $u = (1-\beta)/\rho$, $v = 1/\rho$, and $\omega^\dagger$ is the unique positive root of*

$$\log\!\left[1 + \left(\frac{v}{u} - 1\right)\omega\right] + (v-1)\log(1-\omega) = 0,$$

then $\hat{X} = X^c/(\tau^\dagger)^{1/\beta}$, where $\tau^\dagger = (v/u)(\omega^\dagger)^{1-u}(1-\omega^\dagger)^{v-1} > 1$, so the bound $\hat{X} < X^c$ is immediate. Thus the threshold falls when mass becomes cheaper or more effective, or precision becomes more expensive; its comparative static in β is not unconditional.

In both regimes more force means more mass; the all-mass corner is never optimal, the mass share approaching one only asymptotically.

Strong learning thus changes not just how much cheap mass to buy but *how it arrives*: below $\hat{X}$ a cost-minimizing military correctly buys none, above $\hat{X}$ it re-tools at once. At the calibrated substitution $\sigma = 3$ the regime boundary is $\beta = 1/3$, squarely inside the empirical learning-rate range of Section 10, so both adoption shapes are live possibilities. Feeding the global conditional cost $C(X)$ into the contest, the symmetric equilibrium force solves $X\,C'(X) = RV/4$ on each branch of the arsenal regime; the global best-response scan confirms the equilibrium at moderate learning and its failure in the decisive-race region.

Proposition 10 (Stakes tip every stable equilibrium branch toward mass)**.** *Along any regular locally stable symmetric pure-strategy equilibrium branch, equilibrium force and the mass share increase with the stake V. If $\beta < 1/\sigma$, the tilt is continuous. If $\beta > 1/\sigma$, any transition from the pure-precision branch to the interior mass branch is a discrete jump and must occur no later than*

$$\bar{V}^{\mathrm{ad}} \equiv \frac{4\pi_q \hat{X}}{R} < \frac{4\pi_q X^c}{R}.$$

The bound is sharp for survival of the precision branch, but non-convexity permits a global switch at a lower stake. In the mass-dominated limit, equilibrium force is asymptotically super-linear in the stake, $d\log X^e/d\log V \to 1/(1-\beta) > 1$.

The paradigm bites hardest in large, high-stakes conflicts (Figure 8). Low-stakes disputes are fought with a precision-heavy arsenal (mass has not the scale to be worth learning down); high-stakes wars pull production past the adoption threshold $\hat{X}$ and the arsenal tips to cheap mass, matching the observation that mass drones became decisive in a full-scale war of national survival, not in limited operations, and, in the strong-learning regime, that the switch was rapid rather than gradual. This tip is further *amplified* when quality must be paid on every unit, as in the Becker–Lewis quantity–quality model: the shadow price of quality then rises with fleet size, reinforcing the substitution toward mass (Section 4.2).

Force planning and the mobilization base. The stake-dependence is a planning rule: size the arsenal to the conflict a state expects to fight. A military that might face a large, existential war should hold not a fixed precision inventory but the *latent capacity* to pull production past $\hat{X}$ when the stakes rise, since equilibrium force scales super-linearly with the stake, and, when learning is strong, the optimal switch is a jump rather than a glide. The policy asset is thus the industrial mobilization base and the speed of the use–produce–develop loop, which is also where the first-mover advantage of Proposition 6 is secured; a force sized only for limited operations will be precision-heavy and, in a war of survival, out-massed.

The minimum efficient scale $\hat{X}$ also gives a reading of recent attritable-systems programmes. In the static model, conditional on a fixed mission requirement and abstracting from pilot option value and capacity creation, purchases below $\hat{X}$ leave the scale economy unlocked. The US Replicator initiative, announced in 2023 with the goal of fielding uncrewed systems in the thousands by August 2025 (Sayler, 2026), illustrates the risk of failing to reach operationally meaningful scale: the January 2026 CRS update reports hundreds fielded against the original goal of thousands. Proposals for dedicated government high-volume capacity, sized in the millions of units per year, are the opposite bet: an attempt to arrive above $\hat{X}$ at once rather than to glide toward it. And the comparative static runs in both directions: a cut to the stake term V, as in the reductions to Taiwan's special defence budget reported in 2026, moves the adoption threshold the other way. These examples illustrate the model's geometry; locating $\hat{X}$ empirically requires mission-specific scale and cost data.

Proposition 11 (Phase restrictions)**.** *Under regular locally stable symmetric-branch selection, the* (V, β) *plane has a* precision *region and a* mass *region, separated by an endogenous adoption boundary (located globally), across which the arsenal changes discontinuously when* $\beta > 1/\sigma$*. Independently of that selection, for* $\beta > 1/2$ *there is a* dropout *region where no symmetric interior pure-strategy equilibrium exists. Non-existence holds analytically in the mixed-arsenal model: the symmetric candidate fails once* $\beta\, s_M(X^e) > \frac{1}{2}$*, with* s_M *its mass cost share, so for* $\beta > 1/2$ *failure occurs at all stakes beyond a finite* $\bar{V}$ *(Theorem 2, Appendix C); the pure-mass bound of Proposition 5 is the limit* $s_M \to 1$*, and the global best-response scan locates the boundary at benchmark parameters.*

Propositions 9–11 are what the full model adds: the arsenal mix is endogenous and scale-dependent; cheap-mass adoption is smooth or discrete according to whether learning is weak or strong relative to substitutability (β versus $1/\sigma$); the tip to cheap mass is driven by the stakes; and a decisive-race threshold in the learning rate survives in the joint equilibrium, all in the sharp closed forms the two-good arsenal affords. Appendix C upgrades the arsenal-structure and decisive-race claims from global computations to analytical theorems.

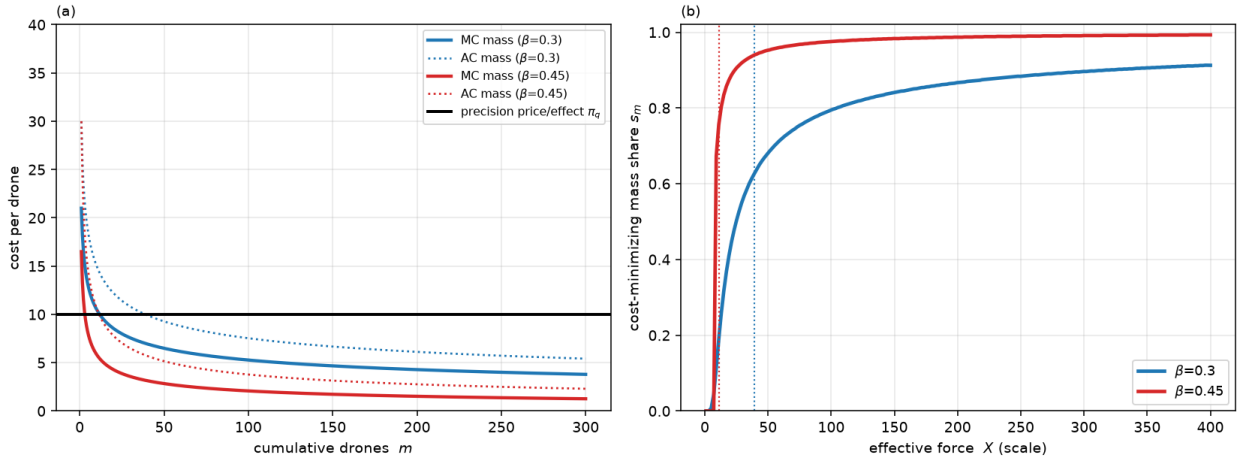


Figure 7: Two-stage cost layer. (a) Learning lowers the marginal (and average) cost of mass below the precision price per unit effectiveness as volume grows. (b) The cost-minimizing mass share rises with the scale of force required.

9 A two-period learning race

Everything so far compressed learning into a static scale economy: cumulative output and fielded quantity coincide. Introducing an experience state separates the investment effect of current production from the evolution of relative leads.

Environment. Each side i carries *inherited experience* $Q_{i,t} > 0$. The marginal cost of the next unit is the Wright curve in the state, $c_0 Q^{-\beta}$, so a production run of m starting from Q costs the integral of the marginal-cost curve,

$$C(Q, m) \;=\; \frac{c_0}{1-\beta}\Big[(Q+m)^{1-\beta} - Q^{1-\beta}\Big], \qquad Q_{i,t+1} = Q_{i,t} + m_{i,t}. \tag{14}$$

There are two periods; each period a Tullock contest ($R = 1$) over a flow stake V resolves on the quantities produced that period; the future is discounted by δ. Outputs are chosen simultaneously each period and observed; I solve for subgame-perfect equilibrium by backward induction. Two consistency checks tie this to the static model: $C(Q \to 0, m) \to \frac{c_0}{1-\beta} m^{1-\beta}$, so the reduced form of Section 6 is exactly the frontier case of zero inherited experience; and at $\beta = 0$ the equilibrium is the static Tullock benchmark $m^* = V/(4c_0)$ in both periods with no intertemporal distortion.

Let $J_i(Q_{i1}, Q_{j1})$ denote side i's period-2 equilibrium continuation payoff at the inherited states. Define the strategic-response component

$$\Sigma_i \equiv V \,\frac{\partial p_2}{\partial m_{j2}} \frac{\partial m^*_{j2}}{\partial Q_{i1}},$$

which captures how an increment to i's experience changes the rival's terminal output.

Proposition 12 (Production is investment near symmetry)**.** *In any regular interior subgame-perfect equilibrium, period-1 output satisfies*

$$\underbrace{\frac{V m_{j1}}{(m_{i1}+m_{j1})^2}}_{\text{contest MB}} \;=\; \underbrace{c_0\,(Q_{i0}+m_{i1})^{-\beta}}_{\text{marginal cost}} \;-\; \delta\,\frac{\partial J_i}{\partial Q_{i1}},$$

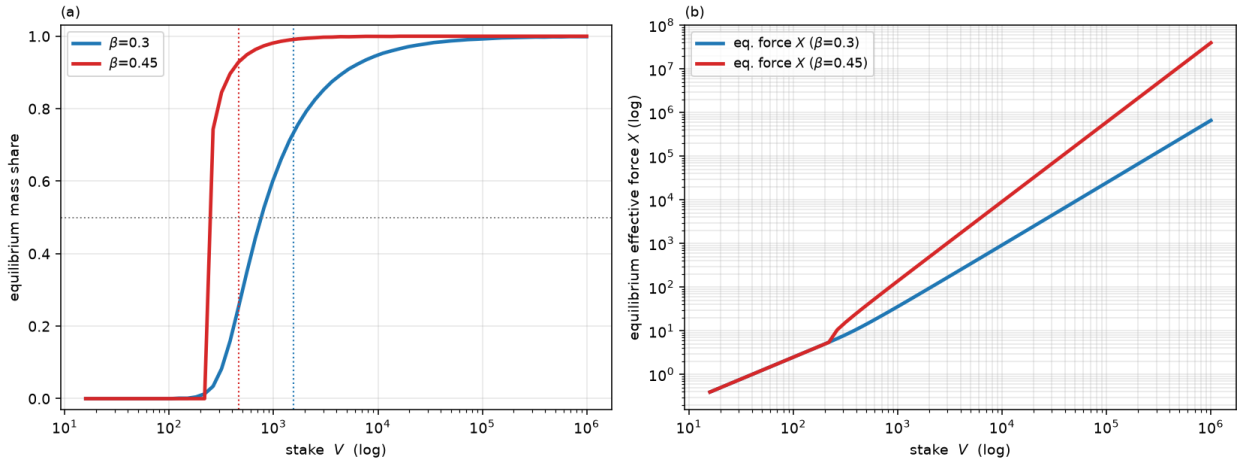


Figure 8: Full two-stage equilibrium. (a) The equilibrium mass share rises with the stake V, tipping the arsenal to cheap mass. (b) Equilibrium force scales super-linearly with the stake.

where

$$\frac{\partial J_i}{\partial Q_{i1}} = \underbrace{c_0\left[Q_{i1}^{-\beta} - (Q_{i1} + m_{i2}^*)^{-\beta}\right]}_{\textit{cost saving (envelope)}} + \underbrace{\Sigma_i}_{\textit{strategic}}.$$

At a symmetric state and symmetric technology, $\Sigma_i = 0$ exactly and

$$\frac{\partial J_i}{\partial Q_{i1}} = c_0\left[Q_{i1}^{-\beta} - (Q_{i1} + m_{i2}^*)^{-\beta}\right] > 0.$$

Consequently symmetric period-1 production strictly exceeds the symmetric myopic equilibrium output for every $\beta > 0$ and $\delta > 0$; by continuity, the conclusion holds in a neighborhood of symmetry. Away from symmetry Σ_i can have either sign, so no global overproduction claim is made. At the benchmark ($\beta = 0.35$, $Q_0 = 2$, $V = 10$, $\delta = 0.9$) the wedge is 28%.

This is the dynamic footing for the use–produce–develop loop: at and near symmetry, current production is priced below the static optimum because each unit buys down the future cost curve. The result is the contest analogue of learning-by-doing pricing in Spence (1981); Fudenberg and Tirole (1983); Cabral and Riordan (1994). The next result establishes how inherited experience leads evolve.

Proposition 13 (A two-period contraction theorem)**.** *Let both sides share (c_0, β), let $\delta \in [0, 1]$, and consider a regular interior subgame-perfect equilibrium satisfying the strict second-order conditions. Write $r_t = Q_{At}/Q_{Bt}$ and suppose $r_0 > 1$. If*

$$0 \le \beta \le \frac{1}{2},$$

then the inherited lead contracts at both dates:

$$1 < r_2 < r_1 < r_0.$$

The result covers the forward-looking period-1 choices, including the strategic effect of current experience on both terminal outputs. In the terminal subgame the simpler identity $m_{A2}/m_{B2} = r_2^{\beta}$ obtains.

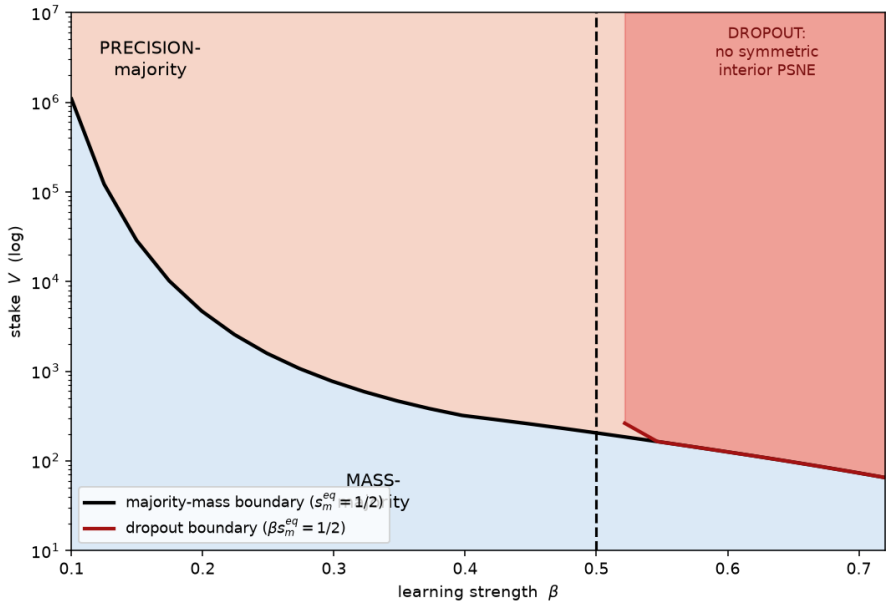


Figure 9: Phase diagram in the stake–learning (V, β) plane. The black curve marks the equilibrium majority-mass boundary. The red curve is the onset of the analytical dropout condition $\beta s_M(X^e) > 1/2$ (equality when the share is continuous); above it the symmetric interior candidate has negative payoff and cannot be an equilibrium. The dropout region is stake-dependent, not the whole half-plane $\beta > 1/2$.

Proposition 13 is stronger than a terminal or myopic argument. Backward induction adds a strategic incentive for the experienced side to invest today, but below $\beta = 1/2$ that incentive is provably too small to overturn catch-up.

Proposition 13 guarantees contraction throughout the weak-learning region $\beta \leq 1/2$. The proof evaluates its strategic bound at $\beta = 1/2$, and numerically the contraction region extends somewhat above that sufficient boundary, making the analytical bound conservative in those computations. The economically relevant asymmetry is sharp: weak learning *cannot* snowball a lead, while strong learning *can*. Under weak learning, additive accumulation dilutes proportional leads. Under strong learning, even an interior race can snowball: at $(\beta, \delta, V, c_0) = (0.988, 0.534, 0.664, 1)$ and $(Q_{A0}, Q_{B0}) = (2.413, 0.772)$, a globally verified interior SPE gives $r_0 = 3.1259 < r_1 = 3.1549$, followed by terminal contraction to $r_2 = 3.1489$. At still larger stakes the interior equilibrium itself can fail (with $Q_A = Q_B = 1$ and $\beta = 0.55$, the terminal subgame has a pure-strategy equilibrium at $V = 5$ but fails the global best-response check at $V = 400$). Thus there are three economically distinct outcomes: catch-up, which is guaranteed below the threshold; interior snowballing, which requires strong learning; and dropout/non-existence, which requires stakes high enough that $\beta s_M(X^e) > \frac{1}{2}$ (Theorem 2).

10 A first calibration: quantitative illustration

Pending a clean within-system drone learning series, I calibrate at a benchmark learning elasticity taken from the products a drone is built from. Battery, solar, and consumer-electronics learning rates cluster near 20% per doubling (Nagy et al., 2013; Farmer and Lafond, 2016; Way et al., 2022), i.e. $\beta = -\log_2(0.8) \approx 0.32$; I use the round benchmark $\beta = 0.30$ (range $[0.20, 0.40]$). The cost

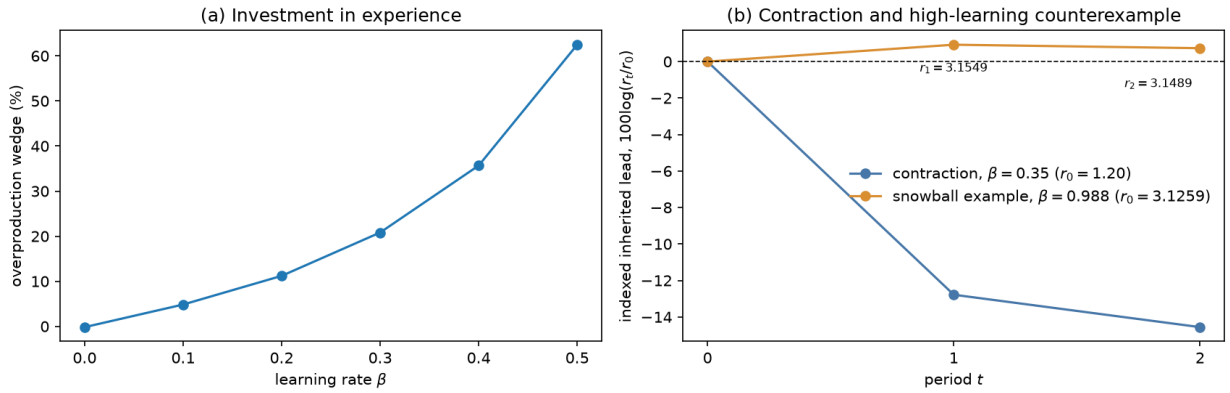


Figure 10: The two-period learning race. (a) At symmetry, period-1 production exceeds the myopic benchmark because experience lowers terminal costs. (b) Indexed inherited leads, $100\log(r_t/r_0)$: the lead contracts at both dates in the $\beta = 0.35$ example, while a globally verified $\beta = 0.988$ interior equilibrium snowballs in period 1 before contracting slightly in period 2. Proposition 13 guarantees contraction for $\beta \leq 1/2$; the high-learning path shows that contraction can fail beyond this sufficient region. Raw initial ratios are $r_0 = 1.20$ and $r_0 = 3.1259$.

primitives come from the assembled panel: an FPV at \$400, a Shahed-class one-way-attack drone at \$59k, a Sting interceptor drone at \$2,100 (the panel median; kill probability 0.68), and a Patriot interceptor at \$3.25m; I set the substitution $\sigma = 3$, the "good-enough" elasticity $b = 0.5$, and $R = 1$. Three of the paper's objects are then quantified; two are pinned by β alone, one by cost ratios, and the arsenal mix by an assumed $\mathcal{E}$ (Figure 11). The exercise is a quantitative illustration disciplined by data: β, σ and $\mathcal{E}$ are assumed, so the numbers below support orders of magnitude and thresholds. Appendix D bounds the survival component of $\mathcal{E}$, which is the only one of the three the available data speak to.

The benchmark lies below the decisive threshold. At $\beta = 0.30$ the arming game has effective decisiveness $1/(1-\beta) = 1.43$ (two-good) or $(1-b\beta)/(1-\beta) = 1.21$ (baseline). Both are below the winner-take-all value of 2, reached only at $\beta = 1/2$ (two-good) or $1/(2-b) = 0.67$ (baseline). So at ordinary electronics learning rates cheap mass is favored and exhibits increasing returns, but the race is *not* winner-take-all: the fast "test-in-the-field" iteration loop would have to lift the effective β toward 1/2 to trigger a decisive one. Equilibrium force scales as $X^* \propto V^{1.43}$, so doubling the stake raises fielded force about 2.7-fold.

Cost-exchange (data). From cost ratios alone, a \$59k Shahed is offense-favored against exquisite interceptors ($A_{\text{off}} \geq 55$ against a Patriot and ≥ 17 against NASAMS, both at $\kappa = 1$ and so lower bounds) but defense-favored against a cheap interceptor drone, $A_{\text{off}} \approx 1/19$ at its reported kill probability (Figure 1).

The arsenal mix is mission-specific. The mix compares systems competing for the *same* mission, so it requires a cheap and an exquisite way of performing one task, and the panel supplies two such pairs. In *tactical* strike, an FPV at \$400 and a Javelin round at \$209 k attack the same armour and point targets, both man-portable, both flown or fired at a target the operator can see.[10] In *deep* strike, a long-range one-way-attack drone at \$70 k and a Tomahawk, JASSM-ER or PrSM at

[10]The Javelin is the closer pair than the air-launched Joint Air-to-Ground Missile the panel also prices, whose \$302 k embeds a helicopter's mission rather than an infantryman's. The figure is the recurring production cost of one more *round*: the reusable launch unit is a separate procurement line and is excluded, exactly as an FPV pilot's goggles and controller are excluded from the \$400 airframe.

a median \$2.6 m deliver a warhead to a fixed target hundreds to thousands of kilometres away.[11] Both gaps are far below the raw ladder of Figure 1a, which spans systems that do not compete for the same task.

The comparison then has a sharp form. Because the relevant price is the price of delivered effect, $\pi_m = c_m/e_m$ against $\pi_q = c_q/e_q$, the arsenal tips toward mass exactly when the precision-effectiveness advantage falls short of the price gap,

$$\mathcal{E} \;<\; c_q/c_m.$$

The central pairings put that threshold at about 521 in tactical strike and about 37 in deep strike. These are benchmark thresholds assembled from a mostly Russian cheap pole and an all-US exquisite pole. A structural same-arsenal estimate would pair both prices with the same side's effectiveness advantage $\mathcal{E}$ (Appendix D). Across defensible pairings the gap runs from about 249 to 1,073 in tactical strike and from about 7 to 55 in deep strike, and Figure 11b shades both bands rather than drawing a line.

The sensitivity range for $\mathcal{E}$ runs to 100, making its upper end *generous to precision*: if cheap mass still dominates a mission when an exquisite round is assumed a hundred times more effective per unit, the conclusion is safe at any smaller advantage. How generous can be checked against the decomposition. $\mathcal{E}$ is survival times accuracy given survival times damage given a hit, and Appendix D measures the survival term at 0.70 for cruise systems and 3.41 for ballistic ones. An $\mathcal{E}$ of 100 therefore requires accuracy and warhead effect together to supply a factor of roughly 30 to 140, depending on the class. That demanding requirement makes 100 a useful upper sensitivity bound.

The sensitivity exercise yields a substantive asymmetry. *Tactical* strike is robust: at the calibrated $\sigma = 3$ the cost-minimizing mass *cost* share stays between 0.86 and 0.99 at $\mathcal{E} = 100$ across the entire range of pairings, so an FPV arsenal is optimal unless a \$200 k missile is two orders of magnitude likelier to deliver its effect. In *deep* strike the crossover moves by a factor of roughly seven across the same exercise, which is precisely the interval in which evidence would have to land to decide the question. At a central $\mathcal{E} = 30$ the mass cost share is near one-half and at $\mathcal{E} = 100$ the arsenal is precision-dominated *by cost*, but where the tipping point actually sits remains open. Within the displayed $\mathcal{E} \in [1, 100]$ sensitivity range and at $\sigma = 3$, tactical strike robustly favors mass across the reported price pairings; deep strike spans both sides of the adoption boundary. Onboard autonomy can move that boundary by lowering $\mathcal{E}$ (Remark 1), and since tactical strike clears its threshold by two orders of magnitude while deep strike does not, autonomy is a prediction about deep strike in particular — the mission whose boundary is close enough to be crossed.

The unit of that statement needs emphasis, because it is easy to misread against what campaigns look like. Throughout, s_m is a share of *spending*, not of rounds. A large price gap makes the two diverge mechanically: an arsenal can be overwhelmingly mass by count and overwhelmingly precision by cost at the same time, and observed deep-strike campaigns are exactly that. Appendix D compares the same attacker's cheap and exquisite systems in the same defended airspace and isolates the survival component of the effectiveness gap.

[11] Exquisite prices here are the three round costs of Cancian and Park (2026). The panel also records cumulative-average procurement costs for Tomahawk, PrSM and JAGM from the FY2027 budget books; those include support and so overstate the marginal round, and pooling the two bases in one statistic would be worse than choosing either. All figures below use round costs.

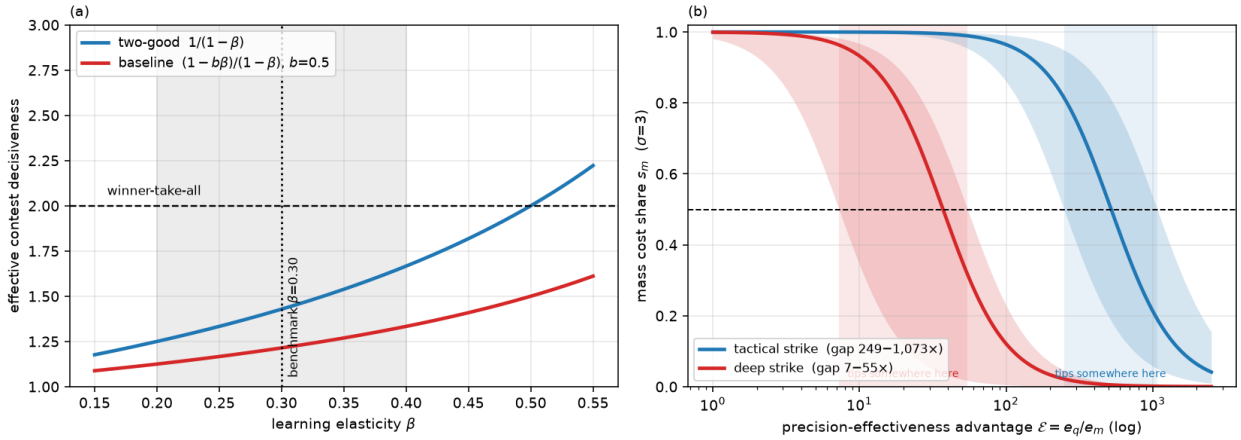


Figure 11: Calibration at benchmark $\beta = 0.30$ (range $[0.20, 0.40]$ shaded) and Ukraine-era cost primitives. (a) Effective contest decisiveness is below the winner-take-all value of 2 at the benchmark; the two-good race becomes winner-take-all only at $\beta = 1/2$. (b) Mission-matched arsenal mix at the calibrated $\sigma = 3$: cost-minimizing mass cost share against the assumed precision-effectiveness advantage $\mathcal{E} = e_q/e_m$, for the two same-mission pairs the panel supports (tactical strike, FPV versus the Javelin round; deep strike, one-way-attack drone versus Tomahawk/JASSM-ER/PrSM at round cost). Each arsenal tips away from mass as $\mathcal{E}$ passes its own price gap. Bands span defensible pairings rather than a single point, since the panel's cheap pole is mostly Russian and its exquisite pole all US: the gap runs 249–1,073× in tactical strike and 7–55× in deep strike. The tactical conclusion is unchanged across its whole band; the deep-strike crossover is not pinned.

11 Bargaining: does cheap mass make war more likely?

The model so far is positive about arming and the contest. Embedding it in the bargaining theory of war (Fearon, 1995; Powell, 2006) converts the changed conflict technology into a statement about the probability and character of war. Two states dispute a divisible stake of value V. They may settle peacefully (state A receives a share $\alpha \in [0, 1]$, worth αV, and B receives $(1 - \alpha)V$) or fight. War is a costly lottery: A prevails with probability p (the contest win probability of the model) and takes V, and each side bears a cost $k_i > 0$. Expected war payoffs are $pV - k_A$ for A and $(1 - p)V - k_B$ for B.

Proposition 14 (Bargaining range, restating Fearon, 1995)**.** *Both states prefer a settlement* $\alpha \in [0, 1]$ *to war iff* $\alpha \in [\max\{0, p - k_A/V\},\ \min\{1, p + k_B/V\}]$. *A peaceful settlement both prefer therefore always exists; the set has* width $\min\{1, p + k_B/V\} - \max\{0, p - k_A/V\} \in (0, 1]$, *equal to* $(k_A + k_B)/V$ *when neither bound binds and saturating at the whole settlement space when both do, and* center $p + (k_B - k_A)/2V$ *in the unconstrained case. Cheap-mass adoption enters through two objects: the win probability* p*, which shifts the* center *toward the more productive side (a better settlement, not war), and the war costs* k_i*, which set the* width.

A peace range of positive width therefore always exists, and a productivity edge from cheap mass slides it toward the adopter (Figure 12). What the paradigm does to the *width* is not free to assume: the war costs k_i are objects the model already determines, and deriving them from it signs the static effect.

Proposition 15 (War costs from the contest: symmetric cheap mass widens the static peace range)**.** *Decompose* $k_i = a_i + d_i$ *into the cost of* mounting *war and the* destruction suffered, *and tie both*

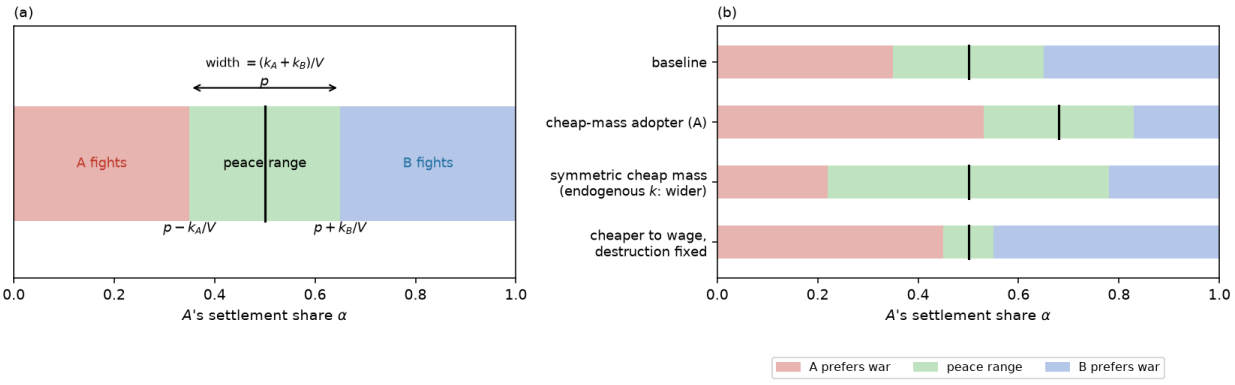


Figure 12: The bargaining range over A's settlement share α. (a) A peace range of width $(k_A + k_B)/V$ separates "A fights" from "B fights." (b) A productivity edge (cheap mass) re-centers the range toward the adopter (a better settlement, not war). Holding destruction fixed, cheaper fighting would narrow the range (top row), but with war costs derived from the model, symmetric cheap mass *widens* it: arming is unchanged while destruction per dollar rises (Propositions 14–15).

to the contest's equilibrium objects under two maintained assumptions: mounting costs scale with the arming that fighting requires, $a_i = \tilde{\lambda} B^*$ with $\tilde{\lambda} \geq 0$ the share not already sunk at the settlement stage, and destruction suffered scales with the adversary's delivered force, $d_i = \chi X_j^*$ with $\chi > 0$. Under symmetric adoption, $B^* = RV/4$ is unchanged as the price of mass falls (Proposition 3) while $X^* = \psi B^*$ rises without bound, so the peace-range width is

$$\frac{k_A + k_B}{V} \;=\; \frac{R}{2}\left(\tilde{\lambda} + \chi\psi\right) \qquad \textit{(truncated at 1, Proposition 14)},$$

strictly increasing as π_m falls, until the range covers the whole settlement space and war is never preferred. The result requires no arming term: it holds for $\tilde{\lambda} = 0$ (all arming sunk), because the driver is destruction. Symmetric cheap mass widens *the static bargaining range: what the paradigm cheapens is the delivery of destruction, which raises the cost of suffering war faster than it lowers the cost of waging it.*

The classic narrowing mechanism, cheaper war consuming the peace range (Fearon, 1995), operates only if fighting becomes cheaper *without* becoming more destructive. Under the maintained assumptions the model signs it the other way: arming budgets do not fall (Proposition 3), and the productivity that rises is destructiveness per dollar, which enters the rival's d_j. This differs from Chang and Luo (2017), where destruction is a strategic choice along a fixed technology; here it rises through the technology itself, at unchanged budgets. The static effect of the paradigm is thus, if anything, stabilizing, and the danger must be dynamic.

The commitment channel, derived. The war risk the paradigm creates is dynamic, and it arrives through the Powell (2006) commitment problem. The two-period model of Section 9 supplies the power-transition law, and the following protocol supplies the bargaining game. Two periods, flow stake V each period, discount δ. At the start of each period the sides may settle on any division of that period's flow (transfers are capped by the flow itself, so neither side can commit future divisions) or fight. War is decisive and absorbing: the winner takes the current and all remaining flows, and each side pays its cost k_i once. Win probabilities are p_1 in a period-1 war and p_2 in a period-2 war (side A's probability, A rising: $p_2 \geq p_1$).

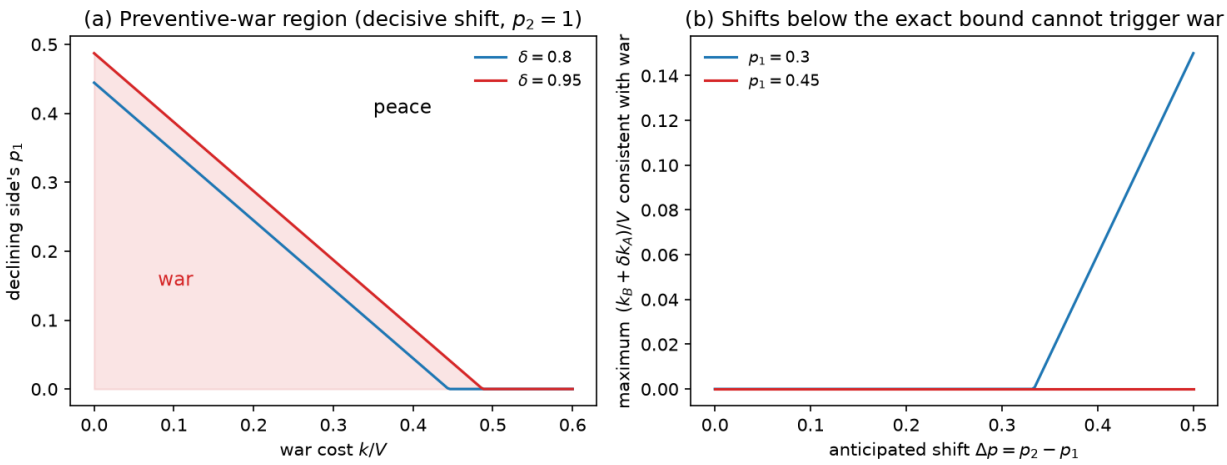


Figure 13: Preventive war under bounded per-period transfers (Proposition 16). (a) Under the commitment-selected illustration ($p_2 = 1$), the declining side fights iff its current strength and the war cost fall in the shaded region, $p_1 < \delta/(1+\delta) - k/V$. (b) Shifts below the exact bound cannot reach the war region: the maximum war cost consistent with fighting is zero until the shift Δp exceeds the exact bound p_1/δ.

Proposition 16 (Preventive war under a bounded-transfer power shift)**.** *In the protocol above, a peaceful path exists in period 1 if and only if*

$$\delta\,(p_2 - p_1) \;\le\; p_1 + \frac{k_B + \delta k_A}{V},$$

and the rising side never blocks peace. Hence:

(i) Small power shifts cannot cause preventive war in this protocol. *If $\delta\Delta p \le p_1$ (equivalently $\Delta p \le p_1/\delta$ when $\delta > 0$), peace survives at every $k_A, k_B \ge 0$.*

(ii) A commitment-selected decisive shift can. *If a separate sequential production commitment selects the deterrence outcome of Proposition 6, so that $p_2 = 1$, and $k_A = k_B = k$, the declining side fights if and only if it is currently strong,*

$$p_1 \;<\; \frac{\delta}{1+\delta} - \frac{k}{V}.$$

The proposition takes the power path as given. The learning race supplies one restriction on that path: for $\beta \le 1/2$, inherited experience ratios contract, but this alone does not bound Δp under asymmetric technologies. Part (i) gives the exact sufficient bound for any smooth or abrupt shift. Part (ii) applies the separate commitment timing of Proposition 6, under which $p_2 = 1$, rather than deriving that outcome from the simultaneous two-period game. Under this timing, a currently dominant power may fight to stop a fast-learning challenger exactly when the cost-adjusted condition holds. (Derivation and verification: Appendix A.)

The character of conflict the paradigm favors is thus conditional. Statically the peace range *widens* (Proposition 15) and a productivity advantage merely re-centers it. Dynamically, the exact preventive-war condition depends on the magnitude of the induced power shift and on the transfer cap. Low learning contracts experience ratios; high learning can snowball them; and sequential

commitment can create the extreme $p_2 = 1$ shift used in part (ii). Cheap mass therefore does not mechanically cause war: the model identifies the power transition large enough to do so under a stated protocol.

12 Discussion

Is cheap mass durable? The model's answer is conditional (Proposition 7): durability turns on whether offensive or defensive costs fall faster with cumulative production. Because drones are electronics-heavy, one might expect steep offensive learning; but counter-drone systems (interceptor drones, autonomy, EW) are also electronics and learn too. General-purpose-technology accounts (Ding and Dafoe, 2023) suggest the deep driver, cheap computation and autonomy, is durable even if specific airframes are not, cautioning against reading today's cost-exchange ratios as permanent.

Deterrence and budgets. Proposition 3 implies that symmetric diffusion of cheap mass raises destructiveness per dollar rather than budgets: deterrence balances can shift sharply with little movement in headline military spending, a measurement caution for empirical work on arms races. Proposition 5 implies that when learning is strong, arming becomes a decisive race, a destabilizing force, because the returns to being first or largest are discontinuous. This yields a decomposition test: in the model, technology (a fall in π_m, a rise in β) moves capability per dollar and the arsenal's composition, while budgets move with stakes and threat perceptions (V). The European rearmament of 2024–26 illustrates the distinction in Proposition 3: budget growth loads on the stake term, while the technology term shows up where the model says it should, in procurement tilting toward volume drones and counter-drone systems rather than proportional platform spending. The European Drone Defence Initiative is an instance: a counter-drone flagship whose stated objectives are affordable technology, mass production, and a drone alliance with Ukraine as the source of iteration experience, with initial operational capability targeted for end-2026 and full functionality for end-2027 (European Commission, 2026, 2025).

Where the boundary falls. The threshold $\mathcal{E} < c_q/c_m$ locates the cheap–exquisite boundary in live procurement, and the two missions of Section 10 show it sitting in different places.

Deep strike. The break-even is the arsenal condition itself: a munition at price c_q is worth buying over a \$70 k one-way-attack drone if and only if $\mathcal{E} > c_q/c_m$. European programmes are now priced close to that line. Under the European Long-Range Strike Approach, six states signed a letter of intent in February 2026 to develop a one-way effector with a range beyond 500 km at a target unit cost below €100,000 (European Long-Range Strike Approach, 2026), a gap of about 1.6×; the UK's Project Nightfall, a ballistic system with a ceiling of £800,000 (UK Ministry of Defence, 2026), is at about 15×. Both sit far inside the 37× the current exquisite inventory would require against the same anchor. The category name is the tell: these programmes pursue cruise-missile range *from the cheap-mass pole*, which is Proposition 10's stake-driven tilt appearing in procurement policy.

The tactical margin. Here the boundary is two orders of magnitude away, and the interesting point is instead a taxonomic one. The exquisite pole is a guided munition such as the Javelin; a dumb 155 mm shell occupies an intermediate position. It costs roughly €2,500–3,000 and needs many rounds per target; an FPV costs \$400 and puts its warhead on the target it is flown at. Ranked by cost per round the FPV is the cheapest of the pair *and* the more accurate, so the wartime move from shells to FPVs is a joint price–quality improvement. It is Proposition 2 in its purest observed form, a fall in $\pi_m = c_m/e_m$ in which both arguments move the same way — which

is why it happened under budget pressure rather than requiring a procurement decision. The two remain imperfect substitutes — artillery delivers volume through jamming, cloud and darkness in a way an FPV does not — which is why σ is finite rather than unbounded at this margin, and why the wartime substitution was partial.

Two cautions carry across both margins. Causal-share estimates are contested and asymmetric: Watling and Reynolds (2025) attribute about two-thirds of Russian losses to tactical drones, while RUSI work also attributes a large majority of *Ukrainian* casualties to artillery. The two findings coexist because the sides face opposite input scarcities, which is the asymmetric adopter of Proposition 2 seen from both ends — but neither supports the claim that drones have replaced artillery. And the price anchors above are not European or Ukrainian respectively, so they inherit the cross-arsenal caveat of Section 10: they locate a boundary rather than estimate it.

Procurement. The increasing returns of Proposition 4 rationalize the shift from slow, exquisite procurement toward fast, iterative, mass production: the strategic payoff is in β and θ, i.e. in the speed and scale of the industrial base, not only in platform quality. Under a hard defense budget (Remark 2), the point sharpens into an allocation rule: with quality paid on every unit and unit costs falling through learning, the budget-maximizing arsenal tilts *toward* cheap mass as the budget grows: optimal quality per unit *falls* in $\bar{B}$ while the number of units rises more than proportionally ($\zeta^* \propto \bar{B}^{-\beta/(1-\beta)}$, $n^* \propto \bar{B}^{1/(1-\beta)}$), and fielded force exhibits increasing returns to the budget. A better-funded military facing learning should therefore buy *more, cheaper* units. This reverses the "spend more, buy more exquisite" instinct and cautions against appropriations that scale platform quality with the budget rather than quantity.

Proposition 6 adds a second procurement reading, and it is uncomfortable. The strategic asset there is the credibility of an early, hard-to-reverse production commitment. Where one side's small-drone output is counted in millions of units a year and rests on a commercial supply chain it also supplies to everyone else, while the other side's flagship attritable-systems programme delivered in the hundreds against a target of thousands (Tollast, 2025; Sayler, 2026), the model says the gap that matters is in commitment capacity rather than in technology. It also says what closing it would require: a production commitment large enough and early enough to be worth pre-empting. Whether any state can make such a commitment credible in peacetime, when the stake V that would justify it has not yet materialized, is precisely the difficulty, and it is the same tension Proposition 16 formalizes from the other side, where a declining power's incentive to act early rises with the size of the shift it anticipates. Applying the result to a live rivalry requires estimating $\delta\Delta p$ against $p_1 + (k_B + \delta k_A)/V$.

Modeling choices and limitations. The paper uses a Tullock ratio-form CSF, a CES aggregator with $\sigma > 1$, learning tied to own cumulative output, and a single divisible stake with linear resource cost. Three parameters are assumed rather than estimated: the learning rate β, taken as a benchmark from the products a drone is built from; the substitution elasticity σ; and the effectiveness advantage $\mathcal{E}$, which enters Figure 11b as an axis for exactly that reason. The calibration is accordingly an illustration disciplined by data, supporting orders of magnitude and threshold locations. The Tullock form has a stochastic microfoundation (Jia, 2008), and the results least dependent on functional form are those the multiplicative baseline reproduces without any CES structure: the productivity-share win probability, capability-not-budgets, the learning threshold, and the stake-driven tilt (Appendix B). The complements case $\sigma < 1$ forces a permanent mix, mass and precision as fire and maneuver rather than substitutes; the paradigm statements survive with damped elasticities, since the price index (7) still falls when π_m falls. Risk neutrality and an

exogenous, fully divisible stake V are likewise maintained hypotheses: risk aversion widens the bargaining range and issue indivisibility narrows it, while leaving the industrial layer unchanged. Each remaining choice is a knob: a difference-form CSF, a microfounded drone-versus-countermeasure duel, learning spillovers beyond the reduced form of Remark 3, and the full asymmetric two-stage equilibrium (beyond the illustration of Section 6.3) are natural extensions. On dynamics, Section 9 now supplies the two-period race with an explicit experience state and Proposition 16 the bargaining embedding; what remains open is the infinite-horizon Markov-perfect version with two-sided offense–defense learning (Q_m and Q_d evolving jointly, Proposition 7) and stochastic innovation arrivals, where the mean-reversion result of Proposition 13 would confront mechanisms that can overturn it (victory-contingent resources, spillover asymmetries, capacity constraints).

13 Concluding remarks

The central result of this paper is that arsenal design and production experience jointly set the curvature of the arming contest. Quantity–quality choice plus learning gives the cost of force increasing returns, and with them an effective decisiveness that crosses into a decisive race at a threshold in the learning rate. Everything else follows from that: the win probability is a productivity share, symmetric adoption raises destructiveness per dollar rather than budgets, adoption is smooth or discrete according to how learning compares with substitutability, and the experience that drives all of it accrues only where weapons are consumed and replaced.

It is worth saying what that would mean if it is right, because the direction is unfamiliar. For most of the modern period, military technology raised the cost of contesting. Gunpowder artillery and the fortifications built to resist it priced war beyond what any but a consolidating fiscal state could sustain, and the state that emerged was shaped by the requirement to pay for it (Parker, 1988; Tilly, 1990). Industrial mobilization, air power and finally nuclear weapons each concentrated decisive capability further, in fewer and richer hands. The tendency was never uniform — cheap small arms, rockets and improvised explosives let poor combatants impose serious costs on rich ones throughout the twentieth century — but what those weapons could not do was destroy a chosen high-value target at distance. That is the capability the drone era has made cheap, and it is why the shift registers as a change in $\pi_m = c_m/e_m$ with both arguments moving at once rather than as one more insurgent's bargain.

The model says where that leads. In equilibrium the win probability is a share of the *productivity* of spending (Lemma 2): what decides the contest is how cheaply a side converts money into delivered effect, not how much money it brings, provided it can afford its own optimum at all. Falling unit costs bring more contestants inside that proviso. Where a defense budget is capped rather than chosen (Remark 2), the fiscal base returns as a direct constraint — so this is a shift in what military power depends on, not an escape from wealth.

Four implications follow, and each is a statement the model supports rather than a forecast. *Arms races become hard to see.* Proposition 3 says symmetric adoption leaves budgets unchanged while destructiveness per dollar rises, so balances can move sharply with no signal in the spending series that scholarship and policy both watch. *Adoption can be lumpy.* In the strong-learning regime of Proposition 9, the cost-minimizing arsenal jumps from pure precision to a mass-heavy mix at the minimum efficient scale. Pilot buys and capacity investments remain separate margins; operational procurement must eventually cross the scale threshold to unlock the production economy. *Reconstitution moves the experience state.* Once a target stock is complete, an unconsumed system generates no further production experience. Sustained consumption and replacement continue to move the cost curve, as does any production programme that keeps cumulative output

growing (Proposition 8). *The danger is the transition.* Statically, cheap mass *widens* the peace range (Proposition 15); what the bargaining literature identifies as war-prone is the fast, potentially decisive shift in relative power that Propositions 5 and 16 describe. On this reading the hazardous period is not the one in which everyone has cheap mass, but the one in which relative productivity is still moving — which is the period now under way.

Taken together these point at something larger than procurement. If the binding constraint on military power moves from the ability to finance a few capital-intensive systems toward the ability to organize a commercial production, feedback and iteration ecosystem, then the advantage accrues to a different kind of state — and to a different part of the state. Cheap mass does not abolish state capacity; it changes what state capacity consists of.

The framework does not settle whether that will happen, and the calibration is deliberately an illustration rather than an estimate. What it does is convert "affordable mass precision" from an observation about one war into an economics of conflict, in which the strategic asset is neither the drone nor the factory nor the budget but the organization that learns fastest from what it spends.

A Proofs

Proof of Lemma 1. Minimize $c_m m + c_q q$ subject to (5) equal to a target X. Cost minimization of a CES technology yields the dual CES price index over the input prices per unit of effectiveness, $\pi_m = c_m/e_m$ and $\pi_q = c_q/e_q$, giving (7); the mass cost share follows from Shephard's lemma, $s_m = \pi_m \, \partial P/\partial \pi_m / P$. Since the technology is constant-returns, total cost is $P \cdot X$, so $X = B/P = \psi B$. □

Proof of Lemma 2. Let $y_i = \psi_i B_i$ denote effective effort, so $B_i = y_i/\psi_i$ and $U_i = y_i^R/(y_i^R + y_j^R)\, V - y_i/\psi_i$. The first-order condition is $R\, y_i^{R-1} y_j^R (y_i^R + y_j^R)^{-2} V = 1/\psi_i$. Dividing the two players' conditions gives $y_A/y_B = \psi_A/\psi_B$, independent of R; hence $X_A/X_B = \psi_A/\psi_B$ and (9). Because $y_i = \psi_i B_i$, the effort ratio implies $B_A = B_B$. For $R = 1$ the conditions solve in closed form to $B_A = B_B = P_A P_B V/(P_A + P_B)^2$; at $\psi_A = \psi_B$ the symmetric condition gives $B^* = RV/4$, independent of the common P. □

Proof of Propositions 2–3. $\partial P/\partial \pi_m > 0$ from (7), so $\psi_i \uparrow$ and, by (9), $p_i \uparrow$ when $\pi_m \downarrow$. Differentiating the mass share, $\partial s_m/\partial \pi_m < 0 \iff \sigma > 1$. For Proposition 3, set $\psi_A = \psi_B$: by Lemma 2(iii) $B^* = RV/4$ is independent of the common P, while $\psi = 1/P \to \infty$ as $\pi_m \to 0$. □

Proof of Propositions 4–5. From (10), spending $B_m = c_{m0} m^{1-\beta}$ inverts to $m = (B_m/c_{m0})^{1/(1-\beta)}$, so $e_m m = \theta B_m^\gamma$ with $\gamma = 1/(1-\beta)$ and $\theta = e_m c_{m0}^{-1/(1-\beta)}$; $\gamma > 1$ gives strict convexity. With both sides all-mass and $R = 1$, $X_i = \theta B_i^\gamma$ and $p_i = B_i^\gamma/(B_i^\gamma + B_j^\gamma)$, a Tullock contest with decisiveness γ. The symmetric first-order condition yields $B^* = \gamma V/4$; existence of a pure-strategy equilibrium in a two-player Tullock contest requires decisiveness ≤ 2, i.e. $\gamma \leq 2 \iff \beta \leq 1/2$ (Pérez-Castrillo and Verdier, 1992; Baye et al., 1994). For general R the effective decisiveness is $R\gamma$. □

Proof of Proposition 6. Spending above V is dominated. Conditional on participation the leader's strategy set is $[\underline{B}, V]$, with concession as an outside action, and the follower's strategy set is $[0, V]$; both are compact. Treating null commitment as concession closes the standard zero-effort discontinuity of the Tullock contest. Given $B_L = b_L \geq \underline{B}$, the follower solves $\max_{x \in [0,V]} f(x; b_L) = V x^\gamma/(b_L^\gamma + x^\gamma) - x$. Continuity gives a nonempty compact best-response set. An active best response with positive payoff is interior, since spending V yields a weakly negative payoff, and

therefore satisfies the first-order condition below. Select dropout when the follower's maximal payoff is zero, as stated in the proposition; among any remaining best-response ties, select the response the leader prefers. Berge's maximum theorem and this boundary selection make the induced leader payoff upper semicontinuous: at the participation boundary the leader receives the upward jump to the dropout payoff. It therefore attains a maximum on $[\underline{B}, V]$; comparing that maximum with concession proves (i).

For (ii), the follower's maximal payoff decreases in b_L. Let $\bar{B}$ solve $\max_x f(x;\bar{B}) = 0$. At this boundary the positive follower optimum x^* satisfies $f(x^*;\bar{B}) = f_x(x^*;\bar{B}) = 0$. The first-order condition gives $V\gamma x^{\gamma-1} b_L^\gamma = (b_L^\gamma + x^\gamma)^2$, while zero payoff gives $Vx^{\gamma-1} = b_L^\gamma + x^\gamma$. Division yields $x^\gamma = (\gamma-1)b_L^\gamma$ and $Vx^{\gamma-1} = \gamma b_L^\gamma$. Eliminating x gives $\bar{B} = (V/\gamma)(\gamma-1)^{(\gamma-1)/\gamma}$. The associated boundary entry is $x^* = V(\gamma-1)/\gamma$. The deterrence level itself is always feasible: with $\bar{B}/V = (1/\gamma)(\gamma-1)^{(\gamma-1)/\gamma}$,

$$\frac{d}{d\gamma}\log(\bar{B}/V) = \frac{\log(\gamma-1)}{\gamma^2},$$

which is positive for $\gamma > 2$, so $\bar{B}/V$ is strictly increasing on $[2,\infty)$ and equals $1/2$ at $\gamma = 2$ (indeed $\gamma = 2$ minimizes it on $(1,\infty)$, since the derivative is negative below 2). Hence $\bar{B} \geq V/2 > \underline{B}$ throughout the winner-take-all region: the floor never binds on the leader, and the tangency and threshold are unchanged.

It remains to compare deterrence with accommodation. At any active follower optimum with nonnegative participation payoff, set $t = x/b_L$. Combining the follower FOC and participation condition gives $t^\gamma \geq \gamma - 1$, strictly when entry yields positive payoff. The leader's gross prize return under accommodation is therefore $V/(1+t^\gamma) \leq V/\gamma$, and its net payoff is strictly lower after deducting $b_L > 0$. For $\gamma \geq 2$,

$$V - \bar{B} \geq \frac{V}{\gamma},$$

because $(\gamma-1)^{(\gamma-1)/\gamma} \leq \gamma - 1$. Hence accommodation is strictly dominated throughout the winner-take-all region. Among deterrent choices the leader selects the least costly, $\bar{B}$, and nets $V - \bar{B} = V\big[1 - (1/\gamma)(\gamma-1)^{(\gamma-1)/\gamma}\big] > 0$, which decreases to 0 as $\gamma \to \infty$. In the simultaneous-move mixed-strategy equilibrium each side's expected payoff is zero (bidding 0 lies in the support), establishing the strict first-mover advantage (Dixit, 1987; Leininger, 1993; Morgan, 2003). □

Proof of Propositions 9–11. Differentiate (13):

$$C_X'(\omega) = c_{m0}(X/e_m)^{1-\beta}\tfrac{1-\beta}{\rho}\,\omega^{(1-\beta)/\rho-1} - \pi_q X\,\tfrac{1}{\rho}\,(1-\omega)^{1/\rho-1}.$$

As $\omega \to 1^-$ the second term vanishes ($1/\rho > 1$) and the first is positive, so the all-mass corner is never optimal and $\omega^* < 1$ always, approaching one asymptotically. (i) If $\beta < 1/\sigma$, i.e. $(1-\beta)/\rho > 1$, the first term vanishes as $\omega \to 0$ while the second tends to $\pi_q X/\rho > 0$: $C_X'(0^+) < 0$, so the precision corner is never optimal. The first-order condition rearranges to $\omega^{(1-\beta)/\rho-1}/(1-\omega)^{1/\rho-1} = \text{const}\cdot X^\beta$, whose left side is strictly increasing in ω from 0 to ∞ (positive exponent in the numerator, decreasing denominator): the interior root $\omega^*(X)$ is unique and strictly increasing in X, from 0 to 1. (ii) If $\beta > 1/\sigma$, $C_X'(0^+) = +\infty$ and the precision corner, with value $\pi_q X$, is a local minimum. Any interior candidate is bounded above by the all-mass cost $c_{m0}(X/e_m)^{1-\beta}$, which is sublinear in X, so the interior branch beats the corner for large X while the corner is global for small X; the switch scale $\hat{X}$ satisfies $\hat{X} < X^c$ because at X^c the two corners tie while a nearby interior mix strictly beats the all-mass corner. At $\hat{X}$ the minimizer jumps from $\omega = 0$ to the interior branch across the intervening local maximum: adoption is discrete. Uniqueness of the switch scale $\hat{X}$, the strict positivity of the jump, monotone comparative statics of the interior share, and the corner conditions are proved in

full in Theorem 1 (Appendix C). For Proposition 10: in the precision regime $C(X) = \pi_q X$, so the symmetric equilibrium ($XC'(X) = RV/4$) is $X = RV/(4\pi_q)$, which reaches $\hat{X}$ at $\bar{V}^{\mathrm{ad}} = 4\pi_q \hat{X}/R$. Above that stake, holding force fixed and switching to the interior arsenal strictly lowers cost, so the precision profile cannot remain an equilibrium; non-convexity can induce the global switch earlier. Along any regular stable branch, implicit differentiation of $F(X, V) = XC'(X) - RV/4 = 0$ gives $dX/dV = (R/4)/F_X > 0$, and Theorem 1 then gives $d\omega^*/dV > 0$. Finally $C(X) \sim c_{m0}(X/e_m)^{1-\beta}$ as $X \to \infty$, so the force–stake elasticity converges to $1/(1-\beta)$. For Proposition 11, the dropout claim is Theorem 2: no symmetric interior pure-strategy equilibrium exists once $\beta\, s_M(X^e) > \frac{1}{2}$, which for $\beta > 1/2$ occurs at all stakes beyond a finite $\bar{V}$, with the pure-mass bound recovered as $s_M \to 1$; the global best-response scan locates the boundary at benchmark parameters (equilibrium at $\beta = 0.30$, none at $\beta = 0.55$). □

Proof of Proposition 12. Backward induction. In period 2 at states (Q_{A1}, Q_{B1}), side i maximizes $Vm_{i2}/(m_{i2} + m_{j2}) - C(Q_{i1}, m_{i2})$, giving the continuation value $J_i(Q_{i1}, Q_{j1})$. In period 1 the payoff is $Vp_1 - C(Q_{i0}, m_{i1}) + \delta J_i(Q_{i0} + m_{i1}, Q_{j0} + m_{j1})$; the displayed first-order condition follows since $\partial C(Q, m)/\partial m = c_0(Q + m)^{-\beta}$. Differentiating J_i and applying the envelope theorem to the own choice m^*_{i2} leaves the direct effect $-\partial C(Q_{i1}, m^*_{i2})/\partial Q_{i1} = c_0[Q_{i1}^{-\beta} - (Q_{i1} + m^*_{i2})^{-\beta}] > 0$ plus the strategic term $\Sigma_i = V \frac{\partial p_2}{\partial m_{j2}} \frac{\partial m^*_{j2}}{\partial Q_{i1}}$. Its sign is not global. At a symmetric state and symmetric terminal equilibrium, however,

$$\frac{\partial}{\partial m_j}\left[\frac{Vm_j}{(m_i + m_j)^2} - c_0(Q_i + m_i)^{-\beta}\right] = \frac{V(m_i - m_j)}{(m_i + m_j)^3} = 0.$$

The terminal first-order-condition Jacobian is therefore diagonal at symmetry, so $\partial m^*_{j2}/\partial Q_{i1} = 0$ and $\Sigma_i = 0$ exactly.

It remains to compare quantities rather than just marginal costs. Let y be symmetric myopic output and x symmetric period-1 output from the dynamic game. Their first-order conditions are

$$\frac{V}{4} = c_0 y(Q_0 + y)^{-\beta}, \qquad \frac{V}{4} = x\{c_0(Q_0 + x)^{-\beta} - \delta D\},$$

where $D = c_0[Q_1^{-\beta} - (Q_1 + m^*_2)^{-\beta}] > 0$. Hence $c_0 x(Q_0 + x)^{-\beta} > c_0 y(Q_0 + y)^{-\beta}$. The function $z \mapsto z(Q_0 + z)^{-\beta}$ is strictly increasing for $\beta < 1$, because its derivative is $(Q_0 + z)^{-\beta}[1 - \beta z/(Q_0 + z)] > 0$. Thus $x > y$. Continuity extends the strict inequality to a neighborhood of symmetry. At $\beta = 0$, J_i does not depend on Q and the wedge is zero. □

Proof of Proposition 13. First solve the terminal subgame. Denote its inherited states by (Q_A, Q_B) and its outputs by (y_A, y_B); let $Z_i = Q_i + y_i$. If $\tilde{r} = Q_A/Q_B > 1$, division of the two terminal first-order conditions gives

$$\frac{y_A}{y_B} = \left(\frac{Z_A}{Z_B}\right)^{\beta}.$$

Writing $r = Z_A/Z_B$, $k = r^\beta$, and $s = y_B/Z_B \in (0, 1)$, so that $\tilde{r}$ is the ratio the subgame inherits and r the ratio after production, gives

$$\tilde{r} = \frac{r - ks}{1 - s}, \qquad 1 < k < r < \tilde{r}. \tag{A.1}$$

The last inequalities follow from the same weighted-average argument as in the one-period model, and establish terminal contraction.

The issue is period 1, whose continuation term invalidates that simple FOC ratio. Implicitly differentiate the two terminal first-order conditions. After scaling by $g = c_0 Z_B^{-\beta}$, define

$$H = r(k+1) - 2\beta ks(r+1) + \beta^2 ks^2(k+1) > 0, \qquad P = \frac{\beta krs(k-1)}{H}.$$

The strict terminal second-order conditions imply $H > 0$. The resulting own-state derivatives of terminal equilibrium values are

$$\frac{J_{A,Q_A}}{g} = (r - ks)^{-\beta} - \frac{1}{k} + \frac{P}{kr}, \qquad \frac{J_{B,Q_B}}{g} = (1-s)^{-\beta} - 1 - \frac{P}{k}. \tag{A.2}$$

Thus the effective period-1 marginal costs $G_i = c_0 Q_i^{-\beta} - \delta J_{i,Q_i}$ satisfy

$$\begin{aligned} \frac{G_A}{g} &= (1-\delta)(r-ks)^{-\beta} + \frac{\delta}{k} - \frac{\delta P}{kr}, \\ \frac{G_B}{g} &= (1-\delta)(1-s)^{-\beta} + \delta + \frac{\delta P}{k}. \end{aligned} \tag{A.3}$$

If $\hat{r} = m_{A1}/m_{B1}$, division of the period-1 FOCs gives $\hat{r} = G_B/G_A$. At $\delta = 0$, this ratio is $\hat{r}(0) = \tilde{r}^{\beta} \in (1, \tilde{r})$; at $\delta = 1$ it is

$$\hat{r}(1) = \frac{k+P}{1 - P/r}.$$

The inequality $\hat{r}(1) < \tilde{r}$ is equivalent to

$$N \equiv (r-k)H - \beta ks(k-1)\{2r - s(r+k)\} > 0. \tag{A.4}$$

I now prove this bound for $\beta \le 1/2$. Holding (r, k, s) fixed,

$$N_{\beta\beta} = 2ks^2(r-k)(k+1) > 0,$$

while, because $k = r^{\beta} \le \sqrt{r}$,

$$N_{\beta}\big|_{\beta=1/2} = 2ks\{k + ks(r-1) - r^2\} < 0.$$

Hence $N(\beta) \ge N(1/2)$. At $\beta = 1/2$, N is convex in s, its derivative at $s = 1$ is non-positive, and

$$N(1/2, 1) = \frac{(r-k)(4r - k^2 - k)}{4} > 0.$$

Indeed, $N_{ss} = k(k^2 + 3kr - 3k - r)/2 \ge 0$ because $r \ge k^2$, and the claimed endpoint derivative follows from $2r^2 + r + k - k^2 - 3kr \ge 0$: it equals $k(k-1)^2(2k+1)$ at $r = k^2$ and is increasing in r. This proves (A.4). It also gives $P < r$, so the denominators are positive and $1 < \hat{r}(1) < \tilde{r}$. For any $\delta \in [0, 1]$, (A.3) makes $\hat{r}$ a positive-denominator-weighted average of $\hat{r}(0)$ and $\hat{r}(1)$; therefore

$$1 < \hat{r} < \tilde{r}. \tag{A.5}$$

Finally, $\tilde{r} = r_1$ is itself the weighted average

$$\tilde{r} = \frac{r_0 Q_{B0} + \hat{r}\, m_{B1}}{Q_{B0} + m_{B1}}.$$

It cannot equal one (then symmetry implies $\hat{r} = 1$, inconsistent with $r_0 > 1$). It cannot be below one either: swapping labels in (A.5) would give $\tilde{r} < \hat{r} < 1$, whereas the displayed weighted average is then strictly above $\tilde{r}$. Hence $\tilde{r} > 1$; (A.5) and the weighted average give $1 < r_1 < r_0$. Applying (A.1) in the terminal subgame yields $1 < r_2 < r_1$, completing the proof. □

Proof of Proposition 16. Period-2 settlements S (to A) are acceptable iff $p_2V - k_A \le S \le p_2V + k_B$; a feasible one exists since war is costly. The declining side B's best feasible peace path takes the whole period-1 flow and the best acceptable period-2 share: $V + \delta \min\{V, (1-p_2)V + k_A\}$. Its war value is $(1-p_1)(1+\delta)V - k_B$. If the transfer cap binds ($k_A > p_2V$), the peace path dominates outright. Otherwise war is strictly preferred iff $(1-p_1)(1+\delta)V - k_B > V + \delta[(1-p_2)V + k_A]$, which rearranges to $\delta p_2 - (1+\delta)p_1 > (k_B + \delta k_A)/V$, i.e. $\delta(p_2 - p_1) > p_1 + (k_B + \delta k_A)/V$. The rising side's constraint is slack for all $p_2 \ge p_1$ (its war value $p_1(1+\delta)V - k_A$ never exceeds its best feasible peace path $V + \delta \min\{V, p_2V + k_B\}$). Part (i): if $\delta \Delta p \le p_1$ the inequality fails at every $k \ge 0$. Part (ii): set $p_2 = 1$, $k_A = k_B = k$ and solve for p_1. □

Proof of Proposition 1. Substitute $n = X\zeta^{-b}$ into B in (4), giving $B = c_{m0}X^{1-\beta}\zeta^{-b(1-\beta)} + wX\zeta^{1-b}$. The first-order condition $\partial B/\partial\zeta = 0$ yields $\zeta^* \propto X^{-\beta/(1-b\beta)}$; hence $n^* = X\zeta^{*-b} \propto X^{1/(1-b\beta)}$. Both terms of B evaluated at ζ^* carry the exponent $(1-\beta)/(1-b\beta)$, so $C(X) \propto X^{(1-\beta)/(1-b\beta)}$, an increasing-returns cost with effective contest decisiveness the reciprocal exponent $(1-b\beta)/(1-\beta)$. Retaining the primitive coefficients in the same first-order condition gives

$$\zeta^* \propto \left(\frac{c_{m0}}{w}\right)^{1/(1-b\beta)} X^{-\beta/(1-b\beta)}, \qquad n^* \propto \left(\frac{w}{c_{m0}}\right)^{b/(1-b\beta)} X^{1/(1-b\beta)},$$

and substitution into either cost term gives the coefficient scaling stated in the proposition. Limits in b follow by inspection. □

Proof of Proposition 14. State A prefers settlement α to war iff $\alpha V \ge pV - k_A$, i.e. $\alpha \ge p - k_A/V$; state B prefers it iff $(1-\alpha)V \ge (1-p)V - k_B$, i.e. $\alpha \le p + k_B/V$. The interval $[\,p - k_A/V,\ p + k_B/V\,]$ has positive length $(k_A + k_B)/V$ and midpoint $p + (k_B - k_A)/2V$. Intersecting with the feasible settlement space $[0,1]$: the point $p \in [0,1]$ lies in both sets, so the intersection is non-empty, with the stated width, capped at 1. □

Proof of Proposition 15. Symmetry gives $B_A^* = B_B^* = B^* = RV/4$ (Lemma 2) and $X_A^* = X_B^* = \psi B^*$. Then $k_i = \tilde{\lambda} B^* + \chi\psi B^*$ and $(k_A + k_B)/V = 2(\tilde{\lambda} + \chi\psi)B^*/V = (R/2)(\tilde{\lambda} + \chi\psi)$, which is increasing in $\psi = 1/P$ (Lemma 1) until the truncation of Proposition 14 binds, beyond which every settlement is preferred to war by both sides. □

B The multiplicative base, result by result

This appendix derives every core result *natively* on the multiplicative baseline of Section 4.2, so the reader can see what the additive two-good formalization adds or hides, and where, if anywhere, the two formulations fork.

Cost of force and contest. With $X = n\zeta^b$ and $B = c_{m0}n^{1-\beta} + wn\zeta$, cost minimization (Proposition 1) gives

$$C(X) = \Upsilon\, X^{\eta}, \qquad \eta \equiv \frac{1-\beta}{1-b\beta} \in (0,1], \qquad \Upsilon = \Lambda(b,\beta)\, c_{m0}^{\frac{1-b}{1-b\beta}}\, w^{\frac{b(1-\beta)}{1-b\beta}},$$

so Υ is Cobb–Douglas in the two input prices with base-cost share $(1-b)/(1-b\beta)$, and the cost of force has increasing returns whenever $\beta > 0$ ($\eta < 1$). Each side chooses X_i to maximize $p_iV - C(X_i)$ under the CSF (8); define *productivity* $\psi_i \equiv \Upsilon_i^{-1/\eta}$.

1. *Win probability (no fork).* First-order conditions give $X_A/X_B = \psi_A/\psi_B$ and hence $p_A = \psi_A^R/(\psi_A^R + \psi_B^R)$, the *same* productivity-share form as the additive Lemma 2, with $\psi = \Upsilon^{-1/\eta}$ in place of $1/P$.

2. *Symmetric arming (level forks).* Symmetric equilibrium spending is $B^* = RV/(4\eta) = \frac{1-b\beta}{1-\beta} \cdot \frac{RV}{4}$: increasing returns raise arming above the additive $RV/4$ by $1/\eta \geq 1$, recovering $RV/4$ as $\eta \to 1$ ($\beta \to 0$).

3. *Paradigm shift (no fork in direction).* A fall in c_{m0} lowers Υ, raises ψ and the win probability, and lowers optimal quality $\zeta^* \propto c_{m0}^{1/(1-b\beta)}$ while raising n^*: "re-tool toward mass" becomes "fewer quality per unit, more units."

4. *Capability, not budgets (no fork).* A symmetric fall in c_{m0} leaves $B^* = RV/(4\eta)$ unchanged (it depends only on b, β) while $\psi \to \infty$: destructiveness per dollar rises without budgets moving (cf. Proposition 3).

5. *Winner-take-all (threshold forks).* The arming game is Tullock with *effective decisiveness* $R/\eta = R(1-b\beta)/(1-\beta)$; a symmetric pure-strategy equilibrium fails when $R/\eta > 2$. For $R = 1$ this is $\beta > 1/(2-b)$, strictly above the additive pure-mass threshold $\beta = 1/2$, which it recovers as $b \to 0$ and which becomes unattainable ($\beta > 1$) as $b \to 1$. The good-enough parameter raises the learning needed for a decisive race.

6. *Stakes and scale (elasticity forks).* Equilibrium force is $X^* \propto V^{1/\eta} = V^{(1-b\beta)/(1-\beta)}$ and quality per unit $\zeta^* \propto V^{-\beta/(1-\beta)}$ falls with the stake, a *continuous* tilt in ζ rather than a discrete majority-mass boundary (Propositions 9–11). The stake elasticity $1/\eta = (1-b\beta)/(1-\beta)$ lies below the additive $1/(1-\beta)$: quality substitution dampens the super-linearity, coinciding as $b \to 0$.

7. *Hard cap, offense–defense, bargaining (no fork).* Under a binding cap the dual gives $X^* \propto \bar{B}^{1/\eta}$, $\zeta^* \propto \bar{B}^{-\beta/(1-\beta)}$ (Remark 2); the offense advantage A_{off} (Section 7) and the bargaining range (Proposition 14) are arsenal-agnostic. The commitment channel (Section 11) invokes the decisive-race threshold and so inherits the $\beta > 1/(2-b)$ fork.

Table 2: The two-good version vs. the multiplicative baseline: where they coincide and where they generalize. "$= \eta \to 1$" marks the constant-returns benchmark ($\beta \to 0$); "$= b \to 0$" marks the multiplicative's pure-mass limit, which reproduces the two-good model's mass-dominated regime.

Result	Additive (two-good)	Multiplicative (baseline)	Relationship
Win probability	$\psi_A^R/\sum$, $\psi = 1/P$	$\psi_A^R/\sum$, $\psi = \Upsilon^{-1/\eta}$	same form
Symmetric spend	$RV/4$	$RV/(4\eta)$	additive $= \eta \to 1$
Winner-take-all (R=1)	$\beta = 1/2$	$\beta = 1/(2-b)$	additive $= b \to 0$
Force–stake elasticity	$1/(1-\beta)$	$(1-b\beta)/(1-\beta)$	additive $= b \to 0$
Tipping to mass	discrete mass-share / phase	continuous $\zeta^*(X)$	structural
Shift, capability, O–D, bargaining	—	—	no fork

Which version for the baseline? The two versions agree wherever they overlap. On the constant-returns benchmark (no learning, $\beta \to 0$, $\eta \to 1$) both give the productivity-share contest with symmetric spend $RV/4$; and the multiplicative's pure-mass limit ($b \to 0$) reproduces the two-good model's mass-dominated regime, the same winner-take-all threshold $\beta = 1/2$ and force–stake elasticity $1/(1-\beta)$. They differ through the "good-enough" parameter b: the multiplicative's

sharpest results, $\beta = 1/(2-b)$ and elasticity $(1-b\beta)/(1-\beta)$, depend on how substitutable quality is, which the fixed two-grade structure holds constant. Because that dependence is economically meaningful, quality substitutability governs how decisive the mass race is, and because a single weapon of chosen number and quality is the more faithful description of the paradigm, the multiplicative is the baseline. The two-good version earns its place for calibration: its price index $P(\pi_m, \pi_q)$ is closed-form, and its two priced grades map onto the distinct systems the cost-exchange data report.

C Global analysis of the learning-CES arsenal: analytical proofs

This appendix establishes the global arsenal structure and dropout boundary analytically. Throughout, write the conditional cost of delivering force X at mass contribution share $\omega \in [0,1]$ (equation 13) as

$$g(\omega) \;=\; A\,\omega^{u} + B\,(1-\omega)^{v}, \qquad A = c_{m0}(X/e_m)^{1-\beta}, \quad B = \pi_q X, \quad u = \frac{1-\beta}{\rho}, \quad v = \frac{1}{\rho} > 1,$$

with $\rho = 1 - 1/\sigma \in (0,1)$. Note $u > 1 \iff \beta < 1/\sigma$. Define $\tau \equiv (A/B)$, which is proportional to $X^{-\beta}$ and strictly decreasing in X; minimizing g is minimizing $f(\omega;\tau) = \tau\omega^u + (1-\omega)^v$.

Theorem 1 (Global arsenal structure)**.** *Let $\sigma > 1$, $\beta \in (0,1)$.*

(i) If $\beta < 1/\sigma$, then f has a unique critical point $\omega^*(\tau) \in (0,1)$, it is the global minimum at every scale, and ω^* is strictly decreasing in τ, hence strictly increasing in X, with $\omega^* \to 0$ as $X \to 0$ and $\omega^* \to 1$ as $X \to \infty$.

(ii) If $\beta > 1/\sigma$, then $\omega = 0$ is a local minimum, f has at most two interior critical points (a local maximum, then a local minimum $\omega_2(\tau)$), and there is a unique *threshold $\hat{X} \in (0, X^c)$ such that the global minimum is the corner $\omega = 0$ for $X < \hat{X}$ and the interior point ω_2 for $X > \hat{X}$; at $\hat{X}$ the optimal share jumps by a strictly positive amount, and ω_2 is strictly increasing in X thereafter. Specifically, the jump share $\omega^\dagger$ is the unique positive root*

$$\ell(\omega) \equiv \log\Big[1 + \Big(\frac{v}{u} - 1\Big)\,\omega\Big] + (v-1)\log(1-\omega) = 0,$$

and, with $X^c = e_m(c_{m0}/(\pi_q e_m))^{1/\beta}$ the all-mass/all-precision average-cost crossover,

$$\tau^\dagger = \frac{v}{u}(\omega^\dagger)^{1-u}(1-\omega^\dagger)^{v-1}, \qquad \hat{X} = X^c/(\tau^\dagger)^{1/\beta} < X^c.$$

(iii) If $\beta = 1/\sigma$, then $u = 1$ and $f(\omega) = \tau\omega + (1-\omega)^v$ is convex. The optimum is $\omega^* = 0$ for $\tau \geq v$; for $\tau < v$ it is interior with $1 - \omega^* = (\tau/v)^{1/(v-1)}$. Adoption enters continuously at $\tau = v$.

Proof. Write the first-order condition $f_\omega = 0$ in log form: $h(\omega) = \ln[v/(u\tau)]$, where $h(\omega) \equiv (u-1)\ln\omega - (v-1)\ln(1-\omega)$.

(i) $u > 1$: $f_\omega(0^+) = -v < 0$ and $f_\omega(1^-) = \tau u > 0$, so a minimum is interior. The function h is strictly increasing from $-\infty$ to $+\infty$, so the critical point is unique and therefore the global minimum. Monotone comparative statics give $f_{\omega\tau} = u\omega^{u-1} > 0$ and $f_{\omega\omega} > 0$ at the minimum, hence $d\omega^*/d\tau = -f_{\omega\tau}/f_{\omega\omega} < 0$. The FOC gives $(1-\omega^*)^{v-1}/{\omega^*}^{u-1} = \tau u/v$, whose left side is a homeomorphism of $(0,1)$ onto $(0,\infty)$. Thus $\omega^* \to 0$ as $\tau \to \infty$ ($X \to 0$) and $\omega^* \to 1$ as $\tau \to 0$ ($X \to \infty$).

(ii) $u < 1$: now $f_\omega(0^+) = +\infty$, so $\omega = 0$ is a local minimum. Since $h'(\omega) = (u-1)/\omega + (v-1)/(1-\omega)$ is strictly increasing, h is strictly quasiconvex with $h \to +\infty$ at both endpoints. A level

set of h therefore has at most two points, giving at most two interior critical points: the smaller is a local maximum and the larger, ω_2, a local minimum.

For the threshold, impose stationarity and equality with the corner $f(0) = 1$. Eliminating τ gives

$$(1-\omega)^{v-1}\left[1+\left(\frac{v}{u}-1\right)\omega\right]=1,$$

or $\ell(\omega) = 0$. Now $\ell(0) = 0$, $\ell'(0) = v(1/u - 1) > 0$, $\ell''(\omega) < 0$, and $\ell(\omega) \to -\infty$ as $\omega \to 1$. Hence there is exactly one positive root $\omega^\dagger$. The maximum of ℓ occurs at $\omega_h = (1-u)/(v-u)$, also the minimum of h; therefore $\omega^\dagger > \omega_h$ lies on the local-minimum branch. Substitution into the stationary condition gives $\tau^\dagger$ and $\hat{X} = X^c/(\tau^\dagger)^{1/\beta}$.

As τ falls (equivalently, as X rises), the larger stationary root ω_2 rises strictly by the implicit-function theorem, so it crosses $\omega^\dagger$ exactly once. Before that crossing its value exceeds the corner; after it, it is below the corner. This proves the unique global switch and the strictly positive jump. At $\tau = 1$ the all-mass and all-precision corners tie, while $\omega = 1$ is not a local minimum; hence an interior point strictly beats both, so $\tau^\dagger > 1$ and $\hat{X} < X^c$. The formula also gives the unambiguous primitive comparative statics: cheaper or more effective mass, or more expensive precision, lowers $\hat{X}$. The effect of β itself is generally ambiguous because it moves u, X^c, $\tau^\dagger$, and the outer exponent simultaneously.

(iii) If $u = 1$, then $f_{\omega\omega} = v(v-1)(1-\omega)^{v-2} > 0$. The derivative at the precision corner is $f_\omega(0) = \tau - v$, so $\omega^* = 0$ when $\tau \geq v$. When $\tau < v$, the unique interior FOC is $\tau = v(1-\omega)^{v-1}$, which gives the stated solution and converges to the corner as $\tau \uparrow v$. □

Lemma 4 (Envelope elasticity of conditional cost)**.** *At an interior optimum, the elasticity of $C(X) = g(\omega^*(X))$ is*

$$\varepsilon_C(X) \;\equiv\; \frac{X\,C'(X)}{C(X)} \;=\; 1-\beta\,s_M(X),$$

where s_M is the mass share of cost, and $1 - s_M(X) = O\big(X^{\beta(1-\sigma)}\big) \to 0$ as $X \to \infty$.

Proof. By the envelope theorem, $C'(X) = \partial g/\partial X$ holding ω^* fixed. The mass component $M = A\omega^u$ has X-elasticity $1-\beta$ and the precision component $P = B(1-\omega)^v$ has X-elasticity 1, so $XC' = (1-\beta)M + P$ and $\varepsilon_C = 1 - \beta M/(M+P)$. For the rate: the FOC gives $(1-\omega^*)^{v-1} \propto \tau{\omega^*}^{u-1}$, so as $\omega^* \to 1$, $(1-\omega^*) \propto \tau^{1/(v-1)}$ and $P/M \propto X^\beta\,(1-\omega^*)^v \propto X^\beta X^{-\beta v/(v-1)} = X^{\beta(1-\sigma)}$, using $v/(v-1) = \sigma$. □

Theorem 2 (The dropout boundary, analytically)**.** *Consider the symmetric arming contest with $R = 1$ and cost $C(X)$. Any symmetric interior pure-strategy equilibrium satisfies $XC'(X) = V/4$ and yields payoff $V/2 - C(X^e) = C(X^e)\big[2\varepsilon_C(X^e) - 1\big]$; dropping out yields 0. The dropout deviation is therefore profitable if and only if $\varepsilon_C(X^e) < \frac{1}{2}$, i.e.*

$$\beta\,s_M(X^e) \;>\; \tfrac{1}{2}.$$

Because $s_M(X) \to 1$ (Lemma 4) and $X^e(V) \to \infty$ as $V \to \infty$, for every $\beta > \frac{1}{2}$ there exists $\bar{V}$ such that no symmetric pure-strategy equilibrium exists for $V > \bar{V}$. The onset is $\bar{V} = \inf\{V : \beta s_M(X^e(V)) > 1/2\}$; when the share is continuous there, its boundary satisfies equality, while a discrete adoption jump may cross the inequality. Conversely, for $\beta \leq \frac{1}{2}$, the dropout deviation is never profitable. This converse rules out dropout, not other possible global deviations, and therefore does not by itself prove equilibrium existence below the boundary.

Proof. The FOC and payoff are standard. Substituting $V = 4X^e C'(X^e)$ into the candidate payoff gives $V/2 - C = 2X^e C'(X^e) - C(X^e) = C(X^e)[2\varepsilon_C(X^e) - 1]$, negative iff $\varepsilon_C < 1/2$, in which case the zero deviation (payoff 0, since $p_i(0, X_j) = 0$ and $C(0) = 0$) strictly improves, ruling out the candidate; the symmetric corner $(0, 0)$ is never an equilibrium since an arbitrarily small positive production wins the whole prize. $X^e(V) \to \infty$ because $XC'(X) = \varepsilon_C C \leq C$ is bounded on bounded sets while $V/4 \to \infty$. By Lemma 4, $\varepsilon_C(X) = 1 - \beta s_M(X)$ with s_M increasing to 1 along the interior branch, which delivers the high-stake non-existence result and the no-dropout converse. □

Theorem 2 establishes the mixed-arsenal dropout condition directly: $\beta s_M > \frac{1}{2}$, non-existence there is proved, and the pure-mass threshold $\beta = \frac{1}{2}$ is recovered exactly as $s_M \to 1$. Below that boundary, the global scans find an equilibrium at the stated benchmark; the theorem establishes non-existence above it.

D Which regime? Evidence from the Russian long-range UAV attack stream

Proposition 8 names three regimes and claims a sign in only two. When a delivered-effect target binds, expenditure is pinned and replacement generates experience. When an inventory requirement binds instead, the $\bar{F}$ terms carry the result. When neither binds, production is a free choice and *no sign is claimed*. Which regime a campaign occupies is therefore an empirical question prior to any estimate, and this appendix answers it for the largest one-way attack campaign on record.

The attack stream rejects the simplest empirical implementation of the first regime: a fixed or smoothly trending delivered-effect target. It therefore lies outside the domain in which the signed countermeasure mechanism can be applied directly. The same data leave a time-varying operational target, a capacity constraint, inventory management and targeting doctrine observationally equivalent.

Data. Ukrainian Air Force daily reports, compiled by Ivaniuk (2026) and cross-checked against the CSIS Russian Firepower Strike Tracker, aggregated to 46 monthly observations from October 2022 to July 2026: **110,882** launches in the Shahed-type stream and **26,469** survivors of air defense. The snapshot is archived with its source, retrieval date and checksum, and revisions to already-published rows are tracked.

Measurement and coverage. Three features of the raw file constrain its use.

First, the series measures the *attack stream*. The `launched` field counts everything reported under the Shahed-type heading, including decoys. Where composition is reported, positively identified Shaheds are **21%** of the stream in 2025 and **35%** in 2026.

Second, launches minus interceptions measures *post-interception survival*. A further `not_reach_goal` category (12,948 against 26,469 survivors, i.e. **49%**) failed short of the aimpoint. Netting it out requires the field to be reported for every episode in a month and to be internally consistent with the other two; only **8 of 46** months satisfy both. The tests below therefore use survivors as an upper bound on delivered effect.

Third, the `not_reach_goal` field changes meaning mid-sample and is excluded from all series work here. Through June 2025 every episode carries a positive value; from August 2025 essentially every episode carries zero, while launches are rising. Failures did not fall from roughly a thousand a month to none. The mechanism is visible in how the data are made: the compiler extracts

figures manually from Air Force social-media posts, so the field records *what the Air Force chose to report*, namely the "locationally lost" category, covering drones that vanish from radar to jamming, poor build quality or fuel exhaustion, rather than a measured quantity. When the phrasing left the reports, the field went to zero. Only four months survive as usable, all in early 2025, which supports a level comparison rather than a time-series comparison.

It does support one *level* statement, and a consequential one. Across those four months, **83%** of the drones that survived interception still failed to reach their target (4,809 of 5,816; monthly range 70–94%). Delivered effect was therefore around a sixth of survivors, so cost per unit of delivered effect is roughly six times cost per survivor. Cost-exchange ratios computed from launch or leakage counts, as they usually are, understate the attacker's true cost per unit of effect by about that factor, in the one window where the wedge is observed.

Test. If a constant delivered-effect target binds, the attacker must fly $y_t = \bar{D}_t/s_t$, so the elasticity of launches with respect to survival is -1. Regressing $\ln y_t$ on $\ln s_t$ with a linear trend gives $+\mathbf{0.112}$ (Newey–West standard error 0.090); in first differences, $-\mathbf{0.057}$ (0.030). With three Newey–West lags, the t statistics for the null of unit negative elasticity are 12.4 and 31.4. The constant-target prediction is rejected in both specifications. A second diagnostic points the same way: if the target bound, delivered effect would be more stable than its components, whereas survivors are the *most* variable series (coefficient of variation 1.23, against 1.02 for launches and 0.73 for survival). Whether the same holds for delivered effect proper cannot be established here: the months in which it is observed straddle the reporting break, so a comparison across them measures that break rather than the underlying quantity.

Reading. The estimated response is far from the one-for-one adjustment implied by a fixed delivered-effect target. The last months of the sample make the point plainly: launches range from 4,985 to 8,207 per month while survivors sit in a narrow 560–701 band and interception holds at 87–92%. A side holding delivered effect to a target would have raised launches one-for-one as survival fell. The observed stream instead combines a large increase in launches with a narrow band of survivors.

This boundary result matters for Proposition 8: the sign in part (v) requires a binding delivered-effect target, whereas this important campaign rejects its fixed-target implementation. Distinguishing a time-varying target from capacity, inventory and doctrine requires production and stock accounting alongside operational plans. The available data establish the boundary of the corollary's application rather than selecting among those alternative mechanisms.

A second use: bounding the survival component of $\mathcal{E}$. The same source answers a different question the calibration needs. The precision-effectiveness advantage $\mathcal{E} = e_q/e_m$ of (6) is a parameter of *one side's* arsenal, so the identifying comparison is the same attacker firing a cheap and an exquisite system into the same defended airspace over the same period. This campaign is exactly that: alongside the Shahed stream, the reports break out Kh-101/555 and Kalibr cruise missiles, Iskander-K, Iskander-M and Kinzhal. Restricting to systems that deliver a warhead to a fixed target at range, and excluding air defense, reconnaissance and loitering munitions as different missions, gives 1,735 episodes over the same window (Table 3).

Trajectory organizes survival more clearly than the price ladder. Pooled, subsonic cruise missiles costing millions leak through *less* often than a \$20–50 k one-way attack drone. That ordering held in 2023, 2024 and 2025, with cruise-to-mass survival ratios of 0.84, 0.61 and 0.65; it reverses in 2026, where the ratio is 1.09. Ballistic systems, by contrast, survive three to four times better throughout,

Table 3: Survival by weapon class, Russian long-range campaign against Ukraine, September 2022–July 2026. $s = 1 -$ destroyed/launched; `destroyed` is the defender's claim. The Shahed total is nine below the attack-stream total because one nine-launch episode lacks a destroyed count and is excluded here.

	System	Launched	Survival s	s/s_{mass}
exquisite cruise	Kh-101/555 ALCM	2,170	0.135	0.57
	Kalibr SLCM	581	0.219	0.92
	Iskander-K GLCM	226	0.350	1.46
cheap mass	Shahed-136/131 one-way attack	110,873	0.239	1.00
exquisite ballistic	Kinzhal ALBM	233	0.725	3.04
	Iskander-M SRBM	581	0.850	3.56

and their advantage widens to 8.1× by 2026. A changing stock of ballistic-capable interceptors, salvo composition, target selection and defender reporting remain observationally equivalent in this series. The comparisons are descriptive associations.

The bearing on the calibration is direct, and the two classes must be kept apart because they face different interceptors and different target sets. Pooled within class, the survival ratio is 0.70 against cruise and 3.41 against ballistic. Since the deep-strike arsenal tips to mass unless $\mathcal{E}$ exceeds the price gap of about 37, accuracy and warhead effect would have to supply the remaining factor on their own: roughly 53× against cruise, or 11× against ballistic. This bounds one component of $\mathcal{E}$. The defender supplies the `destroyed` claims; the missile classes lack a counterpart to the 83% delivered-effect wedge above; accuracy and warhead effect are unmeasured; the classes face different target sets and interceptors; and the attacker chooses the launch mix. Within that scope, the survival channel is small: an exquisite deep-strike munition's case for a thirtyfold price premium must come mainly from accuracy and warhead effect rather than arrival probability.

Counts and cost shares are distinct. This campaign is overwhelmingly cheap mass *by count*: Shahed-type rounds are 110,873 of 114,664 launches in the six classes, or 96.7%. A count-dominated arsenal can nonetheless have a low mass share of *spending*, because a large price gap drives the two measures apart mechanically. At any plausible price for these systems the same campaign is a large majority of rounds and a small minority of expenditure. Observed round counts therefore discipline quantities rather than the cost shares in the model.

Whether the observed mix is what cost minimization would have chosen is a further question, and a harder one: it requires prices for both poles within a single arsenal, a treatment of cruise and ballistic systems that does not pool target classes they do not share, and a position on what was binding on production. That exercise is left to separate work.

Scope of inference. Measuring replacement-driven learning and its elasticity requires production and inventory data; estimating Geran unit cost requires a separate cost source. Recovering $\mathcal{E}$ from an observed arsenal share additionally requires prices for both poles within one arsenal and evidence that the observed mix was freely chosen. For the Russian arsenal, the available price estimates are intelligence assessments with wide dispersion. The saturation regressions are likewise inconclusive: regressing $\ln s_t$ on $\ln y_t$ gives $+1.25$ in levels with a trend but -0.52 in first differences, a disagreement which indicates that the levels estimate is picking up the low-frequency coincidence

of the 2024–25 scale-up with a temporarily overwhelmed defense rather than any within-month salvo-size effect. A direct test of the mechanism calls for linked launch, composition, production and inventory records.

References

Anton, James J. and Dennis A. Yao (1987) "Second Sourcing and the Experience Curve: Price Competition in Defense Procurement," *RAND Journal of Economics*, 18 (1), 57–76.

Armed Conflict Location & Event Data Project (ACLED) (2026) "ACLED Aggregated Conflict-Event Data, Europe and Central Asia," ACLED Europe and Central Asia weekly admin-1 aggregate, `https://acleddata.com`.

Arrow, Kenneth J. (1962) "The Economic Implications of Learning by Doing," *Review of Economic Studies*, 29 (3), 155–173, `10.2307/2295952`.

Batmanghelidj, Esfandyar (2026) "Drones Like Bicycles," Phenomenal World, March, `https://phenomenalworld.org/analysis/cost-of-a-shahed/`.

Baye, Michael R., Dan Kovenock, and Casper G. de Vries (1994) "The Solution to the Tullock Rent-Seeking Game When R > 2: Mixed-Strategy Equilibria and Mean Dissipation Rates," *Public Choice*, 81 (3–4), 363–380, `10.1007/BF01053238`.

——— (1996) "The All-Pay Auction with Complete Information," *Economic Theory*, 8 (2), 291–305, `10.1007/BF01211819`.

Becker, Gary S. (1960) "An Economic Analysis of Fertility," in *Demographic and Economic Change in Developed Countries*, 209–240: Columbia University Press (NBER).

Becker, Gary S. and H. Gregg Lewis (1973) "On the Interaction between the Quantity and Quality of Children," *Journal of Political Economy*, 81 (2, Part 2), S279–S288, `10.1086/260166`.

Becker, Gary S. and Nigel Tomes (1976) "Child Endowments and the Quantity and Quality of Children," *Journal of Political Economy*, 84 (4, Part 2), S143–S162, `10.1086/260536`.

Blanken, Leo J. and Jason J. Lepore (2011) "Slowing Down to Keep the Lead in Military Technology," *Defence and Peace Economics*, 22 (3), 317–334, `10.1080/10242694.2010.491675`.

Brave1 Defense Tech Cluster (2026) "Defense Tech Cluster Brave1," Ministry of Digital Transformation of Ukraine, `https://brave1.gov.ua/en/about`.

Bresnahan, Timothy F. and Manuel Trajtenberg (1995) "General Purpose Technologies: 'Engines of Growth'?" *Journal of Econometrics*, 65 (1), 83–108, `10.1016/0304-4076(94)01598-T`.

Brito, Dagobert L. and Michael D. Intriligator (1995) "Arms Races and Proliferation," in Keith Hartley and Todd Sandler eds. *Handbook of Defense Economics, Vol. 1*, 109–164: Elsevier, `10.1016/S1574-0013(05)80008-0`.

Cabral, Luís M. B. and Michael H. Riordan (1994) "The Learning Curve, Market Dominance, and Predatory Pricing," *Econometrica*, 62 (5), 1115–1140, `10.2307/2951509`.

Cancian, Mark F. and Chris H. Park (2026) "Last Rounds? Status of Key Munitions at the Iran War Ceasefire," CSIS Analysis, April, https://www.csis.org/analysis/last-rounds-status-key-munitions-iran-war-ceasefire.

Center for Strategic and International Studies (2025) "Calculating the Cost-Effectiveness of Russia's Drone Strikes," CSIS Analysis, https://www.csis.org/analysis/calculating-cost-effectiveness-russias-drone-strikes.

Chang, Yang-Ming and Zijun Luo (2017) "Endogenous Destruction in Conflict: Theory and Extensions," *Economic Inquiry*, 55 (1), 479–500, 10.1111/ecin.12354.

Che, Yeon-Koo and Ian L. Gale (1998) "Caps on Political Lobbying," *American Economic Review*, 88 (3), 643–651.

Christensen, Clayton M. (1997) *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*, Harvard Business School Press.

Cohen, Wesley M. and Daniel A. Levinthal (1990) "Absorptive Capacity: A New Perspective on Learning and Innovation," *Administrative Science Quarterly*, 35 (1), 128–152, 10.2307/2393553.

Cornes, Richard and Roger Hartley (2005) "Asymmetric Contests with General Technologies," *Economic Theory*, 26 (4), 923–946, 10.1007/s00199-004-0563-8.

Ding, Jeffrey and Allan Dafoe (2023) "Engines of Power: Electricity, AI, and General-Purpose Military Transformations," *European Journal of International Security*, 8 (3), 377–394, 10.1017/eis.2023.1.

Dixit, Avinash (1987) "Strategic Behavior in Contests," *American Economic Review*, 77 (5), 891–898.

European Commission (2025) "Defence Readiness Roadmap 2030," Directorate-General for Defence Industry and Space, October, https://defence-industry-space.ec.europa.eu/eu-defence-industry/readiness-roadmap-2030_en.

——— (2026) "Action Plan on Drone and Counter-Drone Security," Directorate-General for Defence Industry and Space, February, https://defence-industry-space.ec.europa.eu/commission-publishes-action-plan-drone-and-counter-drone-security-2026-02-11_en.

European Long-Range Strike Approach (2026) "Joint Communiqué by the Defence Ministries of ELSA Nations (One-Way Effector 500 Plus)," Letter of intent, February, https://www.gov.pl/web/national-defence/joint-communiqu-by-the-defence-ministries-of-elsa-nations.

Farmer, J. Doyne and François Lafond (2016) "How Predictable Is Technological Progress?" *Research Policy*, 45 (3), 647–665, 10.1016/j.respol.2015.11.001.

Fearon, James D. (1995) "Rationalist Explanations for War," *International Organization*, 49 (3), 379–414, 10.1017/S0020818300033324.

Fudenberg, Drew and Jean Tirole (1983) "Learning-by-Doing and Market Performance," *Bell Journal of Economics*, 14 (2), 522–530.

Fuhrmann, Matthew and Michael C. Horowitz (2017) "Droning On: Explaining the Proliferation of Unmanned Aerial Vehicles," *International Organization*, 71 (2), 397–418, 10.1017/S0020818317000121.

Garfinkel, Ben and Allan Dafoe (2019) "How Does the Offense-Defense Balance Scale?" *Journal of Strategic Studies*, 42 (6), 736–763, 10.1080/01402390.2019.1631810.

Garfinkel, Michelle R. and Stergios Skaperdas (2007) "Economics of Conflict: An Overview," in Todd Sandler and Keith Hartley eds. *Handbook of Defense Economics, Vol. 2*, 649–709: Elsevier, 10.1016/S1574-0013(06)02022-9.

Gilli, Andrea and Mauro Gilli (2016) "The Diffusion of Drone Warfare? Industrial, Organizational, and Infrastructural Constraints," *Security Studies*, 25 (1), 50–84, 10.1080/09636412.2016.1134189.

Glaser, Charles L. and Chaim Kaufmann (1998) "What Is the Offense-Defense Balance and Can We Measure It?" *International Security*, 22 (4), 44–82, 10.1162/isec.22.4.44.

Hall, Abigail R. and Christopher J. Coyne (2014) "The Political Economy of Drones," *Defence and Peace Economics*, 25 (5), 445–460, 10.1080/10242694.2013.833369.

Hartley, Keith and Todd Sandler eds. (1995) *Handbook of Defense Economics, Vol. 1*, Elsevier.

Hirshleifer, Jack (1989) "Conflict and Rent-Seeking Success Functions: Ratio vs. Difference Models of Relative Success," *Public Choice*, 63 (2), 101–112, 10.1007/BF00153394.

——— (1991a) "The Paradox of Power," *Economics and Politics*, 3 (3), 177–200, 10.1111/j.1468-0343.1991.tb00046.x.

——— (1991b) "The Technology of Conflict as an Economic Activity," *American Economic Review*, 81 (2), 130–134.

Horowitz, Michael C. (2010) *The Diffusion of Military Power: Causes and Consequences for International Politics*, Princeton University Press, 10.1515/9781400835102.

Hwang, Sung-Ha (2012) "Technology of Military Conflict, Military Spending, and War," *Journal of Public Economics*, 96 (1–2), 226–236, 10.1016/j.jpubeco.2011.09.003.

Ivaniuk, Petro (2026) "Massive Missile Attacks on Ukraine," Kaggle dataset piterfm/massive-missile-attacks-on-ukraine.

Jervis, Robert (1978) "Cooperation Under the Security Dilemma," *World Politics*, 30 (2), 167–214, 10.2307/2009958.

Jia, Hao (2008) "A Stochastic Derivation of the Ratio Form of Contest Success Functions," *Public Choice*, 135 (3–4), 125–130, 10.1007/s11127-007-9242-1.

Jia, Hao, Stergios Skaperdas, and Samarth Vaidya (2013) "Contest Functions: Theoretical Foundations and Issues in Estimation," *International Journal of Industrial Organization*, 31 (3), 211–222, 10.1016/j.ijindorg.2012.06.007.

Kremer, Michael (1993) "The O-Ring Theory of Economic Development," *Quarterly Journal of Economics*, 108 (3), 551–575, 10.2307/2118400.

Kress, Moshe (2020) "Lanchester Models for Irregular Warfare," *Mathematics*, 8 (5), 737, 10.3390/math8050737.

KSE Institute (2024) "Ukraine's Drones Industry: Investments and Product Innovations," technical report, Kyiv School of Economics Institute (with Brave1), https://kse.ua/wp-content/uploads/2024/10/241004-Brave1-report-v.1.pdf.

Lanchester, Frederick W. (1916) *Aircraft in Warfare: The Dawn of the Fourth Arm*, Constable and Company.

Leininger, Wolfgang (1993) "More Efficient Rent-Seeking: A Münchhausen Solution," *Public Choice*, 75 (1), 43–62, 10.1007/BF01053880.

Ministry of Defence of Ukraine (2025) "Over 3 Million Strike Drones: Key Ministry of Defence Highlights of 2025," Ministry of Defence of Ukraine, press release, December, https://mod.gov.ua/en/news/45-billion-from-partners-over-3-million-strike-drones-more-ukrainian-weapons-key-ministry-of-defence-highlights.

Morgan, John (2003) "Sequential Contests," *Public Choice*, 116 (1–2), 1–18, 10.1023/A:1024252514806.

Nagy, Béla, J. Doyne Farmer, Quan M. Bui, and Jessika E. Trancik (2013) "Statistical Basis for Predicting Technological Progress," *PLOS ONE*, 8 (2), e52669, 10.1371/journal.pone.0052669.

Nelson, Richard R. (1961) "Uncertainty, Learning, and the Economics of Parallel Research and Development Efforts," *The Review of Economics and Statistics*, 43 (4), 351–364, 10.2307/1927082.

Nelson, Richard R. and Sidney G. Winter (1982) *An Evolutionary Theory of Economic Change*, Cambridge, MA: Harvard University Press.

O'Brien, Phillips P. (2024) "AI and Victory in the Ukraine War," Phillips's Newsletter, January, https://phillipspobrien.substack.com/p/ai-and-victory-in-the-ukraine-war.

——— (2026a) "AI And War Is Being Oversold Right Now," Phillips's Newsletter, August, https://phillipspobrien.substack.com/p/ai-and-war-is-being-oversold-right-now.

——— (2026b) "We Are On A Twin-Track Road To An AI Battlefield," Phillips's Newsletter, June, https://phillipspobrien.substack.com/p/we-are-on-a-twin-track-road-to-an.

Parker, Geoffrey (1988) *The Military Revolution: Military Innovation and the Rise of the West, 1500–1800*, Cambridge University Press.

Pérez-Castrillo, J. David and Thierry Verdier (1992) "A General Analysis of Rent-Seeking Games," *Public Choice*, 73 (3), 335–350, 10.1007/BF00140926.

Plichta, Marcel (2025) "Precise Mass in Action: Assessing Ukraine's One-Way Attack Drone Campaign," *The RUSI Journal*, 170 (4), 42–48, 10.1080/03071847.2025.2527923.

Plichta, Marcel and Ash Rossiter (2024) "A One-Way Attack Drone Revolution? Affordable Mass Precision in Modern Conflict," *Journal of Strategic Studies*, 47 (6–7), 1001–1031, 10.1080/01402390.2024.2385843.

Powell, Robert (2006) "War as a Commitment Problem," *International Organization*, 60 (1), 169–203, 10.1017/S0020818306060061.

Sah, Raaj Kumar and Joseph E. Stiglitz (1986) "The Architecture of Economic Systems: Hierarchies and Polyarchies," *American Economic Review*, 76 (4), 716–727.

Sayler, Kelley M. (2026) "DOD Replicator Initiative: Background and Issues for Congress," Congressional Research Service, In Focus IF12611, January, https://www.congress.gov/crs-product/IF12611.

Simon, Herbert A. (1955) "A Behavioral Model of Rational Choice," *Quarterly Journal of Economics*, 69 (1), 99–118, 10.2307/1884852.

Skaperdas, Stergios (1996) "Contest Success Functions," *Economic Theory*, 7 (2), 283–290, 10.1007/BF01213906.

Spence, A. Michael (1981) "The Learning Curve and Competition," *Bell Journal of Economics*, 12 (1), 49–70, 10.2307/3003508.

Thompson, Peter (2012) "The Relationship Between Unit Cost and Cumulative Quantity and the Evidence for Organizational Learning-by-Doing," *Journal of Economic Perspectives*, 26 (3), 203–224, 10.1257/jep.26.3.203.

Tilly, Charles (1990) *Coercion, Capital, and European States, AD 990–1990*, Blackwell.

Tollast, Robert (2025) "Drones: Decoupling Supply Chains from China," rusi research paper, Royal United Services Institute, https://static.rusi.org/rp-drone-supply-chains-china-nov-2025_0.pdf.

Tullock, Gordon (1980) "Efficient Rent Seeking," in James M. Buchanan, Robert D. Tollison, and Gordon Tullock eds. *Toward a Theory of the Rent-Seeking Society*, 97–112: Texas A&M University Press.

UK Ministry of Defence (2026) "UK to Develop New Deep Strike Ballistic Missile for Ukraine (Project Nightfall)," Press release, January, https://www.gov.uk/government/news/uk-to-develop-new-deep-strike-ballistic-missile-for-ukraine.

Watling, Jack and Nick Reynolds (2025) "Tactical Developments During the Third Year of the Russo–Ukrainian War," RUSI special report, Royal United Services Institute, https://www.rusi.org/explore-our-research/publications/special-resources/tactical-developments-during-third-year-russo-ukrainian-war.

Way, Rupert, Matthew C. Ives, Penny Mealy, and J. Doyne Farmer (2022) "Empirically Grounded Technology Forecasts and the Energy Transition," *Joule*, 6 (9), 2057–2082, 10.1016/j.joule.2022.08.009.

Weitzman, Martin L. (1979) "Optimal Search for the Best Alternative," *Econometrica*, 47 (3), 641–654, 10.2307/1910412.

——— (1998) "Recombinant Growth," *The Quarterly Journal of Economics*, 113 (2), 331–360, 10.1162/003355398555595.

Wright, Theodore P. (1936) "Factors Affecting the Cost of Airplanes," *Journal of the Aeronautical Sciences*, 3 (4), 122–128, 10.2514/8.155.