Making Peace on the Cheap

Robert J. Carroll*

March 23, 2026

Abstract

How should a donor allocate commodity transfers across goods and over time to end an ongoing conflict at minimum cost? I embed the donor's problem in a general equilibrium model of a conflict economy, where commodity transfers affect militarization through endogenous price adjustment—the opportunity cost and rapacity channels identified empirically by Dube and Vargas (2013). When the conflict operates through relative prices, transfers at constant prices have zero effect on militarization: the general equilibrium price response is not merely cheaper but necessary. The optimal schedule exploits this mechanism to achieve peace strictly more cheaply than any policy that must eliminate war as an equilibrium outcome, because the donor need only steer the economy to one peaceful equilibrium rather than ruling out all war equilibria. An optimal schedule always exists, without convexity assumptions or uniqueness of equilibrium. Under an absorptive capacity constraint, optimal aid is always disbursed at the maximum rate; with strictly convex delivery costs, three regimes emerge—pause, graduate, or front-load—determined by an endogenous urgency measure. With n donors, peace is achieved n times more slowly than under cooperation, without any change in total resources: a coordination failure that provides a new, resource-neutral rationale for multilateral peacemaking institutions.

*Department of Political Science, University of Illinois at Urbana-Champaign. Email: rjc@illinois.edu.

1 Introduction

When the international community ships grain to a war economy, local grain prices fall. The fall in grain prices raises real wages for subsistence laborers, increasing the opportunity cost of soldiering. If the donor instead ships fuel, fuel prices fall, lowering the cost of military logistics, and fighting may intensify. The effect of a commodity transfer on conflict runs through equilibrium prices—and its sign depends on which commodity is transferred, in what quantity, and in what sequence. The empirical evidence confirms the mechanism: de Ree and Nillesen (2009) find that a ten-percent increase in foreign aid reduces the continuation probability of an ongoing conflict by roughly eight percentage points; Doyle and Sambanis (2000) document that net current transfers per capita are among the strongest predictors of peacebuilding success. Yet the sign of aid's effect depends on how it is deployed: undirected food transfers can prolong existing conflicts by providing resources to armed factions (Nunn and Qian, 2014). The *design* of transfer schedules—not merely their volume—determines whether aid achieves peace, and this decision remains undertheorized despite its importance.

Existing theoretical frameworks cannot address the design question because they abstract away market structure. The fiscal-instrument view treats aid as government budget augmentation that raises the opportunity cost of rebellion (Savun and Tirone, 2012; Collier and Hoeffler, 2002): it has no relative prices, so it cannot distinguish transfers of oil from transfers of grain. The commitment-device view uses aid as a political incentive lever (Chassang and Padró i Miquel, 2010): it operates through political mechanisms, not commodity markets. In the canonical framework for third-party economic intervention, a donor transfers resources directly to belligerents in a single-good contest; peace is achievable if and only if the donor can close the bargaining gap (Bevia and Corchón, 2010). The theory is clean, but with a single good there is no price mechanism. The limitation is sharp: with multiple goods, when the conflict primitive operates through relative prices, transfers at constant prices have exactly zero effect on militarization (Corollary 6.9). The missing ingredient is the endogenous price response of the conflict economy to commodity transfers.

A growing body of work—theoretical and empirical—establishes that commodity prices affect conflict through general equilibrium channels, motivating the present framework. Dal Bó and Dal Bó (2011) embed conflict in a two-by-two Heckscher-Ohlin economy and identify the Stolper-Samuelson mechanism: increases in the price of labor-intensive goods raise wages and reduce conflict through the opportunity cost channel, while increases in capital-intensive goods raise rents and intensify conflict through the rapacity channel. Dube and Vargas (2013) provide striking empirical confirmation in Colombia: international price increases for oil (a capital-intensive commodity) raise conflict intensity, while price increases for coffee (labor-intensive) reduce it. Garfinkel et al. (2008) embed conflict in a static Walrasian model and characterize how trade regimes affect conflict intensity—the closest precursor to the present framework—but do not study donor intervention or the optimality of transfer schedules. Acemoglu et al. (2012) show that resource price paths shape war incentives dynamically, and Caselli et al. (2015) document that endowment distributions shape conflict incidence across borders. Together, these papers establish that the general equilibrium price mechanism is central to conflict dynamics. A general equilibrium approach is not merely convenient but necessary: the mechanism through which commodity transfers affect conflict—endogenous price adjustment across multiple interdependent markets—is invisible to any framework that holds prices fixed or reduces the economy to a single good.

A natural next question is how a donor should *exploit* that mechanism optimally. More specifically, we ask: what is the optimal dynamic transfer schedule for a donor who understands the general equilibrium of the conflict economy, and how much cheaper is it than the alternatives? The answer has a precise formulation. The *robust peace cost* R^* is the minimum endowment transfer such that every equilibrium at the terminal endowment profile is peaceful. The GE *path cost* P^* is the minimum total transfer variation along any path on the equilibrium manifold from the initial war equilibrium to the peace boundary. Proposition 6.7 establishes $P^* \leq R^*$, with strict inequality whenever the set of equilibria at the terminal endowment generically contains both war and peace outcomes. The title has a precise content: making peace on the cheap means exploiting the equilibrium path to steer the economy to *one* peaceful equilibrium through the market mechanism, without needing to make peace robust to every possible equilibrium selection at the terminal endowment.

The economic question forces a specific mathematical framework. The donor's problem is to choose a path through a space of equilibria—but the equilibrium correspondence is multi-valued and generally discontinuous. To make this a well-posed optimization problem, I model the conflict economy as a *dual competitive equilibrium* (DCE)—a Walrasian trade equilibrium in which belligerents have wealth adjusted by a smooth abstract conflict primitive encoding Nash militarization, conflict payoffs, and aggregate resource drain. Following Balasko (2011), the set of all dual competitive equilibria forms the *equilibrium manifold*, shown to be diffeomorphic to Euclidean space; this converts the multi-valued equilibrium correspondence into a smooth manifold on which optimal control is well defined. The savings $P^* < R^*$ arise because the terminal endowment generically supports both war and peace equilibria—but establishing that this multiplicity is generic requires properness of the natural projection from equilibria to endowments, which the paper proves for the conflict setting. The coordination failure among donors operates through the shared equilibrium path—but defining a "shared path" when equilibria are multiple requires the manifold structure; without it, the public-good nature of the aggregate peace trajectory cannot be formulated. The donor's transfer schedule—an absolutely continuous path in endowment space—induces a path on the equilibrium manifold; the *war region* is defined by a smooth conflict condition, and the donor's problem is to navigate the economy from the initial war equilibrium to the peace boundary at minimum cost.

Three structural results characterize the optimal schedule. An optimal aid schedule always exists (Theorem 5.1), without convexity assumptions, without uniqueness of equilibrium, in the full infinite-dimensional path space—a consequence of the compactness that properness of the natural projection delivers. Under an absorptive capacity constraint, the optimal schedule always disburses at the maximum rate—gradualism is never optimal by choice, requiring strictly convex delivery costs to rationalize (Corollary 6.2). With strictly convex delivery costs, three disbursement regimes emerge as a function of an endogenous urgency measure—the Pontryagin switching function of the associated optimal control problem: the donor pauses when urgency is below a fixed threshold, graduates at an interior rate when urgency is intermediate, and front-loads at the capacity constraint when urgency is high (Proposition 6.3). The urgency measure evolves along the equilibrium path, making disbursement timing endogenous to the economy's trajectory toward peace. Whether a conflict economy should receive front-loaded or graduated aid is not a doctrinal choice but a prediction of the model, determined by the interaction of the economy's position on the equilibrium manifold with the donor's delivery technology.

The paper's most policy-relevant result concerns donor coordination. With n symmetric donors in Nash equilibrium, peace is achieved n times more slowly than under a cooperative arrangement, without any change in total resources (Corollary 7.6). This is a pure coordination failure: the equilibrium path is a public good, and each donor free-rides on the aggregate transfer trajectory. The result yields a new, resource-neutral rationale for multilateral peacemaking institutions—not economies of scale but elimination of the free-rider discount on the peace public good. The donor fragmentation documented by Knack and Rahman (2007) and Acharya et al. (2006) receives a new theoretical explanation: the welfare cost of fragmentation is not duplication of effort but a quantifiable slowdown in the aggregate peace trajectory.

Related literature. The paper contributes to several literatures.

Conflict and general equilibrium. The rationalist approach to war (Fearon, 1995; Powell, 2006) provides the conceptual foundation for the war region, which I take as given. Garfinkel et al. (2008) and Dal Bó and Dal Bó (2011) are the closest theoretical precursors: both embed conflict in general equilibrium, but neither studies donor intervention or the dynamic optimality of transfer schedules. Dube and Vargas (2013) provide the empirical evidence for the heterogeneous price channel that the present model formalizes. Acemoglu et al. (2012) show that resource price dynamics shape war incentives over time, and Caselli et al. (2015) document that endowment geography shapes conflict incidence; both motivate the dynamic general equilibrium formulation here. The present paper adds the donor's problem: dynamic optimal intervention on the equilibrium manifold, which none of these frameworks addresses.

Contest theory. The abstract conflict primitive of section 2 nests the major contest models in the literature: Tullock lotteries (Tullock, 1980), ratio and difference models (Hirshleifer, 1989), axiomatic contest success functions (Skaperdas, 1996), and all-pay auctions (Hodler and Yektaş, 2012); see Konrad (2009, Ch. 2–4) for a unified treatment. The results here depend on smoothness of the conflict equilibrium and an accounting identity, not on the functional form of the contest—the abstraction isolates exactly what the general equilibrium structure needs from the conflict side.

Trade and conflict. Polachek (1980) originates the argument that trade promotes peace through economic interdependence; Martin et al. (2008) formalize this in a model where trade openness raises the opportunity cost of war. The present paper identifies a more specific mechanism: commodity transfers affect conflict through endogenous price adjustment across multiple interdependent markets, and the donor can exploit this price channel optimally.

Aid effectiveness and donor coordination. The empirical literature (de Ree and Nillesen, 2009; Doyle and Sambanis, 2000; Nunn and Qian, 2014; Collier and Hoeffler, 2002; Savun and Tirone, 2012) establishes that aid affects conflict but disagrees on sign, magnitude, and mechanism. Collier and Hoeffler (2002) informally identify equilibrium price adjustment as a channel;

this paper provides the formal theory. The absorptive capacity constraint formalizes Burnside and Dollar (2000), who show that aid effectiveness depends on policy environment and, implicitly, on the economy's capacity to absorb transfers productively.

Burden-sharing and collective action. Olson and Zeckhauser (1966) establish that alliance public goods are undersupplied when members act non-cooperatively—the static logic that the n-donor coordination failure of section 7 dynamizes. The time-to-peace result (Corollary 7.6) is the general equilibrium, continuous-time version of their burden-sharing prediction: fragmentation slows peace not by wasting resources but by diluting each donor's incentive to contribute to the aggregate equilibrium trajectory.

Natural resources and conflict. Ross (2004) surveys the evidence that resource composition—not just resource abundance—shapes conflict; Bazzi and Blattman (2014) find heterogeneous effects of commodity price shocks across resource types. The multi-commodity structure of the present model formalizes this: transfers of different commodities have different effects on conflict through their equilibrium price impacts, and the optimal schedule exploits these differences.

Civil war duration. Fearon (2004) documents that civil wars vary enormously in duration and identifies economic characteristics as key predictors; Collier et al. (2003) argue that conflict traps arise from the interaction of economic and political factors. The model's time-to-peace predictions—and the result that donor fragmentation multiplies duration—provide a structural mechanism for these empirical patterns.

General equilibrium theory. The paper works in the "postmodern" tradition of general equilibrium theory that starts from the set of equilibria and studies its global properties—an approach developed by Debreu (1970) and extended by Balasko (1988, 2011), who establishes that the equilibrium manifold is diffeomorphic to Euclidean space and that the natural projection from equilibria to endowments is proper. I extend these results from exchange economies to conflict economies with a dual competitive equilibrium structure and exploit the manifold to define and solve the donor's optimal control problem—an application of Pontryagin's principle, a standard tool in resource economics, to a novel state space.

The remainder proceeds as follows. Section 2 introduces the conflict economy and the abstract conflict primitive; section 3 constructs the equilibrium manifold and establishes properness; section 4 defines feasible aid schedules and the preference functional; section 5 establishes existence; section 6 derives the Pontryagin characterization with bang-bang and gradualism results; and section 7 extends to multiple donors. Section A verifies the battlefield example satisfies the abstract conflict primitive; section D discusses the topology of loans versus grants.

2 The Model

Four actors—potential belligerents A and B, potential donor C, and the rest of the world W—trade K commodities. We write $I = \{A, B, C, W\}$ for the set of actors, $I_B = \{A, B\}$ for the belligerents, and index commodities by $\{1, \dots, K\}$.

Endowments and utilities. Each actor $i \in I$ holds an initial endowment $\omega_i \in \mathbb{R}^K$ and has a utility function $u_i \colon \mathbb{R}^K \to \mathbb{R}$. Negative entries in ω_i represent liabilities. We collect endowments into the profile $\omega = (\omega_i)_{i \in I} \in \mathbb{R}^{IK} =: \Omega$.

2.1 Assumption (Utility)

For each $i \in I$, u_i is smooth, strictly increasing in each argument, strictly quasiconcave, and has upper contour sets that are closed and bounded below.

Monotonicity ensures budget constraints bind in equilibrium; strict quasiconcavity ensures demand is unique; smoothness and the lower bound together ensure demand is smooth in prices and wealth without requiring a compact choice set. In the conflict setting, these conditions play a specific role beyond standard exchange theory: smoothness of demand is what allows the equilibrium manifold (Proposition 3.2) to inherit smoothness from the consumer side, and uniqueness of demand at each budget ensures that multiplicity of equilibria arises entirely from the interaction of markets and conflict—not from indeterminacy in individual behavior. The lower-bound condition substitutes for compactness of the consumption set, which cannot be imposed when belligerents' effective wealth varies with conflict outcomes.

Prices. A price vector is $p \in \mathbb{R}^K_{++}$. Only relative prices matter, so we work throughout with the numeraire normalization $p^K = 1$; let $\mathcal{P} \coloneqq \mathbb{R}^{K-1}_{++}$ denote the normalized price space.

Conflict. Belligerents A and B interact strategically. Let $I_0 \coloneqq \{A, B, C\}$ and $\Omega_0 \coloneqq \mathbb{R}^{I_0 \times K}$ denote the endowment space of the non-world actors. Rather than commit to a specific contest function—a Tullock lottery (Bevia and Corchón, 2010), an all-pay auction (Hodler and Yektaş, 2012), a

Colonel Blotto game—we characterize the outcome of the militarization game through three smooth objects that encode its reduced form. The abstraction is deliberate: the results of this paper depend on the smoothness of the conflict equilibrium and on an accounting identity (contentious Walras's law, below), not on the functional form of the contest. Smoothness is also easier to assume on the reduced form than to derive from primitives: specific contest functions can produce kinks—corner solutions in all-pay auctions, non-smooth best replies at participation boundaries—that wash out in the equilibrium mapping or that would require ad hoc regularity conditions to handle. The reduced-form approach treats smoothness as a clean, verifiable assumption rather than a derived property that might fail for particular games. The reader may think of the triple $(m^*, (g_i)_{i \in I_B}, r)$ as the "reduced form" of the militarization game, analogous to how Walrasian demand is the reduced form of the consumer's problem: all strategically relevant information is encoded in the three objects

- a *militarization equilibrium* $m^* \colon \mathcal{P} \times \Omega_0 \to \mathbb{R}^2_{\geq 0}$, the unique Nash equilibrium military profile;
- *conflict payoffs* $g_i \colon \mathcal{P} \times \Omega_0 \to \mathbb{R}$ for each $i \in I_B$, giving the net budget impact of equilibrium conflict on belligerent i; and
- a *resource drain* $r \colon \mathcal{P} \times \Omega_0 \to \mathbb{R}^K$, the net resources withdrawn from aggregate supply by conflict.

Each depends on ω_0 but not on ω_W.

2.2 Assumption (Conflict)

The triple $(m^*, (g_i)_{i \in I_B}, r)$ *is smooth on* $\mathcal{P} \times \Omega_0$ *and satisfies* contentious Walras's law*:*

$$\sum_{i \in I_B} g_i(p, \omega_0) \; + \; p \cdot r(p, \omega_0) \; = \; 0 \quad \textit{for all } (p, \omega_0) \in \mathcal{P} \times \Omega_0.$$

Contentious Walras's law is the conflict economy's analogue of the standard accounting identity in exchange economies. Conflict redistributes value among belligerents and destroys resources, but it cannot create value ex nihilo: the budget impacts g_i represent what each belligerent gains or loses in market value from the conflict outcome, and the resource drain r represents what the fighting itself consumes or destroys. The identity requires that every unit of value that disappears from one side of the economy appears on the other side or is accounted for as destroyed supply. This is the consistency condition that allows us to apply Walras's lemma in equilibrium—without it, the market-clearing equations of section 3 would be over-determined. The *adjusted wealth* of belligerent i is

$$w_i(p, \omega) \;\coloneqq\; p \cdot \omega_i + g_i(p, \omega_0).$$

The *contentious demand* of each actor is

$$\widetilde{f}_i(p, \omega) \;\coloneqq\; f_i(p,\; w_i(p, \omega)),$$

setting $g_i \equiv 0$ for $i \notin I_B$. For belligerents, this is Walrasian demand at adjusted wealth; for C and W it reduces to standard Walrasian demand.

2.3 Remark (Battlefield Model)

The following specific model satisfies Assumption 2.2. Each belligerent i chooses militarization $m_i \geq 0$; write $m = (m_A, m_B)$. A conflict technology *$\langle c_i, s_i \rangle$ consists of a cost function $c_i \colon \mathbb{R}_{\geq 0} \times \mathcal{P} \times \Omega_0 \to \mathbb{R}^K_{\geq 0}$ (resources consumed in fielding force m_i) and a battlefield function $s_i \colon \mathbb{R}^2_{\geq 0} \times \mathcal{P} \times \Omega_0 \to \mathbb{R}^K$ (net resource gain or loss from conflict). Belligerent i maximizes $p \cdot [s_i(m, p, \omega_0) - c_i(m_i, p, \omega_0)]$ over $m_i \geq 0$, taking m_{-i} as given. When $\langle c_i, s_i \rangle$ is smooth and satisfies:*

(i) *for each k, $\dfrac{\partial c_i^k}{\partial m_i} \geq 0$ (costs increase weakly in own force);*

(ii) *for each k, $\dfrac{\partial^2 c_i^k}{\partial m_i^2} \geq \left| \dfrac{\partial^2 s_i^k}{\partial m_i^2} \right|$ (cost curvature dominates battlefield curvature—the standard second-order condition ensuring each belligerent's best reply is well-defined);*

(iii) *for each k, $\left| \dfrac{\partial^2 s_i^k}{\partial m_i \, \partial m_{-i}} \right| < \dfrac{\partial^2 c_i^k}{\partial m_i^2} - \dfrac{\partial^2 s_i^k}{\partial m_i^2}$ (cross-effects are strictly small: one belligerent's increase in militarization does not wildly shift the other's incentives, making the game dominance-solvable); and*

(iv) *there exists k such that $\displaystyle\lim_{m_i \to \infty} \frac{\partial c_i^k}{\partial m_i} = \infty$ (some costs eventually become prohibitive, ruling out unbounded militarization at equilibrium);*

then this game has a unique Nash equilibrium $m^*(p, \omega_0)$, smooth in (p, ω_0) (see section A). Setting

$$g_i(p, \omega_0) \;\coloneqq\; p \cdot \big[s_i(m^*, p, \omega_0) - c_i(m_i^*, p, \omega_0)\big],$$
$$r(p, \omega_0) \;\coloneqq\; \sum_{j \in I_B} \big[c_j(m_j^*, p, \omega_0) - s_j(m^*, p, \omega_0)\big]$$

then satisfies Assumption 2.2: smoothness is immediate, and contentious Walras's law holds since $\sum_i g_i + p \cdot r = p \cdot \sum_i (s_i - c_i) + p \cdot \sum_i (c_i - s_i) = 0$.

The contentious economy. We gather the primitives into a *contentious economy*

$$\mathcal{E} \;=\; \big(\omega,\, (u_i)_{i \in I},\, (m^*,\, (g_i)_{i \in I_B},\, r)\big).$$

The endowment profile $\omega \in \Omega$ is our primary parameter; utility functions and conflict primitives are treated as fixed throughout.

3 Equilibrium

We now define equilibrium and establish its global structure. The equilibrium concept combines Walrasian market clearing with the abstract conflict primitive of section 2: non-belligerent actors optimize at posted prices; belligerents consume at prices adjusted for conflict; and markets clear net of the resource drain r.

Non-belligerent demand. Actors C and W are purely economic: they observe prices and choose consumption to maximize utility subject to their budget. For $i \in I \setminus I_B$, the *Walrasian demand* $f_i \colon \mathcal{P} \times \mathbb{R} \to \mathbb{R}^K$ is defined by

$$f_i(p, w) \coloneqq \underset{p \cdot x \le w}{\operatorname{argmax}}\; u_i(x).$$

Under Assumption 2.1, this problem has a unique solution that is smooth in (p, w), and the budget constraint binds: $p \cdot f_i(p, w) = w$.

Belligerent demand. Under Assumption 2.2, belligerent i's equilibrium consumption is the *contentious demand*

$$\widetilde{f}_i(p, \omega) \;\coloneqq\; f_i(p,\; w_i(p, \omega)),$$

Walrasian demand at adjusted wealth $w_i(p, \omega) = p \cdot \omega_i + g_i(p, \omega_0)$. The abstract assumption absorbs all conflict-specific detail; Remark 2.3 verifies that the battlefield model satisfies it.

Dual-competitive equilibrium. The equilibrium concept is dual.

3.1 Definition (Dual-Competitive Equilibrium)

A pair $(p, \omega) \in \mathcal{P} \times \Omega$ *is a* dual-competitive equilibrium *(*DCE*) if:*

(i) (Nash) *militarization is the smooth equilibrium* $m^*(p, \omega_0)$ *from Assumption 2.2; and*

(ii) (Walras) *markets clear:*

$$\sum_{i \in I} \widetilde{f}_i(p, \omega) \;=\; \sum_{i \in I} \omega_i - r(p, \omega_0).$$

The set of all DCE*s is denoted* $\mathcal{D} \subseteq \mathcal{P} \times \Omega$*.*

The right-hand side of the market-clearing condition is aggregate supply: total endowments net of resources consumed by militarization. The left-hand side is aggregate demand: each actor's optimal consumption at equilibrium prices, with belligerents' budgets adjusted for the net value of conflict.

The equilibrium manifold. The set $\mathcal{D}$ has a clean global structure—a property that is essential for what follows. The donor's problem is to choose a path through the set of equilibria, but a priori this set could be badly behaved: disconnected, riddled with singularities, or topologically obstructed. If it were, the optimization problem of section 5 would be ill-posed. The next proposition rules out all such pathologies.

3.2 Proposition (The Equilibrium Manifold)

$\mathcal{D}$ *is a smooth submanifold of* $\mathcal{P} \times \Omega$ *that is diffeomorphic to* $\mathbb{R}^{IK}$*.*

Proof. See section B. ■

We call $\mathcal{D}$ the *equilibrium manifold.* The proposition says that the donor faces a smooth landscape of equilibria that can be parameterized by familiar Euclidean coordinates; the apparent multi-valuedness of the equilibrium correspondence—multiple equilibria at a single endowment profile—is resolved by moving to the richer space of price-endowment pairs. Because $\mathcal{D}$ is diffeomorphic to $\mathbb{R}^{IK}$, it is connected (any two equilibria can be joined by a feasible path), simply connected (no topological obstructions to loans—see section D), and contractible—properties we exploit heavily in the next section.

Equilibrium value. For each actor $i \in I$ and each DCE $e = (p, \omega) \in \mathcal{D}$, we define the *equilibrium utility*

$$V_i(e) \coloneqq u_i\Big(\widetilde{f}_i(p, \omega)\Big),$$

where contentious demand $\widetilde{f}_i$ already incorporates the net impact of conflict on i's wealth. Since $\widetilde{f}_i$ and u_i are both smooth, $V_i \colon \mathcal{D} \to \mathbb{R}$ is smooth, hence continuous.

The natural projection. The equilibrium manifold projects onto endowment space via the *natural projection* $\pi \colon \mathcal{D} \to \Omega$, $\pi(p, \omega) = \omega$. This map plays a central role in section 4.

3.3 Proposition (Properness)

The natural projection $\pi \colon \mathcal{D} \to \Omega$ is smooth and proper: preimages of compact sets are compact.

Proof. See section C. ∎

Smoothness is immediate: π restricts a coordinate projection to the smooth submanifold $\mathcal{D}$. The content is properness, which has a direct economic meaning: bounded endowment transfers produce bounded equilibrium responses. Prices and allocations do not escape to infinity under finite intervention—the donor can be confident that a limited budget produces a limited range of possible outcomes. The argument is standard: if any price diverged, strict monotonicity of preferences would force demand below supply in the corresponding market, violating clearing. It adapts the pure exchange case (Balasko, 2011, Proposition 4.4); the militarization terms remain bounded because m^* is smooth (Assumption 2.2). In the existence proof of section 5, properness is what converts the donor's budget constraint into compactness of the feasible set.

4 Aid Policies

Donor C enters the model as a potential peacemaker. Its instrument is a transfer of endowment to the belligerents: by shifting resources toward A and B, C changes the economic environment in which the militarization game is played, and so alters the equilibrium. This section defines what a transfer policy is, explains how it induces a path on the equilibrium manifold, specifies what makes a policy feasible, and introduces the donor's objective.

Transfer schedules. A *transfer schedule* is an absolutely continuous function

$$\tau\colon [0,1] \longrightarrow \mathbb{R}^K_{\geq 0} \times \mathbb{R}^K_{\geq 0}$$

with $\tau(0) = 0$, where $\tau(t) = (\tau_A(t), \tau_B(t))$ records the cumulative resources transferred from C to A and to B through time t. Non-negativity of τ reflects that C is a donor: it can give resources to the belligerents but cannot extract them. Let $\eta \geq 1$ be an *aid efficiency* parameter: each unit delivered to a belligerent costs C exactly η units, with the residual $\eta - 1$ lost in transit (administrative overhead, leakage, or absorptive friction). The schedule induces an endowment path $\omega\colon [0,1] \to \Omega$ via

$$\omega_A(t) = \omega_A^0 + \tau_A(t), \quad \omega_B(t) = \omega_B^0 + \tau_B(t), \quad \omega_C(t) = \omega_C^0 - \eta\big(\tau_A(t) + \tau_B(t)\big), \tag{1}$$

with $\omega_W(t) = \omega_W^0$ fixed throughout. Aggregate endowments shrink by $(\eta - 1)(\tau_A(t)+\tau_B(t))$ as aid flows; when $\eta = 1$ the transfer is a pure redistribution. The rate of transfer at time t is $\tau'(t)$; its magnitude $\|\tau'(t)\|$ measures how rapidly C is moving resources. Throughout, η is a fixed constant capturing proportional transit losses independent of delivery speed. Speed-dependent delivery costs—the additional inefficiency of rapid disbursement—are modeled separately in section 6 via a convex cost function $C(\|\tau'(t)\|)$ in the donor's objective, following the absorptive-capacity literature (Rajan and Subramanian, 2008).

From transfers to equilibria. Each endowment profile $\omega(t)$ is associated with a set of dual-competitive equilibria $\pi^{-1}(\omega(t)) \subseteq \mathcal{D}$—the *fiber* above $\omega(t)$, consisting of all price-endowment pairs that clear markets at that endowment profile. When equilibria are multiple, the fiber contains more than one point: the same endowment can support different equilibrium prices. A transfer schedule therefore traces a path in endowment space Ω; the question is whether this path lifts uniquely to a path in the equilibrium manifold $\mathcal{D}$.

The answer is yes, generically. At any *regular* point $e_0 = (p^0, \omega^0) \in \mathcal{D}$—one at which the equilibrium responds smoothly to small endowment changes, formally where the differential $D\pi|_{e_0}$ is an isomorphism—the inverse function theorem supplies a smooth local inverse to π near ω^0. Following that local inverse along the endowment path $\omega(t)$ and reapplying the theorem at each step yields a unique smooth path $\gamma\colon [0,1] \to \mathcal{D}$ with $\gamma(0) = e_0$ and $\pi(\gamma(t)) = \omega(t)$ for all t. The only obstruction arises if the path passes through a *critical value* of π—an endowment profile at which equilibria bifurcate or merge, so the local inverse breaks down. Sard's theorem guarantees that such pathologies are rare: the set of critical values has measure zero in Ω, so a generic transfer schedule avoids them entirely and induces a well-defined equilibrium path.

4.1 Definition (Aid Schedule)

An aid schedule *is a Lipschitz path* $\gamma\colon [0,1] \to \mathcal{D}$ *with* $\gamma(0) = e_0$*. We say* γ *is* induced by *the transfer schedule* τ *if* $\pi(\gamma(t)) = \omega(t)$ *for all* t*, where* $\omega(\cdot)$ *is given by* (1)*.*

The economic primitive is the transfer schedule τ; the aid schedule γ is the induced equilibrium trajectory. Because τ is absolutely continuous, the induced endowment path $\omega(\cdot)$ is absolutely continuous, and composing it with the smooth local section supplied by the inverse function theorem yields a Lipschitz γ—the regularity class in which the existence argument of section 5 operates. We work with aid schedules hereafter, since the equilibrium manifold is the natural space for welfare analysis.

The war region. We identify conflict intensity with a smooth function $\phi\colon \mathcal{D} \to \mathbb{R}$ and define the *war region*

$$\mathcal{W} \;:=\; \{e \in \mathcal{D}\colon \phi(e) > 0\}.$$

We treat ϕ as a primitive, requiring only that $\mathcal{W}$ is open with smooth boundary $\partial\mathcal{W} = \phi^{-1}(0)$. The leading example is a militarization threshold: $\phi(e) = \|m^*(e)\| - \bar{m}$ for some $\bar{m} > 0$, so that $\mathcal{W}$ consists of equilibria at which aggregate militarization exceeds $\bar{m}$.

4.2 Assumption (Initial War)

$e_0 \in \mathcal{W}$*: the initial equilibrium lies in the war region.*

Donor C's task is to find an aid schedule that moves the economy out of $\mathcal{W}$ at minimal cost—making peace on the cheap.

Feasibility. An aid schedule γ is *feasible* if three conditions hold.

(i) (Peace) The terminal equilibrium exits the war region: $\gamma(1) \notin \mathcal{W}$, i.e., $\phi(\gamma(1)) \leq 0$.

(ii) (Budget) Donor C cannot transfer resources it does not have. Writing τ for the transfer schedule inducing γ, we require $\eta(\tau_A^k(t)+\tau_B^k(t)) \le \bar{\tau}^k$ for each commodity k and all t, where $\bar{\tau} \in \mathbb{R}_{++}^K$ is an exogenous transfer ceiling (at most $\omega_C^{0,k}$ per commodity).

(iii) (Speed) The aid schedule moves at bounded speed: $\|\gamma'(t)\| \le M$ for almost every t, where $\|\cdot\|$ is the Euclidean norm on $\mathcal{D}$ under the identification of Proposition 3.2, and $M > 0$ is a given constant. Economically, the speed of an aid schedule is controlled by the rate of transfer $\tau'(t)$; bounding it captures the absorptive capacity of the recipient economy: aid delivered too rapidly cannot be productively integrated—ports have throughput limits, bureaucracies have processing capacity, and markets need time to adjust to new supply (Rajan and Subramanian, 2008). The speed bound M is a hard physical ceiling on how fast the equilibrium can move; a separate, softer cost of rapid delivery is modeled via the convex efficiency cost $C(\cdot)$ in section 6.

Write $\mathcal{F}$ for the set of all feasible aid schedules.

4.3 Assumption (Reachability)

$\mathcal{F} \neq \emptyset$: the donor's budget and speed constraints jointly suffice to reach peace.

Assumption 4.3 holds whenever the budget ceiling and time horizon are large enough: formally, whenever $\bar{\tau}^k \ge \eta \cdot \sup_{t\in[0,1]} |\tau^k(t)|$ for each k along some path from e_0 to $\partial\mathcal{W}$ with speed at most M. This is the scope condition that the donor is large enough to matter and the conflict is reachable at the available disbursement rate.

The preference functional. The donor cares about the entire trajectory, not just the terminal equilibrium. A natural way to capture path-dependent preferences is through a 1-form: at each point on the equilibrium manifold, α assigns to each direction of movement a marginal contribution to the donor's welfare, and integrating along the path aggregates these marginal contributions into an overall evaluation. This is just the calculus of line integrals—the same tool used to compute work done by a force field along a trajectory—applied to the donor's preference field over equilibria. Because $\mathcal{D} \cong \mathbb{R}^{IK}$ (Proposition 3.2), the theory of line integrals on Euclidean space applies without modification: for any smooth 1-form α on $\mathcal{D}$ and any aid schedule γ,

$$\mathcal{V}(\gamma) \;:=\; \int_\gamma \alpha \;=\; \int_0^1 \alpha_{\gamma(t)}\big(\gamma'(t)\big)\, dt.$$

We take α as a primitive representing the donor's preferences over trajectories. Different choices accommodate different objectives: $\alpha = -d(p \cdot \omega_C)$ captures the flow cost to C; $\alpha = d(V_A + V_B)$ tracks the welfare gain to the belligerents; more general α can weight both. We do not commit to a particular interpretation here, treating the donor's preferences as given.

The Euclidean structure of $\mathcal{D}$ also clarifies when the path itself matters as opposed to just the destination—an economically meaningful distinction. If $\alpha = dV$ for some smooth $V : \mathcal{D} \to \mathbb{R}$ (i.e., α is the gradient of a function—the analogue of a conservative force field), then

$$\mathcal{V}(\gamma) = V\big(\gamma(1)\big) - V\big(\gamma(0)\big),$$

and the donor's evaluation depends only on the terminal equilibrium, not the transition path. This is natural for a donor whose sole objective is to reduce militarization below a threshold, regardless of what happens along the way. When α is not a gradient—for example, because the donor weighs civilian consumption during the transition, or because delivery costs depend on current equilibrium prices—intermediate equilibria carry independent welfare weight and the path matters. Mathematically, the question reduces to whether the 1-form α is *closed* (its "curl" vanishes): on the simply connected space $\mathcal{D} \cong \mathbb{R}^{IK}$, every closed 1-form is automatically the gradient of a function, so closed and path-independent are synonymous. Preferences are genuinely path-dependent precisely when α is not closed.

Grants and loans. Two special cases deserve mention. An aid schedule γ is a *grant* if $\pi(\gamma(1)) \neq \omega^0$: the terminal endowment profile differs from the initial one, and the transfer is not returned. It is a *loan* if $\pi(\gamma(1)) = \omega^0$: the transfer schedule satisfies $\tau(1) = 0$, so the belligerents' endowments are fully restored at $t = 1$ and C recoups its resources.

A loan traces a loop in endowment space Ω, but the corresponding aid schedule γ need not be a loop in $\mathcal{D}$: the equilibrium price $p(1)$ at the terminal point may differ from p^0, because the fiber $\pi^{-1}(\omega^0)$ can contain multiple equilibria. The continuous path traced through $\mathcal{D}$ determines which equilibrium the economy reaches at the end of the loan—a form of path-dependent equilibrium selection that is invisible when one works in endowment space alone.[1]

Compactness. The existence result of section 5 requires the feasible set to be compact. We verify this now.

4.4 Lemma (Compactness of $\mathcal{F}$)

The set $\mathcal{F}$ of feasible aid schedules is compact in $C^0([0,1],\mathcal{D})$ equipped with the uniform topology.

Proof. Values in a compact set. The budget condition (ii) confines $\omega(t)$ to a bounded region $K \subseteq \Omega$: since $\eta \geq 1$, transfers satisfy $\tau_A^k(t)+\tau_B^k(t) \leq \bar{\tau}^k/\eta \leq \bar{\tau}^k$ componentwise, so each $\omega(t)$ lies in a fixed compact box. Every feasible aid schedule therefore satisfies $\pi \circ \gamma([0,1]) \subseteq K$. By Proposition 3.3, π is proper, so $\pi^{-1}(K)$ is compact in $\mathcal{D}$. Hence every feasible aid schedule takes values in the compact set $\pi^{-1}(K)$.

Equicontinuity. The speed condition (iii) gives $\|\gamma'(t)\| \leq M$ almost everywhere. For any $s, t \in [0,1]$,

$$d(\gamma(s), \gamma(t)) \;\leq\; \int_s^t \|\gamma'(r)\|\, dr \;\leq\; M|t-s|,$$

where d is the Riemannian distance on $\mathcal{D}$. Every feasible aid schedule is therefore M-Lipschitz, hence equicontinuous.

Compactness. By the Arzelà–Ascoli theorem—which states that a uniformly bounded and equicontinuous family of functions is pre-compact—$\mathcal{F}$ is pre-compact in $C^0([0,1],\mathcal{D})$. To conclude compactness, we show $\mathcal{F}$ is closed. Conditions (ii) and (iii) are preserved under uniform limits by continuity of evaluation and integration. For condition (i): if $\gamma^n \to \gamma$ uniformly and $\phi(\gamma^n(1)) \leq 0$ for all n, then $\phi(\gamma(1)) = \lim_n \phi(\gamma^n(1)) \leq 0$ by continuity of ϕ. Hence $\mathcal{F}$ is closed, and compactness follows. ∎

The framework in perspective. Three features of this setup shape the results that follow.

Aid as navigation. The donor steers the economy through a sequence of general equilibria, with prices adjusting at every step. The same resource transfer can reach different equilibria depending on the path through $\mathcal{D}$,

[1] This phenomenon is governed by the monodromy of π: whether a loan returns the economy to e_0 depends on whether its transfer schedule, viewed as a loop in $\Omega \setminus \Sigma$ (the complement of the discriminant locus), acts trivially on e_0 under the monodromy action of $\pi_1(\Omega \setminus \Sigma, \omega^0)$. See section D.

and the equilibrium mechanism can amplify or reverse the direct effect of endowment movements.

Process and outcome. When $\alpha = dV$, the donor evaluates only the terminal equilibrium—any path to the same point on $\partial\mathcal{W}$ is equivalent. When α is not closed, intermediate equilibria carry welfare weight, and the optimal schedule adjusts continuously along the trajectory (section 6).

Finance and topology. A loan traces a loop in endowment space, but the equilibrium the economy occupies after repayment depends on which element of $\pi^{-1}(\omega^0)$ the path terminates at—a topological question governed by the monodromy of π (section D).

5 Existence of an Optimal Aid Schedule

A best aid schedule always exists. The donor's optimization problem is well-posed despite the infinite-dimensional path space, the potential multiplicity of equilibria along the path, and the absence of any convexity assumption on preferences. This result is not obvious: the donor chooses among all feasible paths through the equilibrium manifold, a space far larger than the finite-dimensional strategy sets of standard contest models. It holds because the general equilibrium structure of the conflict economy—specifically, the properness and Euclidean topology of the equilibrium manifold established in section 3—constrains the feasible set tightly enough to guarantee a solution.

The proof strategy has a direct economic interpretation. The donor's budget and absorptive capacity constraints confine the transfer schedule to a bounded region of endowment space; properness of the natural projection (Proposition 3.3) then confines the induced equilibrium path to a compact set—prices and allocations cannot escape to extreme values under finite intervention. Compactness of the feasible set, combined with continuity of the donor's preferences along the equilibrium path, delivers a maximum by the extreme value theorem. The technical content is in establishing compactness and continuity in the appropriate topologies, which we now do.

5.1 Theorem (Existence of an Optimal Aid Schedule)

The donor's optimization problem

$$\max_{\gamma\in\mathcal{F}} \mathcal{V}(\gamma) \qquad \left(\text{and} \quad \min_{\gamma\in\mathcal{F}} \mathcal{V}(\gamma)\right)$$

each have a solution: there exists $\gamma^* \in \mathcal{F}$ attaining the supremum of $\mathcal{V}$ over $\mathcal{F}$, and a $\gamma_* \in \mathcal{F}$ attaining the infimum.

The theorem requires no convexity of the donor's preferences and no uniqueness of equilibrium along the path—conditions that standard welfare theorems rely on but that cannot be expected in a conflict economy with multiple equilibria. The equilibrium manifold structure does all the work: it replaces these assumptions with topological compactness.

Proof. By the extreme value theorem, it suffices to find a topology on $\mathcal{F}$ in which $\mathcal{F}$ is compact and $\mathcal{V}$ is continuous.

Topology. Equip $\mathcal{F}$ with the topology τ of *uniform path convergence and weak-$*$ velocity convergence*: a net $\gamma^\lambda \to \gamma$ in τ if $\gamma^\lambda \to \gamma$ uniformly in $C^0([0,1],\mathcal{D})$ and $(\gamma^\lambda)' \rightharpoonup^* \gamma'$ in $L^\infty([0,1];T\mathcal{D})$. This topology captures the economically relevant notion of convergence: nearby schedules produce nearby equilibrium paths (uniform convergence) with similar disbursement rates (weak-$*$ convergence of velocities). Since $L^1([0,1];T\mathcal{D})$ is separable, the weak-$*$ topology on the M-ball of L^∞ is metrizable; hence τ is metrizable on $\mathcal{F}$ and sequential compactness is equivalent to compactness.

Compactness. Let $\{\gamma^n\}$ be a sequence in $\mathcal{F}$. By Lemma 4.4, $\mathcal{F}$ is compact in C^0, so a subsequence γ^{n_k} converges uniformly to some $\gamma \in \mathcal{F}$. Because every γ^n is M-Lipschitz, the velocities $\{(\gamma^{n_k})'\}$ lie in the M-ball of $L^\infty([0,1];T\mathcal{D})$. By the Banach–Alaoglu theorem—the infinite-dimensional analogue of the Bolzano–Weierstrass theorem, which guarantees that bounded sequences in dual spaces have convergent subsequences—a further subsequence has $(\gamma^{n_{k_j}})' \rightharpoonup^* \gamma'$ in L^∞. Hence $\mathcal{F}$ is sequentially compact, hence compact, in τ. Economically: the set of feasible aid schedules is "closed and bounded" in the relevant sense—the donor cannot escape to an arbitrarily good schedule by taking limits.

Continuity of $\mathcal{V}$. Under the identification $\mathcal{D} \cong \mathbb{R}^{IK}$ of Proposition 3.2, write the preference functional as

$$\mathcal{V}(\gamma) = \int_0^1 \langle a(\gamma(t)),\, \gamma'(t)\rangle \, dt,$$

where $a\colon \mathcal{D} \to \mathbb{R}^{IK}$ is the smooth vector of component functions of α in the coordinates of $\mathcal{D}$, and $\langle \cdot, \cdot \rangle$ is the Euclidean inner product. Suppose $\gamma^n \to \gamma$

in τ. Decompose:

$$\mathcal{V}(\gamma^n) - \mathcal{V}(\gamma) = \int_0^1 \langle a(\gamma),\, (\gamma^n)' - \gamma' \rangle \, dt \\ + \int_0^1 \langle a(\gamma^n) - a(\gamma),\, (\gamma^n)' \rangle \, dt.$$

The first integral tends to zero by weak-$*$ convergence of $(\gamma^n)'$ to γ', tested against the L^1 function $a(\gamma(\cdot))$. The second integral satisfies

$$\left| \int_0^1 \langle a(\gamma^n) - a(\gamma),\, (\gamma^n)' \rangle \, dt \right| \;\le\; \|a(\gamma^n) - a(\gamma)\|_\infty \cdot M \;\to\; 0,$$

since a is continuous and $\gamma^n \to \gamma$ uniformly. Hence $\mathcal{V}(\gamma^n) \to \mathcal{V}(\gamma)$: the donor's welfare varies continuously with the aid schedule.

Conclusion. $\mathcal{F}$ is compact and $\mathcal{V}$ is continuous in τ, so the extreme value theorem delivers γ^* and γ_*. ■

5.2 Remark (Exact preferences)

When the donor evaluates only the terminal equilibrium—that is, when $\alpha = dV$ for some smooth $V\colon \mathcal{D} \to \mathbb{R}$, the case of path-independent preferences (section 4)—the argument simplifies considerably: $\mathcal{V}(\gamma) = V(\gamma(1)) - V(e_0)$ depends only on the terminal equilibrium, so continuity in the C^0 topology is immediate and the Banach–Alaoglu step is unnecessary.

What existence guarantees—and what it does not. Theorem 5.1 establishes that the donor's problem is well-posed: optimality is global—the theorem delivers a supremum attainer, not merely a local optimum. Multiple optimal schedules can coexist; the results below characterize what all optima share rather than singling one out.

6 Structure of Optimal Aid Schedules

Theorem 5.1 guarantees that an optimal aid schedule exists. We now characterize its structure. The donor's problem has the form of a classical optimal control problem, and Pontryagin's maximum principle supplies necessary conditions that every optimal schedule must satisfy. The central finding is that optimal aid is *bang-bang in speed*: the donor always delivers at the maximum rate M, and the interesting choice is direction, not pace.

Pontryagin's maximum principle is the continuous-time analogue of Lagrange multiplier conditions in static optimization. Where Lagrange conditions characterize optima of a function subject to equality constraints, the maximum principle characterizes optimal *trajectories* subject to constraints on the rate of change—exactly the structure of the donor's problem, where the control (the disbursement rate) is bounded by absorptive capacity. The principle introduces a *costate* variable $\psi(t)$—a vector of shadow prices, one for each dimension of the state space—that evolves alongside the state and captures the marginal value of moving the economy in each direction at each moment. The key output is the *switching function* $S(t)$, which combines the direct preference gain with the shadow value to determine the optimal direction and speed of aid at each instant.

The optimal control problem. Under the identification $\mathcal{D} \cong \mathbb{R}^{IK}$ of Proposition 3.2, write the donor's problem as follows. Let $\Phi\colon \mathcal{D} \to \mathbb{R}^{IK}$ denote the coordinate map of the diffeomorphism—the map that labels each equilibrium by its prices, incomes, and non-numeraire endowments (constructed explicitly in section B). The *state* is $x(t) = \Phi(\gamma(t)) \in \mathbb{R}^{IK}$; the *control* is $u(t) = x'(t)$; the *dynamics* are trivially $\dot{x} = u$; and the *objective* is

$$\mathcal{V}(x, u) \;=\; \int_0^1 \langle a(x(t)),\, u(t)\rangle \, dt,$$

where $a\colon \mathbb{R}^{IK} \to \mathbb{R}^{IK}$ collects the component functions of α in the Φ-coordinates. The constraints are: initial condition $x(0) = \Phi(e_0)$; terminal peace condition $\phi(x(1)) \leq 0$; and speed constraint $\|u(t)\| \leq M$ almost everywhere. We assume the budget condition holds in the interior; the case of a binding budget constraint admits an analogous treatment via additional costate multipliers.

The Hamiltonian and switching function. The Pontryagin Hamiltonian is

$$H(x, \psi, u) \;\coloneqq\; \langle a(x) + \psi,\, u\rangle,$$

where $\psi(t) \in \mathbb{R}^{IK}$ is the costate—the donor's shadow value of the economy's position on the equilibrium manifold. The function H is linear in the control u, so its maximizer over the closed ball $\|u\| \leq M$ is always at the boundary. Define the *switching function*

$$S(t) \;\coloneqq\; a\big(x^*(t)\big) + \psi(t).$$

The switching function measures the donor's *urgency*—the marginal value of disbursing one more unit of aid along each commodity dimension, combining the direct preference gain $a(x)$ with the shadow value ψ of moving the economy forward. When urgency is high, the donor front-loads; when urgency is low, the donor pauses. Formally: when $S(t) \neq 0$, the unique maximizer of H is $u^*(t) = MS(t)/\|S(t)\|$, and the maximum speed is always attained.

The optimal aid schedule is characterized by three conditions: (i) the donor's shadow value of the equilibrium state evolves along the optimal path; (ii) at each moment, the donor disburses in the direction that maximizes the instantaneous return to aid, net of delivery costs; and (iii) the shadow value at the terminal time reflects only the value of reaching peace.

6.1 Theorem (Pontryagin Characterization)

Let γ^* be an optimal aid schedule and $x^* = \Phi \circ \gamma^*$ its trajectory in $\mathbb{R}^{IK}$. Then there exist an absolutely continuous costate $\psi\colon [0,1] \to \mathbb{R}^{IK}$ and a multiplier $\lambda \geq 0$ such that for almost every $t \in [0,1]$:

(i) (Costate) $\dot{\psi}(t) = -Da(x^*(t))^\top u^*(t)$, *where Da is the Jacobian of a;*

(ii) (Maximality) $u^*(t) \in \operatorname*{argmax}_{\|u\| \leq M} \langle a(x^*(t)) + \psi(t),\, u\rangle$*;*

(iii) (Transversality) $\psi(1) = \lambda\, \nabla\phi(x^*(1))$, *with $\lambda\, \phi(x^*(1)) = 0$.*

Moreover, the switching function satisfies

$$\dot{S}(t) \;=\; \iota_{u^*(t)}\, d\alpha\big|_{x^*(t)}, \tag{2}$$

Recall that the 1-form α assigns a marginal welfare contribution to each direction of movement at each point on the equilibrium manifold (section 4). Its exterior derivative $d\alpha$ is a 2-form that measures how these marginal contributions vary across directions*: if the welfare gain from transferring grain changes as the economy moves along the fuel dimension, $d\alpha$ is nonzero. When α is a gradient ($d\alpha = 0$), these cross-variations vanish, the switching function is constant, and the optimal commodity mix never changes. When $d\alpha \neq 0$, the relative value of different transfer directions shifts as the economy moves, and the switching function tracks this shift. The interior product $\iota_{u^*}\, d\alpha$ extracts the rate of this shift along the current direction of travel—it is the directional derivative of S along the optimal path.*

Proof. Conditions (i)–(iii) are the Pontryagin maximum principle for an optimal control problem with a terminal inequality constraint; see Liberzon (2011, Ch. 4). The constraint qualification holds because $\nabla\phi(x^*(1)) \neq 0$ ($\partial\mathcal{W}$ is smooth by assumption), giving a normal extremal.

For (2): differentiate $S = a(x^*) + \psi$ and substitute (i) to obtain

$$\dot{S} \;=\; Da\,u^* + \dot{\psi} \;=\; Da\,u^* - Da^\top u^* \;=\; (Da - Da^\top)\,u^*.$$

In coordinates, $[(Da - Da^\top)u^*]_j = \sum_k(\partial a_j/\partial x_k - \partial a_k/\partial x_j)u_k^*$. Since $(d\alpha)_{jk} = \partial a_k/\partial x_j - \partial a_j/\partial x_k$, one computes $(\iota_{u^*}d\alpha)_j = -\sum_k(d\alpha)_{jk}u_k^* = \sum_k(\partial a_j/\partial x_k - \partial a_k/\partial x_j)u_k^*$, which matches. ∎

6.2 Corollary (Bang-Bang)

In the non-singular case—whenever $S(t) \neq 0$—the optimal aid schedule moves at speed exactly M almost everywhere. Singular arcs, intervals on which $S \equiv 0$, require $u^*(t)$ to lie in the characteristic distribution of $d\alpha$ (the kernel of $v \mapsto \iota_v\,d\alpha$). Since $IK = 4K$ is even for any number of commodities K, a generic 2-form $d\alpha$ on $\mathbb{R}^{IK}$ is non-degenerate, its characteristic distribution is trivial, and no singular arcs at positive speed exist. Economically, non-degeneracy means that the donor's urgency generically changes direction as the economy moves through the equilibrium manifold—the donor is never exactly indifferent across all commodity dimensions simultaneously, ruling out intervals of indeterminate speed. In the language of differential forms: a 2-form on an even-dimensional space is generically non-degenerate (analogous to a square matrix being generically invertible), and non-degeneracy means that no direction of movement is "invisible" to the preference twist—the donor's urgency responds to movement in every commodity dimension.

Proof. Maximality gives $u^* = MS/\|S\|$ when $S \neq 0$, so $\|u^*\| = M$. When $S = 0$, (2) forces $\iota_{u^*}d\alpha = 0$, placing u^* in the characteristic distribution. A 2-form on $\mathbb{R}^n$ is non-degenerate—equivalently, symplectic—if and only if n is even and the form has maximal rank; this is the generic case. ∎

The economic content is direct: under generic donor preferences, *optimal aid is never gradual by choice*. Given that the donor is moving the economy, it should do so as fast as the absorptive capacity constraint permits. Pacing aid at less than the maximum rate is suboptimal unless the donor happens to be exactly indifferent at the margin—a non-generic condition on α. The speed constraint M captures absorptive capacity as a physical limit; efficiency grounds for gradualism require that efficiency losses enter the donor's objective, as we now show.

Efficiency costs and interior optima. Suppose the preference functional includes a speed-dependent efficiency cost:

$$\mathcal{V}_C(\gamma) \;\coloneqq\; \int_0^1 \Big[\langle a(\gamma(t)),\, \gamma'(t)\rangle \;-\; C(\|\gamma'(t)\|)\Big]\, dt,$$

where $C\colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is smooth, strictly convex, and increasing, with $C(0) = 0$. The term $C(\|u\|)$ captures the efficiency loss $\eta(\|u\|)$ per unit delivered: faster aid is proportionally more expensive. The Hamiltonian becomes

$$H_C(x, \psi, u) \;\coloneqq\; \langle a(x) + \psi,\, u\rangle \;-\; C(\|u\|) \;=\; \langle S,\, u\rangle - C(\|u\|),$$

which is strictly concave in u (since C is strictly convex). The maximum over $\|u\| \leq M$ is therefore at an interior point whenever $\|S\| < C'(M)$, and the first-order condition gives the following.

6.3 Proposition (Efficiency Costs and Gradualism)
Let C be strictly convex and increasing. The optimal speed $\|u^*(t)\|$ and direction $u^*(t)/\|u^*(t)\|$ satisfy:

(i) *the optimal direction is $S(t)/\|S(t)\|$, unchanged from the bang-bang case;*

(ii) *the optimal speed is determined by*

$$\|u^*(t)\| \;=\; \begin{cases} 0 & \|S(t)\| \leq C'(0), \\ (C')^{-1}(\|S(t)\|) & C'(0) < \|S(t)\| < C'(M), \\ M & \|S(t)\| \geq C'(M). \end{cases}$$

The optimal speed is strictly increasing in $\|S(t)\|$ and strictly decreasing in C' at every interior point.

Proof. Fixing direction $v = u/\|u\|$, the reduced problem maximizes $\|S\| \cdot s - C(s)$ over $s \in [0, M]$. The objective is strictly concave in s; the unconstrained maximizer is at $C'(s) = \|S\|$, i.e., $s = (C')^{-1}(\|S\|)$, which is feasible when $C'(0) \leq \|S\| \leq C'(M)$. Outside this range, the optimum is at the boundary $s = 0$ or $s = M$. The optimal direction maximizes $\langle S, v\rangle$ over unit vectors, giving $v = S/\|S\|$. ∎

Proposition 6.3 delivers a three-regime characterization of optimal aid plans. Take the leading special case $C(s) = \eta_0 s + \kappa s^2$ (linear η, quadratic cost), so $C'(s) = \eta_0 + 2\kappa s$ for parameters $\eta_0 > 0$ (fixed delivery cost) and $\kappa > 0$ (marginal inefficiency of speed). Then:

(i) (Pause) if $\|S(t)\| \le \eta_0$: no aid is delivered. The preference gain is insufficient to cover the fixed efficiency cost η_0.

(ii) (Gradual) if $\eta_0 < \|S(t)\| < \eta_0 + 2\kappa M$: aid is delivered at interior speed $\|u^*(t)\| = (\|S(t)\| - \eta_0)/(2\kappa)$. Higher η_0 or κ implies slower optimal pace; stronger preference intensity $\|S\|$ implies faster pace.

(iii) (Rapid) if $\|S(t)\| \ge \eta_0 + 2\kappa M$: the speed constraint binds and the schedule is bang-bang at rate M.

The regime a particular donor operates in—and hence whether the optimal plan is short or long—is determined by the interplay of the switching function $S(t)$ (endogenous to the equilibrium path) and the efficiency parameters (η_0, κ) (primitives of the delivery technology). Countries where peace is urgently needed—high $\|S\|$—attract rapid transfers; those with lower urgency receive graduated aid; marginal cases receive nothing.

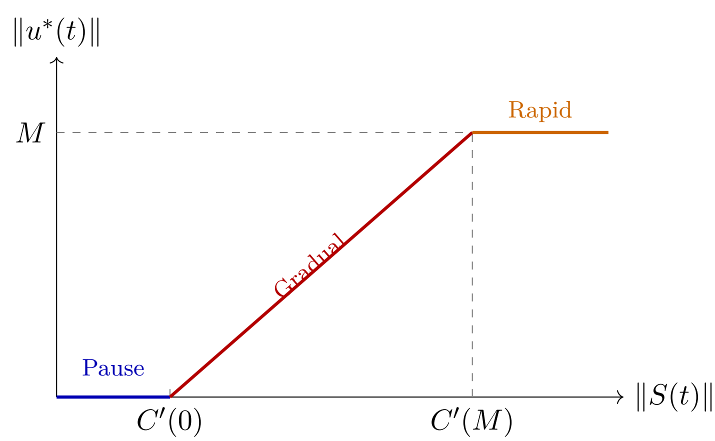

Figure 1: Optimal disbursement speed as a function of the switching-function magnitude $\|S(t)\|$ under quadratic efficiency costs $C(s) = \eta_0 s + \kappa s^2$. Below $C'(0) = \eta_0$, the preference gain does not cover the fixed cost and no aid flows. Between $C'(0)$ and $C'(M)$, the donor graduates at an interior rate. Above $C'(M)$, the speed constraint binds.

6.4 Remark (Cheapest $\neq$ Fastest)

The results above jointly imply that the welfare-optimal aid schedule is generally neither the shortest nor the fastest path to peace. Two distinct forces drive the wedge.

Direction. *Even in the bang-bang case, the switching function $S(t)$ governs direction, and S evolves according to the donor's preference form α via $\dot{S} = \iota_{u^*}\, d\alpha$—not according to the geometry of $\mathcal{W}$. A donor who values a particular type of peace (one that restores specific trade flows, say, or that leaves a preferred balance of power) follows a curved trajectory through $\mathcal{D}$, incurring a longer time-to-peace in exchange for arriving at a preferred point on $\partial\mathcal{W}$. The fastest path would head straight for the nearest point on $\partial\mathcal{W}$; the optimal path heads for the best point.*

Speed. *Efficiency costs introduce a second departure: the optimal speed falls strictly below M in the gradual regime. Absorptive-capacity constraints ($\kappa > 0$) make rapid delivery costly, so the donor deliberately slows down—extending the schedule to reduce waste. The "cheapest" plan in the welfare sense is therefore longer in calendar time than a pure time-minimizing plan would be.*

Together, these two forces mean that cross-country variation in aid schedules reflects heterogeneity in donor preferences and *in delivery technology, not just variation in how far each economy is from peace.*

6.5 Corollary (Geodesics under Exact Preferences)

If $\alpha = dV$ for some smooth $V : \mathcal{D} \to \mathbb{R}$, then $d\alpha = 0$ and the switching function is constant along every optimal path. The optimal aid schedule is a geodesic*: it moves in a fixed direction at speed M, connecting e_0 to the terminal point*

$$x^*(1) \in \operatorname*{argmax}_{x \in \partial\mathcal{W}} V(x) \quad (\textit{resp.}\ \operatorname{argmin}).$$

The optimal allocation of aid between A and B, and across commodities, is therefore constant *throughout the schedule.*

Proof. With $a = \nabla V$ one has $Da = D^2 V$ (symmetric), so $Da - Da^\top = 0$ and $\dot{S} = 0$. The switching function $S_0 = S(t)$ is constant, giving $u^*(t) = MS_0/\|S_0\|$ (constant direction, speed M). The transversality condition identifies $S_0 = \nabla V(x^*(1)) + \lambda \nabla\phi(x^*(1))$, so $x^*(1)$ is the constrained optimum of V on $\partial\mathcal{W}$ by the Lagrange multiplier theorem. ∎

Corollary 6.5 is the simplest possible aid schedule: pick the best peace equilibrium on $\partial\mathcal{W}$ and go straight there at full speed, never adjusting the commodity mix or the split between A and B. It arises when the donor cares only about where the economy ends up, not about the transition—the path-independent case of section 4. This is a useful benchmark. Non-exact preferences break the constancy: the switching function evolves as the economy moves through $\mathcal{D}$ (via (2)), the relative value of different transfer directions shifts, and the optimal schedule curves through the equilibrium manifold rather than following a straight line. The geodesic case predicts a rigid aid program—fixed commodity shares, constant allocation across belligerents—while the non-exact case predicts an adaptive one.

6.6 Remark (Singular arcs)

A singular arc *is an interval during which the switching function vanishes identically: the donor is exactly indifferent about speed, neither preferring to pause nor to go faster. On a singular arc, the bang-bang result (Corollary 6.2) does not pin down the optimal speed—any speed is equally good, and the solution is not uniquely determined. This indeterminacy is non-generic: it requires the donor's preference form α to be tuned so that urgency vanishes over an entire interval, a condition that fails for "almost all" α. The efficiency-cost formulation of Proposition 6.3 resolves the issue completely: with C strictly convex, the Hamiltonian H_C is strictly concave in u, so the optimal speed is always uniquely determined and singular arcs cannot arise. In practice, this means the three-regime characterization of Proposition 6.3—pause, graduate, front-load—is the complete description of optimal aid timing whenever delivery costs are convex.*

Call a transfer $\Delta = (\Delta_A, \Delta_B) \geq 0$ a *robust peace transfer* if every equilibrium at the resulting endowment profile is peaceful:

$$\pi^{-1}\big(\omega_A^0 + \Delta_A,\ \omega_B^0 + \Delta_B,\ \omega_C^0 - \eta(\Delta_A + \Delta_B)\big)\ \subseteq\ \mathcal{D} \setminus \mathcal{W}.$$

Robust peace requires peace regardless of which equilibrium the economy coordinates on—appropriate for a donor who cannot influence equilibrium selection. Define the *robust peace cost*

$$R^*(e_0)\ =\ \eta \cdot \inf\big\{\|\Delta\| : \Delta \text{ is a robust peace transfer}\big\}.$$

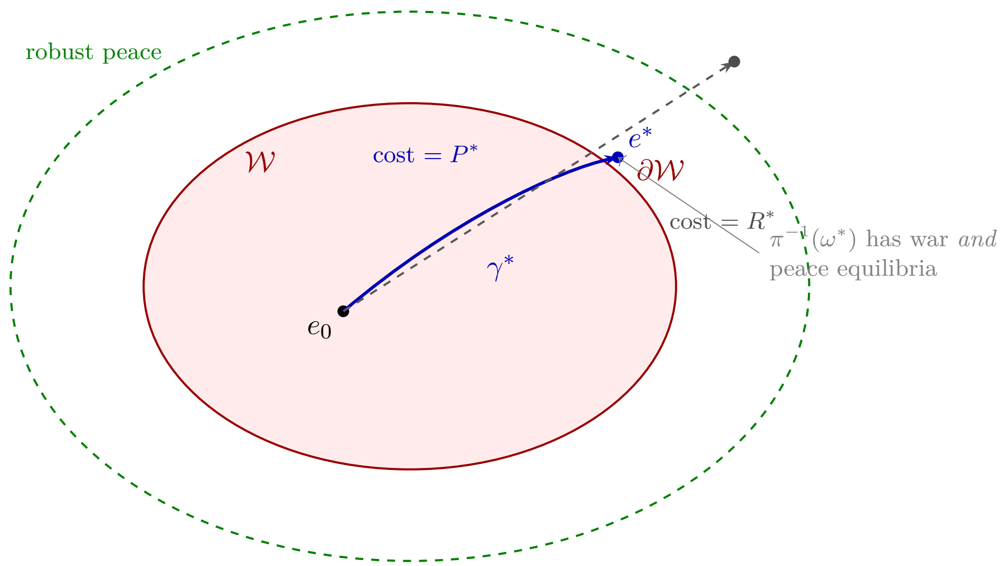

Figure 2: The optimal GE path γ^* reaches peace at $e^* \in \partial\mathcal{W}$ at cost P^*. Because the fiber $\pi^{-1}(\omega^*)$ still contains war equilibria, ω^* is not a robust peace endowment, and the robust peace cost R^* requires a strictly larger transfer that clears all war equilibria from the terminal fiber.

6.7 Proposition (Making Peace on the Cheap)

$P^*(e_0) \le R^*(e_0)$. *The savings* $\mathcal{S}(e_0) := R^*(e_0) - P^*(e_0)$ *is strictly positive whenever the terminal fiber* $\pi^{-1}(\omega^*)$ *contains both war and peace equilibria—a condition that is generic by properness of* π.

Proof. Weak inequality. Any robust peace transfer Δ provides a straight-line path $\tau(t) = t\Delta$ that crosses $\partial\mathcal{W}$ at some time $T \le 1$, with cost $\eta \int_0^T \|\tau'(t)\|\, dt \le \eta\|\Delta\|$. Taking the infimum over all such Δ gives $P^*(e_0) \le R^*(e_0)$.

Strictness under multiplicity. Let γ^* be an optimal GE path terminating at $e^* \in \partial\mathcal{W}$, and set $\omega^* = \pi(e^*)$. Suppose $\pi^{-1}(\omega^*)$ also contains a war equilibrium $e' \in \mathcal{W}$. Then ω^* is not a robust peace endowment, so every robust peace transfer Δ satisfies $\omega^0 + \Delta \neq \omega^*$; in particular, $\|\Delta\| > \|\omega^* - \omega^0\|/\eta$. Since $P^*(e_0)$ is at most the cost of reaching e^* via γ^*, which requires endowment displacement at most $\|\omega^* - \omega^0\|$, we have $P^*(e_0) < \eta\|\Delta\|$ for every robust peace transfer Δ, and hence $P^*(e_0) < R^*(e_0)$.

Genericity of mixed fibers. The key economic condition for strict savings is that the terminal endowment profile ω^* supports both war and peace equilibria simultaneously—the economy *could* be at war or at peace at the same endowment, depending on which equilibrium it coordinates on. This is the generic situation: in general equilibrium models, multiple equilibria at a single endowment profile are the rule rather than the exception, and global uniqueness is a non-generic knife-edge condition (Balasko, 2011, Ch. 5). Formally: since $\mathcal{D}$ is connected (Proposition 3.2) and $\mathcal{W}$, $\mathcal{D} \setminus \overline{\mathcal{W}}$ are non-empty open subsets, both $\pi(\mathcal{W})$ and $\pi(\mathcal{D} \setminus \overline{\mathcal{W}})$ are open and non-empty in Ω. At any regular value ω in their intersection, the fiber $\pi^{-1}(\omega)$ contains equilibria from both $\mathcal{W}$ and $\mathcal{D} \setminus \overline{\mathcal{W}}$. The intersection is non-empty whenever π fails to be globally injective near $\partial\mathcal{W}$—the generic situation. ∎

6.8 Remark (Cheap for Whom?)

"Cheap" in Proposition 6.7 refers to the donor's cost. Whether the optimal schedule also improves belligerent welfare depends on the preference form α. When α incorporates belligerent utility—e.g., $\alpha = d(V_A + V_B) - \eta\, d(p \cdot \omega_C)$—the optimal schedule maximizes a weighted combination of peace and belligerent welfare, and Pareto improvement is built in. When α reflects only the donor's strategic or financial interest, the optimal schedule may steer to a peace that is cheap for C but not Pareto-improving for A and B. The framework accommodates both; the distinction is a choice of α.

The savings in Proposition 6.7 have a sharp source: the general equilibrium price adjustment does work that redistribution at constant prices cannot. The next result isolates this mechanism in the battlefield model, and the numerical illustration that follows quantifies it.

6.9 Corollary (PE Transfers Cannot Reduce Conflict)

In the battlefield model of Remark 2.3, suppose the conflict prize is denominated in good j and militarization costs are denominated in good $k \neq j$. Then the Nash militarization $m^*(p, \omega_0)$ depends on endowments only through the equilibrium price ratio p^j/p^k. At fixed prices, any commodity transfer from C to the belligerents leaves the conflict equilibrium unchanged: m^* is invariant to endowment redistribution conditional on p.

Proof. With prize in good j and cost in good k, belligerent i maximizes $p^j s_i^j(m) - p^k c_i^k(m_i)$ over $m_i \geq 0$. The first-order condition depends on p but not on ω. At fixed p, transfers shift only the budget constraints, not the militarization game. Hence $m^*(p, \omega_0) = m^*(p, \omega_0')$ for any two endowment profiles ω_0, ω_0' sharing the same prices. ∎

Corollary 6.9 makes the role of the price mechanism precise. When prize and cost are denominated in the same good ($j = k$), PE transfers can affect militarization through direct wealth effects—the single-good mechanism of Bevia and Corchón (2010). The GE price channel is then additive rather than exclusive; Corollary 6.9 isolates the case where it is the *only* channel. When the conflict primitive operates through market prices (the opportunity-cost and rapacity channels of Dal Bó and Dal Bó 2011; Dube and Vargas 2013), direct redistribution at constant prices has zero effect on militarization. Only the *general equilibrium* price adjustment—the endogenous response of p to the change in aggregate supply—reduces conflict.

6.10 Remark (Numerical Illustration)

Figure 3 illustrates Corollary 6.9 in a calibrated two-good economy ($K = 2$, good 2 numeraire) with Cobb-Douglas preferences, a Tullock contest over a prize $R = 2$ of good 1 with quadratic cost $\bar{c} = 3$ in good 2, and peace threshold $\bar{m} = 0.55$. The symmetric Nash militarization $m_i^* = \sqrt{p^1 R/4\bar{c}}$ depends on the price p^1 but not on the endowment distribution. Under the GE path (solid), each transfer lowers p^1 through market clearing, driving $\|m^*\|$ below $\bar{m}$ at a cumulative transfer of $\delta^* = 2.33$ (29% of C's good-1 endowment). Under the PE benchmark (dashed), the same transfers at fixed p^1 leave $\|m^*\|$ unchanged: the line is flat.

---

7 Multiple Donors

The single-donor analysis of sections 5 and 6 treats C as a unitary actor. In practice, foreign aid flows from many sources—bilateral donors, multilateral institutions, NGOs—each with their own budget, efficiency, and preferences. We now extend the model to $n \geq 1$ donors $C_1, \ldots, C_n$ and ask: does a Nash equilibrium among donors exist, what does it look like, and how does it compare to the cooperative optimum of earlier sections?

The multi-donor setting introduces a fundamental strategic complication absent from the single-donor case. Peace—more precisely, the equilibrium trajectory that brings the economy from e_0 to $\partial\mathcal{W}$—is a public good: all donors benefit from any progress toward peace regardless of who paid for it, because each donor's preference functional $\mathcal{V}_j(\gamma)$ depends on the *aggregate* equilibrium path induced by all transfers combined. Each donor therefore has an incentive to free-ride on others' contributions, and the Nash equilibrium reflects this collectively: aggregate momentum toward peace is un-

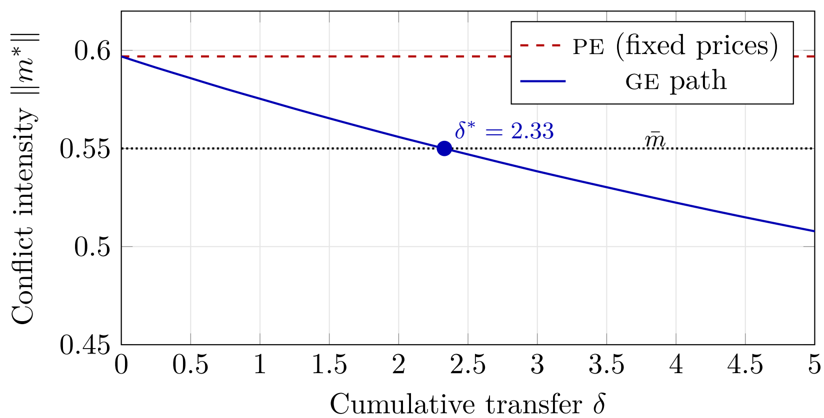

Figure 3: Equilibrium trajectory of a two-good conflict economy under commodity transfers. The GE path (solid) reduces conflict intensity through equilibrium price adjustment; the PE benchmark (dashed) shows that the same transfers at fixed prices have zero effect on militarization. Peace ($\|m^*\| \le \bar{m}$) is achieved at $\delta^* = 2.33$—29% of the donor's good-1 endowment—entirely through the price channel.

dersupplied relative to the cooperative optimum. The model quantifies this undersupply precisely.

Setup. Each donor $j \in \{1, \ldots, n\}$ independently chooses a transfer schedule $\tau_j : [0,1] \to \mathbb{R}^K_{\ge 0} \times \mathbb{R}^K_{\ge 0}$, with $\tau_j(0) = 0$, budget ceiling $\bar{\tau}_j$, efficiency parameter $\eta_j \ge 1$, and speed limit $M_j > 0$. Write $\mathcal{T}_j$ for the set of feasible transfer schedules for donor j—defined exactly as in section 4 with j-specific parameters—and $\mathcal{T} = \mathcal{T}_1 \times \cdots \times \mathcal{T}_n$ for the joint strategy space.

The aggregate transfer $\tau = \sum_j \tau_j$ induces the joint endowment path

$$\omega_A(t) = \omega_A^0 + \textstyle\sum_j \tau_{j,A}(t), \qquad \omega_B(t) = \omega_B^0 + \textstyle\sum_j \tau_{j,B}(t),$$

and the equilibrium path $\gamma : [0,1] \to \mathcal{D}$ is determined by the aggregate, exactly as in Definition 4.1. Donor j's payoff combines the shared equilibrium benefit and their private efficiency cost:

$$V_j(\tau_1, \ldots, \tau_n) \;\coloneqq\; \mathcal{V}_j(\gamma) \;-\; \int_0^1 C_j(\|\tau_j'(t)\|)\, dt,$$

where $\mathcal{V}_j(\gamma) = \int_\gamma \alpha_j$ is donor j's preference functional (section 4) and C_j is a smooth, strictly convex, increasing cost function with $C_j(0) = 0$.

7.1 Definition (Nash Equilibrium among Donors)
A profile $(\tau_1^*, \dots, \tau_n^*) \in \mathcal{T}$ *is a* Nash equilibrium *if for each* j*,*

$$V_j(\tau_j, \tau_{-j}^*) \ \leq \ V_j(\tau_j^*, \tau_{-j}^*) \quad \textit{for all } \tau_j \in \mathcal{T}_j.$$

7.2 Assumption (Quasi-concave payoffs)
For each donor j *and each fixed profile* $\tau_{-j} \in \mathcal{T}_{-j}$*, the payoff* $V_j(\tau_j, \tau_{-j})$ *is quasi-concave in* τ_j*.*

Quasi-concavity ensures that each donor's best response is well-behaved: the set of optimal transfer schedules is convex, and small changes in other donors' strategies produce small changes in each donor's best response—there are no discontinuous jumps between distant optima. Two sufficient conditions for Assumption 7.2 are worth noting. First, if all goods are *gross substitutes* for every actor, then equilibrium is globally unique at each endowment profile (Arrow–Block–Hurwicz), the map $\tau_j \mapsto \gamma$ is a smooth function, and V_j inherits quasi-concavity from the concavity of $-C_j$. Second, for *small transfers*—when $\|\sum_j \tau_j\|$ is small relative to initial endowments—the equilibrium path remains in a neighborhood of e_0 where the inverse function theorem guarantees $\tau_j \mapsto \gamma$ is approximately affine, and quasi-concavity holds. More generally, Assumption 7.2 holds locally near any regular point of $\mathcal{D}$; we impose it globally as a maintained assumption.

7.3 Theorem (Existence of Nash Equilibrium)
Under Assumption 7.2, the multi-donor game has a Nash equilibrium.

Proof. Each $\mathcal{T}_j$ is convex (budget and speed constraints are linear) and compact (the argument of Lemma 4.4 applies componentwise to each donor's strategy space). The payoff V_j is continuous in (τ_j, τ_{-j}) in the product topology τ of uniform convergence and weak-$*$ velocity convergence (the same argument as in Theorem 5.1 applies). By Assumption 7.2, each V_j is quasi-concave in τ_j for fixed τ_{-j}. The best-response correspondence $\beta_j(\tau_{-j}) = \operatorname{argmax}_{\tau_j \in \mathcal{T}_j} V_j(\tau_j, \tau_{-j})$ therefore has nonempty convex compact values and is upper hemicontinuous by the Berge maximum theorem. The joint best-response $\beta = \beta_1 \times \cdots \times \beta_n \colon \mathcal{T} \to 2^{\mathcal{T}}$ satisfies the hypotheses of the Fan–Glicksberg fixed-point theorem, and a fixed point is a Nash equilibrium. ∎

Pontryagin characterization of Nash equilibria. Each donor's problem, taking τ^*_{-j} as given, is an optimal control problem of the same form as in section 6. The aggregate velocity $u(t) = \sum_k \tau_k'(t)$ drives the state $\dot{x} = u$. Applying Theorem 6.1 to each donor's problem yields a coupled system.

7.4 Proposition (Nash–Pontryagin Conditions)

At a Nash equilibrium $(\tau_1^*, \ldots, \tau_n^*)$*, there exist costates* $\psi_j : [0,1] \to \mathbb{R}^{IK}$ *and multipliers* $\lambda_j \geq 0$ *such that for almost every* t*:*

(i) (Costates) $\dot{\psi}_j(t) = -Da_j(x^*(t))^\top u^*(t)$ *for each* j*, where* $u^* = \sum_k (\tau_k^*)'$ *is the aggregate velocity;*

(ii) (Switching functions) $S_j(t) = a_j(x^*(t)) + \psi_j(t)$ *evolves as* $\dot{S}_j = \iota_{u^*} d\alpha_j$*;*

(iii) (Optimality) $(\tau_j^*)'(t)$ *maximizes* $\langle S_j(t), u_j \rangle - C_j(\|u_j\|)$ *over* $\|u_j\| \leq M_j$*, giving* $C_j'(\|(\tau_j^*)'(t)\|) = \|S_j(t)\|$ *at interior optima;*

(iv) (Transversality) $\psi_j(1) = \lambda_j \nabla\phi(x^*(1))$ *with* $\lambda_j\, \phi(x^*(1)) = 0$*.*

The key feature of Proposition 7.4(i) is that *each donor's costate is driven by the aggregate velocity*, not their own contribution. Donor k's actions shift donor j's switching function through the state $x^*(t)$, and thereby alter j's optimal direction and pace. This is the channel through which strategic complementarities and substitutabilities propagate across donors.

The free-rider problem. When donors share the same preference (*common preferences*: $\alpha_j = \alpha$ and $C_j = C$ for all j), the equilibrium benefit $\mathcal{V}(\gamma)$ is a public good—all donors enjoy the same equilibrium trajectory—while each donor bears only their private efficiency cost. This creates a free-rider problem.

7.5 Proposition (Free-Rider Under Common Preferences)

Suppose $\alpha_j = \alpha$*,* $C_j = C$*, and* $M_j = M/n$ *for all* j *(symmetric donors with equal budget shares). At the symmetric Nash equilibrium, each donor contributes* $u_j^*(t) = u^*(t)/n$ *and the aggregate switching function satisfies*

$$C'(\|u^*(t)\|/n) \;=\; \|S^*(t)\|,$$

while the cooperative optimum solves

$$C'(\|u^{\mathrm{co}}(t)\|/n) \;=\; n\,\|S^{\mathrm{co}}(t)\|.$$

Since $n > 1$, the cooperative aggregate speed strictly exceeds the Nash aggregate speed at every t where $\|S\| > C'(0)$.

Proof. In the symmetric Nash equilibrium, all donors share the same costate (since $\alpha_j = \alpha$ and costates satisfy the same ODE driven by the common x^*) and hence $S_j(t) = S^*(t)$ for all j. The optimality condition (iii) gives $C'(\|u_j^*\|) = \|S^*\|$ for each j, and symmetry gives $u_j^* = u^*/n$, so $C'(\|u^*\|/n) = \|S^*\|$.

For the cooperative problem, the social planner maximizes $\sum_j V_j = n \cdot \mathcal{V}(\gamma) - n \cdot \int C(\|u^{\mathrm{co}}\|/n)\, dt$ (by symmetry, equal contributions are optimal). The Pontryagin condition for the aggregate gives $C'(\|u^{\mathrm{co}}\|/n) = n\|S^{\mathrm{co}}\|$.

Since $(C')^{-1}$ is strictly increasing and $n > 1$, we have $\|u^{\mathrm{co}}\|/n = (C')^{-1}(n\|S\|) > (C')^{-1}(\|S\|) = \|u^*\|/n$ at the same switching-function value, so $\|u^{\mathrm{co}}\| > \|u^*\|$. ∎

7.6 Corollary (Time-to-Peace)

With quadratic costs $C(s) = \kappa s^2$ ($\kappa > 0$) and common preferences, the cooperative aggregate speed is exactly n times the Nash aggregate speed at each t:

$$\|u^{\mathrm{co}}(t)\| \;=\; n\,\|u^*(t)\|.$$

Cooperative peacemaking therefore reaches any target on $\partial\mathcal{W}$ in $1/n$ the time of Nash equilibrium. In particular, if Nash equilibrium reaches peace at time $T_{\mathrm{Nash}} \le 1$, cooperation reaches peace by T_{Nash}/n; and Nash equilibrium may fail to achieve peace within the horizon $[0, 1]$ even when cooperation would.

Proof. For $C(s) = \kappa s^2$, $C'(s) = 2\kappa s$ and $(C')^{-1}(x) = x/(2\kappa)$. Nash: $\|u^*\|/n = \|S\|/(2\kappa)$, so $\|u^*\| = n\|S\|/(2\kappa)$. Cooperative: $\|u^{\mathrm{co}}\|/n = n\|S\|/(2\kappa)$, so $\|u^{\mathrm{co}}\| = n^2\|S\|/(2\kappa) = n\|u^*\|$. Since speed determines time-to-target by the inverse relationship $T = \mathrm{dist}/\mathrm{speed}$, $T_{\mathrm{co}} = T_{\mathrm{Nash}}/n$. ∎

Corollary 7.6 gives a sharp characterization of the cost of donor fragmentation: the time-to-peace scales linearly with the number of independent donors. Each additional donor, acting non-cooperatively, slows the collective response by a factor of n relative to a unified actor with the same aggregate budget. The mechanism is not budget inadequacy—the total resources are identical—but the failure to internalize the benefit that one's own contribution confers on the other $n-1$ donors. Institutional arrangements that align donor behavior (joint conditionality, burden-sharing agreements, multilateral coordination) restore the cooperative speed without requiring additional resources.

The cost of fragmentation. The mechanism is precise: each donor's costate is driven by the aggregate velocity (Proposition 7.4(i)), so each internalizes only $1/n$ of the social benefit of acceleration. The binding constraint on peacemaking in fragmented donor environments is therefore coordination failure, not budget adequacy—adding resources to a fragmented system amplifies the inefficiency. Coordination mechanisms (joint conditionality, multilateral institutions with unified disbursement authority) restore the cooperative speed without additional resources. The empirical implications are developed in section 8.

8 Conclusion

This paper develops a general equilibrium theory of optimal aid for ongoing conflict. The equilibrium manifold—the set of all Walrasian trade equilibria of a conflict economy, shown to be diffeomorphic to $\mathbb{R}^{IK}$—converts the donor's problem into a well-posed optimal navigation problem and guarantees existence of an optimal aid schedule. Three structural results characterize that solution: bang-bang disbursement under absorptive capacity constraints, a three-regime characterization of graduated versus front-loaded aid when delivery is costly, and a pure coordination failure in which n Nash donors deliver peace n times more slowly than a cooperative arrangement. Together they answer the question the title poses: "cheap" has a precise meaning (Proposition 6.7: $P^* \leq R^*$), and the savings come from steering the economy to a specific peaceful equilibrium via the market mechanism rather than eliminating war as an option altogether.

The framework generates a cluster of empirical implications that are, in principle, testable with existing data.

Donor fragmentation and conflict duration. The starkest prediction is Corollary 7.6: conditional on total aid volume and conflict characteristics, the number of active donors should be a significant positive predictor of conflict duration. The mechanism is not transactions costs, bureaucratic competition, or preference heterogeneity—it operates even when donors share identical objectives and resources. It is the public-good character of the equilibrium path: each donor discounts the marginal value of its contribution by the free-ride it gets on the others' transfers. With quadratic delivery costs the time-to-peace scales linearly in n, so the fragmentation effect is large and grows with the donor pool.

Preference alignment amplifies free-riding. Proposition 7.5 establishes the free-rider result under common preferences; the externality is strongest precisely when donors agree on which peace to pursue. This inverts the naive intuition that donor agreement should facilitate coordination. Empirically, high-salience conflicts—those attracting many donors with strong and convergent interests—should exhibit a *larger* fragmentation penalty than low-visibility conflicts where donor objectives are diffuse or heterogeneous. A conditional test would interact donor count with a measure of interest alignment (e.g., shared IGO membership, geographic proximity of donor capitals to the conflict, or common-language indicators) and expect a positive coefficient on the interaction.

Front-loading, graduation, and the urgency gradient. Proposition 6.3 gives a three-regime characterization of disbursement timing: pause when $\|S(t)\| \leq C'(0)$, graduate at interior speed when $C'(0) < \|S(t)\| < C'(M)$, and front-load at the capacity constraint when $\|S(t)\| \geq C'(M)$. The switching function $\|S(t)\|$ is the model's urgency measure—it tracks how much closer to peace the economy is relative to the donor's preferences. The model predicts that aid to conflicts exhibiting high urgency proxies should be front-loaded relative to those that are structurally distant from the peace boundary, controlling for total committed ODA. Crucially, Corollary 6.2 implies that graduated aid in the absence of observable delivery constraints is a symptom of institutional inefficiency, not optimal design.

Sequencing and composition, not just volume. The path-dependence result (Section 4) implies that two conflicts receiving identical total transfers over the same period should have different outcomes if the delivery sequence or commodity composition differs—a prediction that standard aid-effectiveness regressions, which control for ODA totals, cannot capture. In commodity dimensions, transfers in markets tightly linked to conflict financing (minerals, fuel, staple food) have outsized peace-promoting effects relative to their endowment value, because the equilibrium price adjustment in those markets directly alters conflict payoffs g_i and the resource drain r.

The GE savings: trade creation as a proxy. The savings $P^* < R^*$ measure the value of exploiting the equilibrium path rather than brute-forcing robust peace. A reduced-form proxy for these savings is the bilateral trade flow that commodity transfers generate between former belligerents—gains from trade that arise as a byproduct of the price adjustment through which aid operates. The Camp David Accords (1978) provide a suggestive illustration. Bilateral trade between Egypt and Israel was effectively zero throughout the conflict era (1952–1979); it appeared immediately after the treaty and grew steadily, reaching roughly eight percent of total U.S. aid to the two countries by the late 2000s (COW Dyadic Trade data; USAID Greenbook). The Carter administration's emphasis on trade creation as a strategic objective—including infrastructure investment to link the two economies—is consistent with the model's mechanism: the donor does not merely buy peace but restructures the economy so that peace is self-reinforcing through the market. The trade flows represent welfare from exchange that required no additional donor resources; in the language of the model, they are precisely the general equilibrium response through which the path cost P^* falls below the robust peace cost R^*.

Trade potential, peacemaking incidence, and peacemaking success. The Camp David example suggests a general prediction. Define a dyad's *trade potential* as the residual from a standard gravity model of bilateral trade: dyads that trade far less than distance, GDP, and institutional similarity predict have high unrealized gains from exchange. In the model, high trade potential means the conflict economy is far from its efficient Walrasian allocation, so commodity transfers generate large endogenous price adjustments—precisely the channel through which P^* falls below R^*. Two hypotheses follow. First, peacemaking interventions should be more likely in high-potential dyads, because the gap $R^* - P^*$ is larger and the cost of peace is lower: donors face a cheaper problem and are more willing to attempt it. Second, conditional on intervention, peacemaking should be more successful in high-potential dyads. The mechanism is not that the conflict is milder but that the market has more leverage: for a fixed donor budget and time horizon, the equilibrium path covers more distance toward $\partial\mathcal{W}$ when the price channel is powerful than when the economy is already near its trade-efficient allocation and the donor is effectively in the brute-force R^* regime. Both hypotheses are testable in dyadic data on third-party peacemaking episodes, using gravity residuals computed from pre-conflict trade patterns.

Loans, grants, and peace durability. A loan traces a loop in endowment space; whether the economy returns to the pre-aid equilibrium after repayment depends on the monodromy of the natural projection π (section D). When the equilibrium manifold has nontrivial monodromy—when the economy passes through a region of multiple equilibria during the loan period, so that the "return trip" lands on a different equilibrium branch—peace achieved by conditional transfers can be durable even after repayment: the economy settles on a peaceful equilibrium that persists without continued aid. When monodromy is trivial, loan-based peace collapses at repayment. The prediction is testable against data on conflict recurrence following debt relief, conditional aid programs, and peace agreements with donor-enforced economic conditionality: recurrence rates following grant-financed peace should be systematically lower than those following loan-financed peace of equivalent face value, conditional on settlement type.

The framework has natural extensions. The present model restricts the donor's instrument to commodity transfers—endowment shifts that operate through the price mechanism. In practice, third parties also reshape the conflict primitive itself: military aid alters the cost functions c_i and battlefield functions s_i of Remark 2.3, arms embargoes constrain the feasible militarization set, and security guarantees change the payoff to fighting directly. Extending the donor's control to the triple $(m^*, (g_i), r)$—allowing the third party to jointly navigate the equilibrium manifold and deform the war region $\mathcal{W}$—would capture the interaction between economic and military aid and characterize when each instrument is substitutable or complementary. Production, incomplete information about the conflict primitive, and stochastic equilibrium paths are further natural directions. A fully dynamic bargaining model that endogenizes $\mathcal{W}$ alongside the donor's navigation problem would integrate the Fearon barriers to bargaining with the GE mechanism studied here. The present paper establishes that the mechanism exists and is quantitatively significant; those extensions determine when it dominates.

A Verification of Remark 2.3

We verify that the battlefield model of Remark 2.3 satisfies Assumption 2.2: the game Γ_{p,ω_0} has a unique Nash equilibrium $m^*(p,\omega_0)$ that is smooth in (p,ω_0), and the induced triple $(m^*, (g_i)_{i\in I_B}, r)$ satisfies contentious Walras's law.

Strict concavity and best responses. Fix $(p,\omega_0) \in \mathcal{P} \times \Omega_0$. Each player $i \in I_B$ maximizes $G_i(m_i, m_{-i}) \coloneqq p \cdot [s_i(m,p,\omega_0) - c_i(m_i,p,\omega_0)]$. By condition (iii) of Remark 2.3, for each k,

$$\frac{\partial^2 c_i^k}{\partial m_i^2} - \frac{\partial^2 s_i^k}{\partial m_i^2} > \left| \frac{\partial^2 s_i^k}{\partial m_i \, \partial m_{-i}} \right| \geq 0,$$

so $\partial^2 G_i / \partial m_i^2 = \sum_k p^k [\partial^2 s_i^k/\partial m_i^2 - \partial^2 c_i^k / \partial m_i^2] < 0$: G_i is strictly concave in m_i. Combined with coerciveness (condition (iv)), each player has a unique best response $B_i(m_{-i})$, which is smooth by the implicit function theorem applied to the first-order condition. The slope of B_i satisfies

$$|B_i'(m_{-i})| \;=\; \frac{\left| \displaystyle\sum_k p^k \frac{\partial^2 s_i^k}{\partial m_i \, \partial m_{-i}} \right|}{\displaystyle\sum_k p^k \Big[\frac{\partial^2 c_i^k}{\partial m_i^2} - \frac{\partial^2 s_i^k}{\partial m_i^2} \Big]} \;<\; 1,$$

where the bound follows from condition (iii) and the triangle inequality.

Uniqueness. Suppose $\widetilde{m}$ and $\overline{m}$ are both Nash equilibria. By the mean-value theorem applied to B_A and B_B:

$$\begin{aligned} |\overline{m}_A - \widetilde{m}_A| &= |B_A'(\xi)| \cdot |\overline{m}_B - \widetilde{m}_B| < |\overline{m}_B - \widetilde{m}_B|, \\ |\overline{m}_B - \widetilde{m}_B| &= |B_B'(\xi')| \cdot |\overline{m}_A - \widetilde{m}_A| < |\overline{m}_A - \widetilde{m}_A|, \end{aligned}$$

for intermediate points ξ, ξ', a contradiction unless $\overline{m} = \widetilde{m}$. Smoothness of $m^*(p,\omega_0)$ follows from the implicit function theorem applied to the joint first-order system; the Jacobian is nonsingular because $|B_i'| < 1$ implies strict diagonal dominance.

Contentious Walras's law. With g_i and r as defined in Remark 2.3,

$$\sum_{i\in I_B} g_i + p\cdot r \;=\; p\cdot \sum_{i\in I_B} \big[s_i(m^*,p,\omega_0) - c_i(m_i^*,p,\omega_0) \big] + p\cdot \sum_{i\in I_B} \big[c_i(m_i^*,p,\omega_0) - s_i(m^*,p,\omega_0) \big] \;=\; 0.$$

■

B Proof of Proposition 3.2

The following proof constructs explicit coordinates on the equilibrium manifold—roughly, a way to label every possible equilibrium by a point in Euclidean space, so that the donor's optimization problem can be stated as a standard control problem on $\mathbb{R}^{IK}$. We adapt the coordinate-map technique of Balasko (2011) to the contentious economy. The proof has two steps: we first show $\mathcal{D}$ is a smooth submanifold of $\mathcal{P} \times \Omega$ of the right dimension, then exhibit an explicit diffeomorphism $\mathcal{D} \cong \mathbb{R}^{IK}$.

Notation. Recall $I_0 = \{A, B, C\}$, $\Omega_0 = \mathbb{R}^{I_0 \times K}$, and write $\omega_0 = (\omega_i)_{i \in I_0} \in \Omega_0$ for the endowments of the non-world actors. For $p \in \mathcal{P}$ and $i \in I_0$, write $\widehat{p}$ for the first $K-1$ components of p (recall $p^K = 1$) and $\widehat{\omega}_i$ for the first $K-1$ components of ω_i, so that $p \cdot \omega_i = \widehat{p} \cdot \widehat{\omega}_i + \omega_i^K$.

Step 1: $\mathcal{D}$ is a smooth submanifold of dimension IK. Define the aggregate excess demand map

$$z \colon \mathcal{P} \times \Omega \longrightarrow \mathbb{R}^{K-1}$$

whose k-th component ($k < K$) is the excess demand for good k: total demand minus total supply in that good, with $m^*(p, \omega_0)$ substituted for militarization throughout. Formally,

$$z^k(p, \omega) \;:=\; \sum_{i \in I} \widetilde{f}_i^k(p, \omega) \;-\; \sum_{i \in I} \omega_i^k \;+\; r^k(p, \omega_0),$$

where $\widetilde{f}_i$ denotes contentious demand (uniform across all i, with $g_i \equiv 0$ for $i \notin I_B$). The map z is smooth: each $\widetilde{f}_i$ is smooth by Assumptions 2.1 and 2.2, and r is smooth by Assumption 2.2. Only $K-1$ components of excess demand appear because the contentious Walras's law—which holds irrespective of the optimality of militarization—ensures the K-th market clears whenever the first $K-1$ do.

We claim $0 \in \mathbb{R}^{K-1}$ is a regular value of z, so that $\mathcal{D} = z^{-1}(0)$ is a smooth submanifold of codimension $K-1$ in $\mathcal{P} \times \Omega$, hence of dimension $(K-1) + IK - (K-1) = IK$. To verify regularity, it suffices to show the partial Jacobian $\partial z / \partial p$ has rank $K-1$ at every point of $\mathcal{D}$. By strict quasiconcavity (Assumption 2.1), each actor's Slutsky matrix is negative semi-definite on the budget hyperplane and negative definite on its kernel; the aggregate Slutsky matrix therefore has the same property. The militarization terms contribute smooth corrections to $\partial z / \partial p$ but do not affect its rank, since

m^* is smooth and the militarization-induced shifts are dominated by the curvature conditions of Assumption 2.2. The rank condition follows exactly as in the pure exchange case; see Balasko (2011, Proposition 4.9).

Step 2: $\mathcal{D} \cong \mathbb{R}^{IK}$ via explicit coordinates. The idea is to label each equilibrium not by its endowments but by its prices and incomes. This works because, in a Walrasian economy, prices and incomes—together with utility maximization and market clearing—determine everything else: each agent's demand, the implied endowments, and the Nash militarization. The construction below makes this bijection explicit and verifies that it is smooth in both directions.

We construct smooth maps $\Phi\colon \mathcal{D} \to \mathcal{P} \times \mathbb{R}^I \times \mathbb{R}^{(K-1)(I-1)}$ and $\theta\colon \mathcal{P} \times \mathbb{R}^I \times \mathbb{R}^{(K-1)(I-1)} \to \mathcal{D}$ with $\Phi \circ \theta = \mathrm{id}$ and $\mathrm{im}(\theta) = \mathcal{D}$, which establishes that Φ is a diffeomorphism onto its codomain. Since $\mathcal{P} \times \mathbb{R}^I \times \mathbb{R}^{(K-1)(I-1)}$ has dimension $(K-1) + I + (K-1)(I-1) = IK$, this gives $\mathcal{D} \cong \mathbb{R}^{IK}$.

The forward map Φ. Define

$$\Phi(p, \omega) \;\coloneqq\; \Big(p,\; (p \cdot \omega_i)_{i \in I},\; (\widehat{\omega}_i)_{i \in I_0}\Big).$$

That is, Φ records the price vector, the income of every actor, and the non-numeraire endowment components of every actor except W. It is evidently smooth.

The inverse map θ. Given $(p,\, w = (w_i)_{i \in I},\, \widehat{\omega}_0 = (\widehat{\omega}_i)_{i \in I_0}) \in \mathcal{P} \times \mathbb{R}^I \times \mathbb{R}^{(K-1)(I-1)}$, define $\theta(p, w, \widehat{\omega}_0) = (p, \omega)$ where $\omega = (\omega_i)_{i \in I}$ is constructed as follows.

(1) For each $i \in I_0$: set $\widehat{\omega}_i$ as given, and set

$$\omega_i^K \;\coloneqq\; w_i - \widehat{p} \cdot \widehat{\omega}_i.$$

This recovers the numeraire component of ω_i from the budget identity $p \cdot \omega_i = \widehat{p} \cdot \widehat{\omega}_i + \omega_i^K = w_i$.

(2) Compute the Nash equilibrium militarization $m^*(p, \omega_0)$, which depends only on p and $\omega_0 = (\omega_i)_{i \in I_0}$, now fully determined.

(3) Set ω_W by the contentious market-clearing identity:

$$\omega_W \;\coloneqq\; \sum_{i \in I} \widetilde{f}_i(p,\, w,\, \widehat{\omega}_0) \;+\; r(p, \omega_0) \;-\; \sum_{i \in I_0} \omega_i.$$

Each step is smooth in the inputs, so θ is smooth.

$\Phi \circ \theta = \mathrm{id}$. By construction, $\theta(p, w, \widehat{\omega}_0)$ produces a pair (p, ω) with $p \cdot \omega_i = w_i$ for all $i \in I_0$ (by step 1) and $\widehat{\omega}_i$ as given for $i \in I_0$. Applying Φ recovers $(p, (p \cdot \omega_i)_{i \in I}, (\widehat{\omega}_i)_{i \in I_0}) = (p, w, \widehat{\omega}_0)$. ✓

$\mathrm{im}(\theta) \subseteq \mathcal{D}$. We must show the output of θ satisfies market clearing. By step 3,

$$\sum_{i \in I} \omega_i = \sum_{i \in I_0} \omega_i + \omega_W = \sum_{i \in I} \widetilde{f}_i + r(p, \omega_0),$$

which is precisely the contentious market-clearing condition. The Nash condition is satisfied because m^* was computed as the Nash equilibrium in step 2. ✓

$\mathcal{D} \subseteq \mathrm{im}(\theta)$. For any $(p, \omega) \in \mathcal{D}$, set $w_i = p \cdot \omega_i$ and $\widehat{\omega}_i = (\omega_i^1, \ldots, \omega_i^{K-1})$ for each i. Then $\theta(p, w, \widehat{\omega}_0) = (p, \omega)$ by construction. ✓

Since $\Phi \circ \theta = \mathrm{id}$ and $\mathrm{im}(\theta) = \mathcal{D}$, the map $\Phi|_{\mathcal{D}}$ is a smooth bijection onto $\mathcal{P} \times \mathbb{R}^I \times \mathbb{R}^{(K-1)(I-1)}$ with smooth inverse θ, hence a diffeomorphism. This completes the proof. ■

C Proof of Proposition 3.3

The proof shows that if the donor's transfers remain bounded, so do the resulting equilibrium prices and allocations—prices cannot escape to the boundary of the price simplex under finite intervention. This is what converts the donor's budget constraint into compactness of the feasible set in the existence proof of section 5.

Smoothness of π is immediate: it restricts a coordinate projection to the smooth submanifold $\mathcal{D}$. For properness, let $K \subseteq \Omega$ be compact and let $\{(p^n, \omega^n)\}_{n=1}^{\infty} \subseteq \pi^{-1}(K)$ be a sequence. Since $\omega^n \in K$, the endowments $\{\omega^n\}$ are bounded. We claim prices $\{p^n\}$ are also bounded.

Suppose for contradiction that some $p^{k,n} \to \infty$ along a subsequence (since $p^{K,n} = 1$ throughout, escaping to infinity means some $p^{k,n} \to \infty$ for $k < K$). Market clearing at $(p^n, \omega^n) \in \mathcal{D}$ requires

$$\sum_{i \in I} \widetilde{f}_i^k(p^n, \omega^n) \;=\; \sum_{i \in I} \omega_i^{k,n} - r^k(p^n, \omega_0^n).$$

As $p^{k,n} \to \infty$, good k becomes infinitely expensive relative to all other goods. Under Assumption 2.1, demand for good k collapses: $\widetilde{f}_i^k(p^n, \omega^n) \to 0$ for each i (for belligerents, contentious demand is Walrasian demand at adjusted wealth, and the same argument applies since adjusted wealth is bounded when endowments and m^* are bounded). The left-hand side therefore goes to zero. The right-hand side, however, remains bounded away from zero: $\sum_i \omega_i^{k,n}$ is bounded (since $\omega^n \in K$) and $r^k(p^n, \omega_0^n)$ is bounded (since r is smooth and its arguments are bounded by Assumption 2.2). This contradicts market clearing. The case $p^{k,n} \to 0$ is symmetric: demand diverges while supply is bounded.

Hence $\{p^n\}$ is bounded. Combined with $\{\omega^n\}$ bounded, the sequence $\{(p^n, \omega^n)\}$ has a convergent subsequence; its limit lies in $\mathcal{D}$ since $\mathcal{D} = z^{-1}(0)$ is closed. This adapts Balasko (2011, Proposition 4.4) to the contentious economy; the militarization corrections are bounded by smoothness of m^* and r (Assumption 2.2). ■

D Monodromy and Loan Reversibility

This section establishes when a loan—a transfer that is eventually repaid—returns the economy to its original equilibrium. The answer has policy content: if the economy passes through a region of multiple equilibria during the loan period, repayment of the loan need not restore the original prices, militarization levels, or welfare. The economy may settle on a different equilibrium branch even after all resources are returned. The formal tool is the monodromy of π, which characterizes path-dependent equilibrium selection in terms of the topology of the equilibrium manifold above the endowment space.

The path-dependent equilibrium selection described for loans in section 4 has a precise topological characterization in terms of the monodromy of π.

The discriminant locus. Let $\Sigma \subset \Omega$ denote the set of critical values of π—points $\omega \in \Omega$ at which $D\pi|_e$ fails to be an isomorphism for some $e \in \pi^{-1}(\omega)$, so that equilibria bifurcate or merge. By Sard's theorem applied to the smooth map $\pi\colon \mathcal{D} \to \Omega$, the set Σ has measure zero in Ω.

Covering space structure. A covering space is a continuous map that "unfolds" the multi-valued equilibrium correspondence into something single-valued and trackable. Above each regular endowment profile, there are finitely many equilibria; the covering structure ensures that as endowments change smoothly, each equilibrium moves smoothly and can be followed individually—no equilibrium suddenly appears, disappears, or collides with another, as long as the endowment path avoids the critical values Σ. This is the mathematical structure that makes path-dependent equilibrium selection well-defined: given a starting equilibrium and a smooth path of endowment changes, there is a unique equilibrium path that continues it.

The projection $\pi\colon \pi^{-1}(\Omega \setminus \Sigma) \to \Omega \setminus \Sigma$ is a covering space. We verify the three conditions.

Fibers are finite. Fix $\omega \in \Omega \setminus \Sigma$. The fiber $\pi^{-1}(\omega)$ is discrete: at each $e \in \pi^{-1}(\omega)$, regularity means $D\pi|_e$ is an isomorphism, so the implicit function theorem supplies an open neighborhood $U_e \ni e$ on which π is injective, giving $U_e \cap \pi^{-1}(\omega) = \{e\}$. The fiber is also compact: by Proposition 3.3, π is proper, so $\pi^{-1}(\{\omega\})$ is compact. A discrete compact set is finite.

Local homeomorphisms. At each $e \in \pi^{-1}(\Omega \setminus \Sigma)$, the implicit function theorem gives a neighborhood U_e of e and $V_e = \pi(U_e)$ of $\pi(e)$ on which $\pi|_{U_e}\colon U_e \to V_e$ is a diffeomorphism.

Even covering. Fix $\omega \in \Omega \setminus \Sigma$ with fiber $\{e_1, \ldots, e_n\}$. The sets $U_{e_1}, \ldots, U_{e_n}$ from the previous step can be chosen pairwise disjoint (shrink each to exclude the other fiber points, which are isolated). Set $V = V_{e_1} \cap \cdots \cap V_{e_n}$ and $\widetilde{U}_j = \pi^{-1}(V) \cap U_{e_j}$. Then $\pi^{-1}(V) = \widetilde{U}_1 \sqcup \cdots \sqcup \widetilde{U}_n$, and each $\pi|_{\widetilde{U}_j}\colon \widetilde{U}_j \to V$ is a homeomorphism. So V is an evenly covered neighborhood of ω.

Monodromy action. The question of whether a loan returns the economy to its original equilibrium reduces to a topological question about the path of transfers: does the loop in endowment space, when "unfolded" via the covering structure, come back to the same sheet? The answer is encoded in the fundamental group of the space of regular endowments, which classifies loops up to continuous deformation.

Formally: the fundamental group $\pi_1(\Omega \setminus \Sigma, \omega^0)$ acts on the fiber $\pi^{-1}(\omega^0)$ by monodromy: a loop $[\ell] \in \pi_1(\Omega \setminus \Sigma, \omega^0)$ sends $e \in \pi^{-1}(\omega^0)$ to the endpoint of the unique lift of ℓ to $\mathcal{D}$ starting at e. This is well-defined (lifts of homotopic loops share endpoints) and defines a group homomorphism from $\pi_1(\Omega \setminus \Sigma, \omega^0)$ to the symmetric group on $\pi^{-1}(\omega^0)$.

Loan reversibility. A loan has transfer schedule τ with $\tau(1) = 0$, so its endowment path $\omega(\cdot)$ is a loop in Ω based at ω^0. If τ avoids Σ—which holds generically by Sard—the induced aid schedule γ is the unique lift of this loop in $\mathcal{D}$ starting at e_0. The loan returns the economy to e_0 if and only if the monodromy action of $[\tau] \in \pi_1(\Omega \setminus \Sigma, \omega^0)$ fixes e_0. Loans whose transfer paths wind around a point of Σ may permute the fiber $\pi^{-1}(\omega^0)$, landing on a different equilibrium branch even after all resources are restored—a change in prices, militarization, and welfare that persists beyond the repayment date.

References

Acemoglu, Daron, Mikhail Golosov, Aleh Tsyvinski, and Pierre Yared, "A Dynamic Theory of Resource Wars," *Quarterly Journal of Economics*, 2012, *127* (1), 283–331.

Acharya, Arnab, Ana Teresa Fuzzo de Lima, and Mick Moore, "Proliferation and Fragmentation: Transactions Costs and the Value of Aid," *Journal of Development Studies*, 2006, *42* (1), 1–21.

Balasko, Yves, *Foundations of the Theory of General Equilibrium*, Boston: Academic Press, 1988.

_ , *The Equilibrium Manifold: Postmodern Developments in the Theory of General Economic Equilibrium*, Cambridge, MA: MIT Press, 2011.

Bazzi, Samuel and Christopher Blattman, "Economic Shocks and Conflict: Evidence from Commodity Prices," *American Economic Journal: Macroeconomics*, 2014, *6* (4), 1–38.

Bevia, Carmen and Luis C. Corchón, "Peace Agreements Without Commitment," *Games and Economic Behavior*, 2010, *68* (2), 469–487.

Bó, Ernesto Dal and Pedro Dal Bó, "Workers, Warriors, and Criminals: Social Conflict in General Equilibrium," *Journal of the European Economic Association*, 2011, *9* (4), 646–677.

Burnside, Craig and David Dollar, "Aid, Policies, and Growth," *American Economic Review*, 2000, *90* (4), 847–868.

Caselli, Francesco, Massimo Morelli, and Dominic Rohner, "The Geography of Interstate Resource Wars," *Quarterly Journal of Economics*, 2015, *130* (1), 267–315.

Chassang, Sylvain and Gerard Padró i Miquel, "Conflict and Deterrence Under Strategic Risk," *Quarterly Journal of Economics*, 2010, *125* (4), 1821–1858.

Collier, Paul and Anke Hoeffler, "Aid, Policy and Peace: Reducing the Risks of Civil Conflict," *Defence and Peace Economics*, 2002, *13* (6), 435–450.

_ **, Lani Elliott, Håvard Hegre, Anke Hoeffler, Marta Reynal-Querol, and Nicholas Sambanis**, *Breaking the Conflict*

Trap: Civil War and Development Policy, Washington, DC: World Bank and Oxford University Press, 2003.

de Ree, Joppe and Eleonora Nillesen, "Aiding Violence or Peace? The Impact of Foreign Aid on the Risk of Civil Conflict in Sub-Saharan Africa," *Journal of Development Economics*, 2009, *88* (2), 301–313.

Debreu, Gérard, "Economies with a Finite Set of Equilibria," *Econometrica*, 1970, *38* (3), 387–392.

Doyle, Michael W. and Nicholas Sambanis, "International Peacebuilding: A Theoretical and Empirical Analysis," *American Political Science Review*, 2000, *94* (4), 779–801.

Dube, Oeindrila and Juan F. Vargas, "Commodity Price Shocks and Civil Conflict: Evidence from Colombia," *Review of Economic Studies*, 2013, *80* (4), 1384–1421.

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

_ , "Why Do Some Civil Wars Last So Much Longer than Others?," *Journal of Peace Research*, 2004, *41* (3), 275–301.

Garfinkel, Michelle R., Stergios Skaperdas, and Constantinos Syropoulos, "Globalization and Domestic Conflict," *Journal of International Economics*, 2008, *76* (2), 296–308.

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

Hodler, Roland and Hadi Yektaş, "All-Pay War," *Games and Economic Behavior*, 2012, *74* (2), 526–540.

Knack, Stephen and Aminur Rahman, "Donor Fragmentation and Bureaucratic Quality in Aid Recipients," *Journal of Development Economics*, 2007, *83* (1), 176–197.

Konrad, Kai A., *Strategy and Dynamics in Contests*, Oxford: Oxford University Press, 2009.

Liberzon, Daniel, *Calculus of Variations and Optimal Control Theory: A Concise Introduction*, Princeton, NJ: Princeton University Press, 2011.

Martin, Philippe, Thierry Mayer, and Mathias Thoenig, "Make Trade Not War?," *Review of Economic Studies*, 2008, *75* (3), 865–900.

Nunn, Nathan and Nancy Qian, "US Food Aid and Civil Conflict," *American Economic Review*, 2014, *104* (6), 1630–1666.

Olson, Mancur and Richard Zeckhauser, "An Economic Theory of Alliances," *The Review of Economics and Statistics*, 1966, *48* (3), 266–279.

Polachek, Solomon William, "Conflict and Trade," *Journal of Conflict Resolution*, 1980, *24* (1), 55–78.

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

Rajan, Raghuram G. and Arvind Subramanian, "Aid and Growth: What Does the Cross-Country Evidence Really Show?," *The Review of Economics and Statistics*, 2008, *90* (4), 643–665.

Ross, Michael L., "What Do We Know About Natural Resources and Civil War?," *Journal of Peace Research*, 2004, *41* (3), 337–356.

Savun, Burcu and Daniel C. Tirone, "Exogenous Shocks, Foreign Aid, and Civil War," *International Organization*, 2012, *66* (3), 363–393.

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

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