# Making Peace on the Cheap

Robert J. Carroll*

March 15, 2026

**Abstract**

This paper develops a general equilibrium theory of optimal foreign aid for ongoing conflict. A donor's transfer schedule induces a Lipschitz path on the equilibrium manifold $\mathcal{D}$—the set of all Walrasian equilibria of a conflict economy, shown to be diffeomorphic to $\mathbb{R}^{IK}$—converting a navigation problem on a multi-valued equilibrium correspondence into a well-posed smooth optimal control problem. The donor's problem is to navigate from an initial war equilibrium to the peace boundary at minimum cost. The title refers to a precise inequality: the general equilibrium path cost $P^*$ is weakly less than the robust peace cost $R^*$—the minimum transfer that eliminates all war equilibria at the terminal endowment—with strict savings whenever the terminal fiber contains both war and peace equilibria, a generic condition. The savings arise because the market price mechanism does part of the work: the donor need only steer to one peaceful equilibrium, not eliminate war as an option. Under an absorptive capacity constraint, the optimal schedule always disburses at maximum rate—gradualism is never optimal by choice. With strictly convex delivery costs, three regimes emerge as a function of an endogenous urgency measure: pause, graduate, or front-load. With $n$ symmetric donors in Nash equilibrium, peace is achieved $n$ times more slowly than under cooperation, without any change in total resources—a pure coordination failure that provides a resource-neutral rationale for multilateral peacemaking institutions.

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

# 1 Introduction

When a donor transfers commodities to a conflict economy, equilibrium prices adjust. Those price adjustments alter the conflict payoffs of both belligerents continuously along the path of the transfer schedule, potentially achieving peace at strictly lower cost than any approach that ignores them. Little work formalizes this mechanism or asks how a donor should exploit it optimally. This paper does.

The empirical case for aid as a peacemaking instrument is robust: 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. Even the sign of the effect depends on design: undirected food transfers can prolong existing conflicts (Nunn and Qian, 2014), a finding that reverses under improved identification (Mary, 2026). Existing theoretical frameworks treat aid as a fiscal instrument—government budget augmentation that raises the opportunity cost of rebellion (Savun and Tirone, 2012; Collier and Hoeffler, 2002)—or as a political commitment device (Chassang and Padró i Miquel, 2010). Neither uses the market structure of the conflict economy.

In the canonical framework for third-party economic intervention, a donor transfers a lump sum to one or both belligerents, changing their outside options at *fixed prices*; peace is achievable if and only if the donor can close the bargaining gap (Bevia and Corchón, 2010). The theory is clean, but it omits the price channel. When aid takes the form of commodity transfers, equilibrium prices adjust; those adjustments alter the conflict payoffs of both parties continuously along the path of the transfer schedule, potentially achieving peace at strictly lower total cost than any partial-equilibrium calculation would predict. 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. Dal Bó and Dal Bó (2011) formally establish, and Dube and Vargas (2013) empirically confirm in Colombia, that commodity prices affect conflict incentives through the opportunity cost and rapacity channels—exactly the general equilibrium price mechanism our model formalizes and a donor can exploit. This paper asks the missing question: 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?

I model a conflict economy as a *dual competitive equilibrium* (DCE)—

a Walrasian trade equilibrium in which belligerents have wealth adjusted by a smooth abstract conflict primitive $(m^*, g_i, r)$ encoding Nash militarization, net conflict payoffs, and aggregate resource drain. The set of all DCEs forms the *equilibrium manifold* $\mathcal{D}$ (Balasko, 2011), shown to be diffeomorphic to $\mathbb{R}^{IK}$. The manifold structure does three specific things for the analysis. First, the diffeomorphism $\mathcal{D} \cong \mathbb{R}^{IK}$ converts a potentially ill-posed navigation problem on a multi-valued equilibrium correspondence into a well-posed smooth optimal control problem; without it, existence of an optimal aid path is not guaranteed. Second, properness of the natural projection $\pi\colon \mathcal{D} \to \Omega$ makes generic equilibrium multiplicity a theorem: the terminal fiber $\pi^{-1}(\omega^*)$ generically contains both war and peace equilibria, which is precisely what makes the savings inequality $P^* < R^*$ strict rather than merely weak. Third, the shared equilibrium path through which the multi-donor coordination failure operates is only well-defined as a smooth curve on $\mathcal{D}$. A donor's transfer schedule—an absolutely continuous path in endowment space—induces a Lipschitz path on $\mathcal{D}$; the *war region* $\mathcal{W} \subset \mathcal{D}$ is defined by a smooth conflict condition, and the donor's problem is to navigate the economy from the initial war equilibrium $e_0 \in \mathcal{W}$ to the peace boundary $\partial\mathcal{W}$ at minimum cost. The title has a precise content. Define the *robust peace cost* $R^*(e_0)$ as the minimum endowment transfer such that every equilibrium at the terminal endowment profile is peaceful. Define the *GE path cost* $P^*(e_0)$ as the minimum total transfer variation along any path in $\mathcal{D}$ from $e_0$ to $\partial\mathcal{W}$. Proposition 6.7 establishes $P^*(e_0) \leq R^*(e_0)$, with strict inequality whenever the terminal fiber $\pi^{-1}(\omega^*)$ contains both war and peace equilibria—a generic condition by properness of the natural projection $\pi\colon \mathcal{D} \to \Omega$. The savings from the equilibrium path arise because the donor need only steer the economy to *one* peaceful equilibrium through the market mechanism, without eliminating war as a possible outcome at the terminal endowment.

Three structural results characterize the optimal schedule. First, 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. Second, under an absorptive capacity constraint $\|u\| \leq M$, 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). Third, with strictly convex delivery costs $C$, three disbursement regimes emerge as a function of an endogenous urgency measure $\|S(t)\|$: 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 $S(t)$—the switching function of the associated Pontryagin problem—evolves along the equilibrium path, making disbursement timing endogenous to the economy's trajectory toward peace. An optimal aid schedule always exists (Theorem 5.1).

The paper connects two literatures. *General equilibrium and conflict*: Garfinkel et al. (2008) embed conflict in a static Walrasian model and study how trade regimes affect conflict intensity; Dal Bó and Dal Bó (2011) identify the Stolper-Samuelson mechanism through which commodity prices affect conflict incentives via opportunity cost and rapacity; Dube and Vargas (2013) provide empirical confirmation in Colombia. The present paper studies dynamic optimal *intervention* on the equilibrium manifold—the donor's problem none of these frameworks addresses. Bevia and Corchón (2010) compute the minimum static transfer for peace at fixed prices; the GE path cost $P^*$ is strictly lower, with the savings attributable to equilibrium price adjustment along the transfer schedule—the channel their framework omits. Acemoglu et al. (2012) show that the path of resource prices shapes war incentives in a dynamic setting, motivating the dynamic formulation; Caselli et al. (2015) document that the distribution of endowments across potential belligerents shapes conflict incidence. I take the existence of the war region as given; for rationalist foundations see Fearon (1995) and Powell (2006). *Aid effectiveness and donor coordination*: the empirical evidence for aid's effectiveness in shortening conflicts (de Ree and Nillesen, 2009; Doyle and Sambanis, 2000) motivates the donor's problem I formalize; Nunn and Qian (2014) find that undirected food aid prolongs conflict, but Mary (2026) overturns this under improved identification—the sensitivity of the finding to design choices makes the optimality of the transfer schedule the central question. Collier and Hoeffler (2002) identify equilibrium price adjustment as an informal mechanism linking aid to conflict reduction, and I provide the formal theory for this channel. The donor fragmentation documented by Knack and Rahman (2007) and Acharya et al. (2006) receives a new theoretical explanation: even without administrative frictions or talent-poaching, $n$ non-cooperating donors deliver peace $n$ times more slowly than a cooperative arrangement, with the free-rider externality operating through the shared equilibrium path rather than through administrative channels.

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 $\omega_i$ 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.

**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. We characterize conflict equilibrium through three smooth objects on $\mathcal{P} \times \Omega_0$:

- 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 $\omega_0$ but not on $\omega_W$.

---

**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 ties the budget impacts $(g_i)$ to the resource drain $r$, ensuring that aggregate spending power equals aggregate endowment value net of the drain—a consistency condition that will allow us to apply Walras's lemma in equilibrium. 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);*

(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); and*

(iv) *there exists* $k$ *such that* $\lim_{m_i \to \infty} \dfrac{\partial c_i^k}{\partial m_i} = \infty$ *(some costs eventually become prohibitive);*

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

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

*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

$$f_i(p, w) \coloneqq \operatorname*{argmax}_{p \cdot x \leq w} \; u_i(x).$$

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.

---

**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*. Because it is diffeomorphic to $\mathbb{R}^{IK}$, it is connected, path-connected, simply connected, 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*

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

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: $\pi$ restricts a coordinate projection to the smooth submanifold $\mathcal{D}$. The content is properness, which says equilibrium prices cannot escape to zero or infinity when endowments stay bounded. If any price diverged, strict monotonicity of preferences would force demand below supply in the corresponding market, violating clearing. The argument is the same as in pure exchange (Balasko, 2011, Proposition 4.4); the militarization terms remain bounded because $m^*$ is smooth (Assumption 2.2).

## 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 $\tau$ 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 $\eta$ 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, $\eta$ 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 $\pi^{-1}(\omega(t)) \subseteq \mathcal{D}$. A transfer schedule therefore traces a path in endowment space $\Omega$; 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 differential $D\pi|_{e_0}$ is an isomorphism—the inverse function theorem supplies a smooth local inverse to $\pi$ near $\omega^0$. 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 $\pi$, where equilibria can bifurcate or merge; by Sard's theorem, the set of critical values has measure zero in $\Omega$, so a generic transfer schedule 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* $\gamma$ *is* induced by *the transfer schedule* $\tau$ *if* $\pi(\gamma(t)) = \omega(t)$ *for all* $t$*, where* $\omega(\cdot)$ *is given by* (1).

---

The economic primitive is the transfer schedule $\tau$; the aid schedule $\gamma$ is the induced equilibrium trajectory. Because $\tau$ 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 $\gamma$—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*

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

We treat $\phi$ 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 $\gamma$ 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 $\tau$ for the transfer schedule inducing $\gamma$, we require $\eta\big(\tau_A^k(t)+\tau_B^k(t)\big) \leq \bar{\tau}^k$ 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: $\|\gamma'(t)\| \leq M$ 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 (Rajan and Subramanian, 2008).

Write $\mathcal{F}$ for the set of all feasible aid schedules. We assume $\mathcal{F} \neq \emptyset$ throughout.

**The preference functional.** The donor cares about the entire trajectory, not just the terminal equilibrium. 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 $\alpha$ on $\mathcal{D}$ and any aid schedule $\gamma$,

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

We take $\alpha$ 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 $\alpha$ 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. Since $\mathcal{D} \cong \mathbb{R}^{IK}$ is simply connected, its first de Rham cohomology vanishes: $H^1(\mathcal{D}) = 0$. Every closed 1-form on $\mathcal{D}$ is therefore exact. If $\alpha = dV$ for some smooth $V : \mathcal{D} \to \mathbb{R}$, then

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

and the donor's evaluation depends only on the endpoints—the path is irrelevant. Preferences are genuinely path-dependent precisely when $\alpha$ is not closed.

**Grants and loans.** Two special cases deserve mention. An aid schedule $\gamma$ 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 $\Omega$, but the corresponding aid schedule $\gamma$ 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.

---

**4.3 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, $\pi$ 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 $s, t \in [0,1]$,

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

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, $\mathcal{F}$ is pre-compact in $C^0([0,1], \mathcal{D})$. To conclude compactness, we show $\mathcal{F}$ is closed. Conditions (ii) and (iii) are

[1] This phenomenon is governed by the monodromy of $\pi$: 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.

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 $\phi$. Hence $\mathcal{F}$ is closed, and compactness follows. ∎

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

*Aid as navigation.* The equilibrium manifold framework recasts aid design as a navigation problem in a precisely defined space. The donor does not simply move resources from $C$ to $A$ and $B$; it steers the economy through a sequence of general equilibria, with prices adjusting at every step to clear markets under the prevailing conflict equilibrium. The same resource transfer can move the economy to very different equilibria depending on which path is taken through $\mathcal{D}$—which commodities are transferred when, to whom, in what sequence—and the welfare consequences depend on both path and destination. This is the sense in which the approach here differs from treating aid as a wealth transfer: the equilibrium mechanism amplifies, shapes, and sometimes reverses the direct effect of endowment movements. It is also the sense in which the central result of the paper—that optimal peace is achievable at strictly lower cost than brute-force approaches require—has a precise foundation.

*Process and outcome.* Whether the trajectory through $\mathcal{D}$ matters, or only the terminal equilibrium, depends on the donor's preference form $\alpha$. When $\alpha = dV$—exact preferences—the donor is a pure outcome-optimizer: any two schedules that arrive at the same point on $\partial\mathcal{W}$ are equally good regardless of the intervening path. This captures a donor focused solely on whether peace is achieved, and on which peace, but indifferent to the trajectory of conflict, trade, and consumption in the interim. When $\alpha$ is not closed, intermediate equilibria carry welfare weight: the same terminal peace is more or less valuable depending on the sequence of prices, consumption levels, and conflict intensities along the way. Donors that care about civilian welfare during the transition, about which industries benefit from realigned trade, or about the credibility and legitimacy of the peace process are modeled by non-exact $\alpha$. The Pontryagin analysis of section 6 shows that the two cases have qualitatively different optimal schedules: straight-line geodesics in the exact case, and continuously adjusting trajectories in the non-exact case. The distinction has normative content: it formalizes whether *how* peace is made matters, or only that peace is made.

*Finance and topology.* The grants-versus-loans distinction turns out to be topological rather than merely financial. A loan restores the terminal endowment to $\omega^0$; a grant does not. But the equilibrium the economy occupies after a loan's repayment depends on which element of the fiber $\pi^{-1}(\omega^0)$ the equilibrium path terminates at—and this depends on the entire path traced through $\mathcal{D}$, not just the terminal endowment. If the transfer schedule, viewed as a loop in endowment space, acts nontrivially under the monodromy of $\pi$, the economy arrives at a different equilibrium at the same endowment $\omega^0$ after repayment. Peace achieved by a loan may therefore not survive repayment, not because of any financial incentive to rearm, but because the economy has been guided to a new equilibrium that happens to be peaceful while the loan is active and reverts to conflict once it ends. Whether loan-based peacemaking is durable is thus a topological question about the conflict environment—one that financial analysis alone cannot answer (see section D).

# 5 Existence of an Optimal Aid Schedule

The donor's problem is to choose a feasible aid schedule $\gamma \in \mathcal{F}$ that optimizes the preference functional $\mathcal{V}(\gamma) = \int_\gamma \alpha$. This is an optimization over a set of paths in an infinite-dimensional function space, and existence of a solution is not automatic. Standard finite-dimensional arguments do not apply. The challenge is to find a topology on $\mathcal{F}$ strong enough that the objective is continuous, yet weak enough that the strategy space is compact—a balance that requires exploiting both the general equilibrium structure of $\mathcal{D}$ and the regularity imposed by the speed constraint. The equilibrium manifold machinery of sections 3 and 4 makes this tractable: properness of $\pi$ converts the budget constraint into compactness of the path values, the speed constraint gives equicontinuity, and the Euclidean identification $\mathcal{D} \cong \mathbb{R}^{IK}$ allows a product topology under which both compactness and continuity hold. The extreme value theorem then delivers a solution.

**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.*

*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 $\tau$ of *uniform path convergence and weak-$*$ velocity convergence*: a net $\gamma^\lambda \to \gamma$ in $\tau$ if $\gamma^\lambda \to \gamma$ uniformly in $C^0([0,1],\mathcal{D})$ and $(\gamma^\lambda)' \rightharpoonup^* \gamma'$ in $L^\infty([0,1];T\mathcal{D})$. Since $L^1([0,1];T\mathcal{D})$ is separable, the weak-$*$ topology on the $M$-ball of $L^\infty$ is metrizable; hence $\tau$ is metrizable on $\mathcal{F}$ and sequential compactness is equivalent to compactness.

**Compactness.** Let $\{\gamma^n\}$ be a sequence in $\mathcal{F}$. By Lemma 4.3, $\mathcal{F}$ is compact in $C^0$, so a subsequence $\gamma^{n_k}$ converges uniformly to some $\gamma \in \mathcal{F}$. Because every $\gamma^n$ 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, a further subsequence has $(\gamma^{n_{k_j}})' \rightharpoonup^* \gamma'$ in $L^\infty$. Hence $\mathcal{F}$ is sequentially compact, hence compact, in $\tau$.

**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 $\alpha$ in the coordinates of $\mathcal{D}$, and $\langle \cdot,\cdot\rangle$ is the Euclidean inner product. Suppose $\gamma^n \to \gamma$ in $\tau$. Decompose:

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

The first integral tends to zero by weak-$*$ convergence of $(\gamma^n)'$ to $\gamma'$, 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 $\mathcal{V}(\gamma^n) \to \mathcal{V}(\gamma)$, and $\mathcal{V}$ is continuous in $\tau$.

**Conclusion.** $\mathcal{F}$ is compact and $\mathcal{V}$ is continuous in $\tau$, so the extreme value theorem delivers $\gamma^*$ and $\gamma_*$. ∎

**5.2 Remark (Exact preferences)**

*When $\alpha = dV$ for some smooth $V : \mathcal{D} \to \mathbb{R}$—the case of path-independent preferences (section 4)—the argument simplifies: $\mathcal{V}(\gamma) = V(\gamma(1)) - V(e_0)$ 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: there is always a best aid schedule, and the donor need not settle for an arbitrarily good approximation. Optimality here is global—the theorem delivers a supremum attainer, not merely a local optimum or stationary point—and it holds without convexity of preferences or uniqueness of equilibria along the path.

The theorem is deliberately silent on two issues that the subsequent analysis addresses. First, *uniqueness*: multiple optimal schedules can coexist whenever the peace boundary $\partial\mathcal{W}$ contains a region that $\alpha$ evaluates uniformly, or when the equilibrium manifold has multiple geodesics of equal preference-integral value connecting $e_0$ to $\partial\mathcal{W}$. In general equilibrium models with multiple equilibria, this is the rule rather than the exception, and the results below characterize what all optima share rather than singling one out. Second, *attainability*: the theorem assumes $\mathcal{F} \neq \emptyset$—that some feasible peace-achieving schedule exists—without characterizing when this is satisfied. The condition amounts to asking whether the budget and speed constraints are sufficient to steer the economy from $e_0$ to $\partial\mathcal{W}$ at all, a question whose answer depends on the geometry of $\mathcal{W}$ relative to the donor's resources and on whether $e_0$ and $\partial\mathcal{W}$ lie in the same connected component of $\mathcal{D}$.

The proof's product topology $\tau$ is also worth interpreting. Uniform convergence of paths ensures that the terminal peace condition is preserved under limits: a sequence of peace-achieving schedules converges to a peace-achieving schedule. Weak-$*$ convergence of velocities ensures that the preference integral $\int_\gamma \alpha$ is continuous: a key property that fails under the strong $L^\infty$ norm (too demanding for bounded-speed paths) and also fails under uniform path convergence alone (which does not control velocity integrals). The topology $\tau$ is thus not a technical convenience but the minimal structure that accommodates both the terminal constraint and the path-integral objective simultaneously.

# 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.

**The optimal control problem.** Under the identification $\mathcal{D} \cong \mathbb{R}^{IK}$ of Proposition 3.2, write the donor's problem as follows. 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

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

where $a\colon \mathbb{R}^{IK} \to \mathbb{R}^{IK}$ collects the component functions of $\alpha$ in the $\Phi$-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 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*

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

When $S(t) \neq 0$, the unique maximizer of $H$ is $u^*(t) = MS(t)/\|S(t)\|$, and the maximum speed is always attained.

---

### 6.1 Theorem (Pontryagin Characterization)

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

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

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

*Moreover, the switching function satisfies*

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

*where $\iota_v\, d\alpha$ denotes the interior product of the 2-form $d\alpha$ with $v$.*

---

*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.*

---

*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 $\alpha$. 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 $\eta$, 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 $\|S(t)\| \leq \eta_0$: no aid is delivered. The preference gain is insufficient to cover the fixed efficiency cost $\eta_0$.

(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 $\eta_0$ or $\kappa$ implies slower optimal pace; stronger preference intensity $\|S\|$ implies faster pace.

(iii) (Rapid) if $\|S(t)\| \geq \eta_0 + 2\kappa M$: 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 $(\eta_0, \kappa)$ (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.

---

**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 $\alpha$ 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*

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

*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 $Da = D^2V$ (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 connects to the discussion of path-independence in section 4: when $\alpha = dV$, the donor evaluates only the terminal equilibrium, and the optimal strategy is to reach the best point on $\partial\mathcal{W}$ as fast as possible along a straight line. Non-exact preferences break this constancy—the switching function rotates as the economy moves through $\mathcal{D}$, and the optimal direction of aid adjusts continuously according to (2).

---

**6.6 Remark (Singular arcs)**

*Singular arcs in the bang-bang problem (Corollary 6.2) represent indifference in the speed dimension and are non-generic. The efficiency-cost formulation of Proposition 6.3 resolves this: with $C$ strictly convex, the Hamiltonian $H_C$ is strictly concave in $u$ and the optimal speed is always uniquely determined, eliminating singular arcs entirely.*

---

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*

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

---

**6.7 Proposition (Making Peace on the Cheap)**

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

---

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

For strictness: if $\pi^{-1}(\omega^*)$ contains both war and peace equilibria, the optimal path steers the economy to the peaceful equilibrium by guiding prices along $\mathcal{D}$, while any robust peace transfer must eliminate the war equilibrium entirely, requiring a strictly larger transfer. To see that this is the generic case, note that properness of $\pi$ (Proposition 3.3) implies fibers are finite, but finiteness alone does not entail multiplicity. Multiplicity follows from a separate argument: since $\mathcal{D}$ is connected 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 non-empty in $\Omega$, and at any regular value in their intersection the fiber contains equilibria from both regions. The intersection is non-empty and open whenever $\pi$ fails to be injective near $\partial\mathcal{W}$—which is the generic situation, since global injectivity of $\pi$ would require a globally unique equilibrium at every endowment profile, a condition that is non-generic in the sense of Balasko (2011, Ch. 5). The price mechanism supplies the savings: as the path moves from $e_0$ to $e^*$, equilibrium prices shift, altering conflict payoffs $g_i$ and resource drain $r$, so that peace is achieved at strictly lower endowment cost than $R^*$ demands.■

**The theory of aid timing in summary.** The results of this section jointly constitute a theory of aid timing—when to deliver how much, in which direction, and why—that is grounded in equilibrium rather than convention. The central findings can be stated compactly. Optimal aid is *directional*: the switching function $S(t)$ continuously tracks the marginal value of movement in $\mathcal{D}$, rotating as the economy evolves, and the optimal schedule always moves in the direction of $S$. Optimal aid is *urgent when urgency is warranted*: absent efficiency costs, the donor always moves at maximum speed—gradual aid is generically suboptimal, not a policy choice—and gradualism enters only when the efficiency cost of rapid delivery is explicitly priced in the objective. When it is, the degree of graduation is pinned down by a first-order condition ($C'(\|u^*\|) = \|S\|$), not by convention or institutional capacity alone. And optimal aid is *cheap*: the equilibrium path exploits price adjustment to achieve peace at strictly lower cost than an approach that ignores the general equilibrium mechanism.

These results have direct implications for how donors should think about aid design. The ongoing debate between front-loading (deliver rapidly and early) and graduation (phase aid to match institutional capacity) dissolves into a question about primitives: front-loading is optimal when $\|S(t)\| \geq C'(M)$, graduation is optimal when $C'(0) < \|S(t)\| < C'(M)$, and suspension is optimal when $\|S(t)\| \leq C'(0)$. These thresholds depend jointly on the donor's preferences over equilibria (encoded in $\alpha$, which drives $S$) and on the delivery technology (encoded in $C$). No a priori presumption for or against speed is correct across all contexts; the model instead identifies the conditions under which each regime is optimal. Cross-country variation in observed aid timelines is predicted to reflect variation in urgency $\|S\|$ and delivery costs $(C', M)$—not bureaucratic inertia, donor caution, or arbitrary conditionality cycles.

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

**Setup.** Each donor $j \in \{1, \ldots, n\}$ independently chooses a transfer schedule $\tau_j \colon [0,1] \to \mathbb{R}^K_{\geq 0} \times \mathbb{R}^K_{\geq 0}$, with $\tau_j(0) = 0$, budget ceiling $\bar{\tau}_j$, efficiency parameter $\eta_j \geq 1$, 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 $\gamma \colon [0,1] \to \mathcal{D}$ 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* $(\tau_1^*, \ldots, \tau_n^*) \in \mathcal{T}$ *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 $\tau_j$.*

---

Assumption 7.2 holds when the map $\tau_j \mapsto \gamma(\tau_j + \sum_{k \neq j} \tau_k^*)$ is affine, which is the case locally near any regular point of the equilibrium manifold by the inverse function theorem. We impose it as a maintained assumption rather than deriving it from primitives; the alternative—working with mixed strategies on $\mathcal{T}$—would substantially complicate the framework. Whether Assumption 7.2 holds globally, and what relaxing it implies for the coordination failure results, are left for future work.

---

**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 4.3 applies componentwise to each donor's strategy space). The payoff $V_j$ is continuous in $(\tau_j, \tau_{-j})$ in the product topology $\tau$ 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 $\tau_j$ for fixed $\tau_{-j}$. 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 theorem, and a fixed point is a Nash equilibrium. ∎

**Pontryagin characterization of Nash equilibria.** Each donor's problem, taking $\tau_{-j}^*$ 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 $\psi_j \colon [0,1] \to \mathbb{R}^{IK}$ 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* $\|u_j\| \le M_j$*, 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 $T_{\mathrm{Nash}} \leq 1$, cooperation reaches peace by $T_{\mathrm{Nash}}/n$; 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 and the case for coordination.** The results of this section carry implications that cut across the normative, institutional, and empirical dimensions of foreign aid.

*The free-rider discount on peace.* The Nash–Pontryagin conditions make the mechanism precise. Each donor $j$'s costate $\psi_j$ is driven by the *aggregate* velocity $u^* = \sum_k (\tau_k^*)'$ (Proposition 7.4(i))—not by $j$'s own contribution alone. Every donor observes the peace process advancing, adjusts its urgency measure $S_j$ accordingly, and then optimizes as if it were alone at the margin. Each donor correctly perceives how the aggregate equilibrium path evolves but prices only its own fraction of the cost, internalizing only $1/n$ of the social benefit of acceleration. The result is a collectively rational but individually suboptimal equilibrium: the aggregate speed is $1/n$ of the cooperative optimum, not because donors are irrational or uninformed, but because peace is a non-excludable, non-rival good that markets cannot price.

*Coordination without additional resources.* The time-to-peace result of Corollary 7.6 is stark in a specific sense: it implies that the binding constraint on peacemaking in fragmented donor environments is not budget adequacy but coordination failure. Two donors with combined budget $B$ cooperating achieve peace in $1/2$ the time of the same two donors acting non-cooperatively, and in $1/n$ the time of $n$ non-cooperating donors with the same total $B$. Adding more money to a fragmented system amplifies this inefficiency: more donors, each free-riding on the others, converge on a Nash equilibrium that is slower as $n$ grows. Coordination mechanisms—joint conditionality, burden-sharing agreements, lead-donor frameworks with binding commitments, or multilateral institutions with unified disbursement authority— restore the cooperative speed *without additional resources.* The institutional design question is therefore not primarily about mobilizing more aid but about restructuring who delivers it and under what decision rule.

*The standard case for multilateralism, and a new one.* The conventional argument for multilateral aid institutions—the World Bank, the UN Peacebuilding Commission, regional development banks—appeals to economies of scale, informational pooling, and the credibility that size and permanence confer. Corollary 7.6 adds a distinct and quantitative argument: multilateral institutions reduce the free-rider discount on the peacemaking public good. A unified multilateral donor internalizes the full social benefit of advancing the equilibrium path toward $\partial\mathcal{W}$; a fragmented set of $n$ bilateral donors internalizes only $1/n$ of that benefit each. The welfare gain from multilateralization is therefore not measured in dollars—total resources are unchanged—but in time: the duration of conflict avoided, and all the economic, human, and institutional costs that accumulate during that period.

*Empirical implications.* The multi-donor model generates a set of predictions that distinguish it from single-donor analyses and that are, in principle, testable against existing cross-national data on conflict duration and aid composition. First, conditional on total aid volume, donor count should be a significant positive predictor of conflict duration: more donors means a slower aggregate peace trajectory, holding resources fixed. Second, the free-rider distortion is sharpest under common preferences—when all donors value the same equilibrium path—so the duration effect of donor count should be larger for conflicts with high donor salience (salient to many donors simultaneously) than for idiosyncratic or low-visibility conflicts. Third, heterogeneous preferences among donors introduce a partially offsetting effect: donors targeting different parts of $\partial\mathcal{W}$ do not compete for the same marginal contribution, partially internalizing their complementarity. The net prediction is that preference alignment among donors amplifies the free-rider problem, so coordination benefits are largest precisely when donors agree on what peace should look like—a result that inverts the naive intuition that agreement makes coordination easier to achieve without requiring it to be formalized.

## 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. The UCDP conflict dataset, AidData, and the OECD Creditor Reporting System jointly provide the right variables—conflict duration, disaggregated bilateral flows by donor, and total ODA by episode—to estimate this relationship, with conflict fixed effects absorbing the obvious confounders.

**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. Observable proxies for urgency include conflict intensity (casualty rates, displacement), proximity to a negotiated settlement, and the trajectory of recent fighting. 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 $r$. Whether this amplification is empirically recoverable depends on disaggregation of aid by commodity and timing within episodes, which is increasingly available in project-level datasets.

**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 $\pi$ (section D). When the conflict manifold has nontrivial monodromy—when the topology of the equilibrium structure makes certain paths irreversible—peace achieved by conditional transfers can be durable even after repayment. 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. Production, incomplete information about the conflict primitive, and stochastic equilibrium paths are the most immediate. A fully dynamic bargaining model that endogenizes the war region $\mathcal{W}$ 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 $\Gamma_{p,\omega_0}$ has a unique Nash equilibrium $m^*(p,\omega_0)$ that is smooth in $(p,\omega_0)$, 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 $G_i(m_i, m_{-i}) := p \cdot [s_i(m,p,\omega_0) - c_i(m_i,p,\omega_0)]$. 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 $\xi, \xi'$, 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

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 $\omega_i$, 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,

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

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.** 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 $\Phi$ 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 $\Phi$.* Define

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

That is, $\Phi$ 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 $\theta$.* 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 $\omega_i$ 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 $\omega_W$ 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 $\theta$ is smooth.

$\Phi \circ \theta = \mathrm{id}$. By construction, $\theta(p, w, \widehat{\omega}_0)$ produces a pair $(p, \omega)$ 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 $\Phi$ 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 $\theta$ 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. ✓

$\mathcal{D} \subseteq \text{im}(\theta)$. 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 $\Phi \circ \theta = \text{id}$ and $\text{im}(\theta) = \mathcal{D}$, 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 $\theta$, hence a diffeomorphism. This completes the proof. ■

## C Proof of Proposition 3.3

Smoothness of $\pi$ 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

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

**The discriminant locus.** Let $\Sigma \subset \Omega$ denote the set of critical values of $\pi$—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 $\Sigma$ has measure zero in $\Omega$.

**Covering space structure.** 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 $\pi$ is injective, giving $U_e \cap \pi^{-1}(\omega) = \{e\}$. The fiber is also compact: by Proposition 3.3, $\pi$ 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 $\{e_1, \dots, e_n\}$. The sets $U_{e_1}, \dots, U_{e_n}$ 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 $\omega$.

**Monodromy action.** 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 $\ell$ 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 $\tau$ with $\tau(1) = 0$, so its endowment path $\omega(\cdot)$ is a loop in $\Omega$ based at $\omega^0$. If $\tau$ avoids $\Sigma$—which holds generically by Sard—the induced aid schedule $\gamma$ 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 $\Sigma$ 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**, *The Equilibrium Manifold: Postmodern Developments in the Theory of General Economic Equilibrium*, Cambridge, MA: MIT Press, 2011.

**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.

**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.

**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.

**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.

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

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

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

**Mary, Sébastien**, "Revisiting US Food Aid and Civil Conflict," *American Economic Review*, 2026. forthcoming.

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

**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.

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