War and Peace in the Marketplace

Robert J. Carroll

University of Illinois at Urbana-Champaign

Draft — March 26, 2026

Contents

Chapter 1

Introduction

States acquire resources in two ways: through exchange and through force. The first creates markets; the second creates armies. When both are possible simultaneously—as they routinely are in international relations—prices and military deployments are jointly determined. A good's market price already reflects the ease with which it can be seized; a state's military posture already responds to the profitability of what lies across the border. Neither can be analyzed in isolation without losing the feedback that runs through the other.

International relations theory has largely studied these two phenomena with separate tools. Trade and economic interdependence are the province of international political economy (e.g., Keohane, 1984); war and the threat of force belong to international security (e.g., Waltz, 1979). The equilibrium concepts most common in each tradition reflect the split: Walrasian equilibrium for markets, Nash equilibrium for strategic military interaction. Each captures something real, but each misses what is determined by the other. A Walrasian analysis of trade that takes military posture as exogenous cannot explain how the threat of force affects market prices; a game-theoretic analysis of conflict that takes economic conditions as given cannot explain how trade opportunities shape the incentive to fight.

The cost of this separation is not merely aesthetic. Some of the most consequential questions in international political economy depend on exactly the interaction that the separated analyses suppress. Why does resource wealth so often generate conflict rather than prosperity (Collier and Hoeffler, 2004; Le Billon, 2013)? How does economic interdependence affect the probability of war (Polachek, 1980; Copeland, 2015)? When does the arrival of a new great power destabilize the international order (Gilpin, 1981)? Each question turns on how market prices respond to changes in military deployments and how deployment incentives respond to changes in market prices. That response is invisible when the two are analyzed separately.

This book proposes a framework that takes the joint determination seriously. Exchange and conflict are both competitive activities: self-interest operating through different social conventions—market prices in the first case, Nash best response in the second (Hirshleifer, 2001). The claim is not that one is prior to the other but that a satisfying theory must model them simultaneously. The appropriate equilibrium concept is a *dual-competitive equilibrium* (DCE): a tuple of commodity prices and military deployments that simultaneously satisfies Walrasian market clearing and Nash optimality in deployments. Prices are determined by supply and demand given military choices; military choices are determined by Nash incentives given prices. The two conditions hold together, or neither holds.

The central analytical result (Chapter 3) is that the set of all dual-competitive equilibria forms a smooth manifold E that is diffeomorphic to Euclidean space: $E \cong \mathbb{R}^{IL}$, where I is the number of states and L the number of commodities. This is a strong structure theorem. Its implications run through the remainder of the manuscript.

First, $E \cong \mathbb{R}^{IL}$ means that equilibria respond smoothly and predictably to changes in underlying parameters. Chapter 4 develops the comparative statics: a marginal change in a state's endowment produces a well-defined, computable change in all equilibrium prices and deployments. The signs and magnitudes of these responses are analytically determinate because the equilibrium mapping is smooth and its derivative is meaningful. The international system is, in a precise sense, continuous: small changes in the distribution of resources lead to small and predictable changes in the equilibrium configuration.

Second, $E \cong \mathbb{R}^{IL}$ enables a welfare analysis with precise objects. Chapter 5 establishes that any equilibrium with positive military deployment is generically Pareto inefficient and introduces an *efficiency gap* Δ^* that measures the welfare cost of conflict in terms of forgone market value. The Second Welfare Theorem fails: the market mechanism alone cannot decentralize the efficient allocation when states find exploitation profitable. The smooth structure of E is what makes Δ^* a well-defined, computable object rather than an ordinal claim about rankings.

Third, $E \cong \mathbb{R}^{IL}$ is the topological condition that makes consistent institutional design possible. The space of *institutional paths*—smooth trajectories through E connecting the current equilibrium to a desired target—is contractible. Chichilnisky's (1982) theorem then guarantees the existence of a continuous, anonymous social choice function over institutional alternatives, escaping the impossibility that Arrow (1963) establishes for ordinal aggregation. Chapter 8 develops this result, constructs gradient fields on E as policy instruments, and connects the framework to the empirical rational-design literature (Koremenos, Lipson and Snidal, 2001).

Chapter 6 develops three specific applications. The first concerns *systemic conflict*: the arrival of a new state with a large non-lootable endowment drives up the prices of contested goods, altering the deployment incentives of all incumbents and potentially destabilizing a previously peaceful system. The second identifies a *pathology in a standard measure of militarization*: the GDP share of military spending is neither monotone nor single-valued as a function of endowments in the DCE, which calls into question empirical strategies that rely on it as a proxy for military effort. The third formalizes the trade expectations mechanism of Copeland (2015): a state's expectation of future trade opportunities directly affects its current willingness to fight, and the DCE framework gives this intuition exact content as a comparative statics result on the peace condition.

Chapter 7 extends the framework in two directions. A numerical section shows how DCEs can be computed and illustrates the comparative statics graphically. A dynamic stochastic general equilibrium extension embeds the static model in an infinite-horizon environment with stochastic endowment shocks, establishes existence of a recursive DCE, and shows that dynamic trade expectations are amplified by a factor that depends on the persistence of the endowment process and the states' discount factors.

The manuscript is organized as follows. Chapter 2 sets up the individual state's optimization problem: the choice of military deployments and consumption given prices and the peace condition. Chapter 3 defines the DCE, proves existence, and establishes $E \cong \mathbb{R}^{IL}$. Chapters 4 and 5 develop the positive and normative theories. Chapter 6 presents three applications. Chapter 7 develops numerical methods and the DSGE extension. Chapter 8 analyzes social choice and institutional design. Chapter 9 concludes.

Chapter 2

The State's Problem

This chapter develops the individual state's decision problem—the microeconomic foundation on which the entire framework rests. We lay out the primitives (§2.1), state the optimization problem (§2.2), derive the first-order conditions and establish that the problem has a unique, smooth solution (§2.3), and collect the properties of the resulting demand function (§2.4). We close by studying the topology of the space of all possible behaviors under threat (§2.5).

2.1 Primitives

2.1.1 States, Goods, and Prices

We take as primitive a list of *states*, indexed $i = 1, \ldots, I$, where $I \geqslant 2$, and a list of *goods*, indexed $\ell = 1, \ldots, L$, where $L \geqslant 2$. States are unitary actors. Goods may be differentiated by physical characteristics, time, place of delivery, and the state of the world.

Each good ℓ is assigned a price $p^\ell > 0$. The vector of prices is $p = (p^1, \ldots, p^L) \in S := \mathbb{R}^L_{++}$. Since only relative prices matter, we normalize by setting good L as numéraire:

$$p \in S_N := \left\{ s \in S \;\middle|\; s^L = 1 \right\}.$$

2.1.2 Bundles, Endowments, and the Consumption Space

State i's consumption *bundle* is a vector $x_i = (x_i^1, \dots, x_i^L)$. The set of all possible consumption bundles is

$$X_i := \prod_{\ell=1}^{L} (b_i^\ell, \infty),$$

where $b_i^\ell < 0$ for each i and ℓ. This is an open, convex subset of $\mathbb{R}^L$, bounded below but not above, containing the origin and the entire strictly positive orthant.

The lower bounds b_i^ℓ are state-and-good-specific. A state rich in oil can sustain larger oil losses than a landlocked state; a state with deep capital markets can absorb larger financial shocks. The bounds encode the structural limits of each state in each commodity: the floor below which losses become catastrophic and the state, as a functioning entity, ceases to operate. We collect these bounds into a vector $b_i = (b_i^1, \dots, b_i^L) \in \mathbb{R}^L_{--}$.

States enter the world with an initial *endowment* $\omega_i = (\omega_i^1, \dots, \omega_i^L) \in \mathbb{R}^L_{++}$; endowments are strictly positive and lie well inside X_i. An *allocation* is a list $x = (x_1, \dots, x_I)$; the initial allocation is $\omega = (\omega_1, \dots, \omega_I)$. The set of all possible allocations is $X := \prod_i X_i$, and we write $\Omega_0 := \mathbb{R}^{IL}_{++}$ for the set of strictly positive endowment profiles. The initial allocation $\omega \in \Omega_0$ is the model's main parameter: the distribution of resources across the globe.

2.1.3 Preferences

State i's preferences over bundles are represented by a utility function $u_i : X_i \to \mathbb{R}$.

2.1 Assumption (Preferences)
For each i, the utility function $u_i : X_i \to \mathbb{R}$ is:

1. *smooth (C^∞) on X_i;*
2. *smoothly strictly quasiconcave (the bordered Hessian of u_i is nonsingular with sign $(-1)^L$ for all $x_i \in X_i$);*
3. *smoothly monotone ($Du_i(x_i) \gg 0$ for all $x_i \in X_i$);*
4. *boundary blow-up: $u_i(x_i) \to -\infty$ as $x_i^\ell \to (b_i^\ell)^+$ for any ℓ.*

These are the standard assumptions of smooth general equilibrium theory (Debreu, 1972; Mas-Colell, 1985; Balasko, 2009, 2011), adapted to the bounded-below domain X_i. Conditions (1)–(3) are identical to the classical setting on the strictly positive orthant. Condition (4) replaces the role played by the boundary of $\mathbb{R}^L_{++}$ in the standard theory: as any component of consumption approaches its lower bound, utility diverges to $-\infty$, ensuring that optimal bundles remain in the interior. We denote the class of utility functions satisfying Assumption 2.1 by $\mathcal{U}_i$.

2.1.4 Where Utility Comes From

States are not individuals; the meaning of u_i requires comment. We take the unitary actor assumption seriously: each state acts *as if* it maximizes a single objective over bundles. This is a modeling commitment, not a claim about the internal politics of states. Three readings are available, in increasing order of ambition.

The most conservative reading treats u_i as the objective of a representative consumer. If the state's population has homogeneous and homothetic preferences, individual demands aggregate into a single demand function that can be rationalized by a representative utility (Mas-Colell, Whinston and Green, 1995, Chapter 4). Under this reading, u_i inherits its properties from the common individual preferences, and the consumption bundle x_i is the aggregate consumption of the state's population.

A less restrictive reading allows heterogeneous preferences but imposes that the state redistributes internally so as to maximize a social welfare function. If this function is smooth and strictly quasiconcave in aggregate consumption—as it will be under standard conditions on individual utilities and the redistribution mechanism—then it serves as u_i. The state's internal politics determine the *shape* of u_i (whose welfare is weighted how much) but not its existence.

The most permissive reading is purely behavioral: u_i is whatever smooth, strictly quasiconcave, monotone function rationalizes the state's observed demand behavior. Under this reading, the content of u_i is entirely in its implications for f_i, and the question of what generates it—domestic politics, institutional design, the preferences of an autocrat—is deferred. This is the reading most consonant with the unitary actor assumption and with the general approach of absorbing the state's internal structure into the functional objects (τ_i, κ_i, and now u_i) rather than modeling it explicitly.

2.1.5 Walrasian and Marshallian Demand

Whichever reading one adopts, the demand function f_i is *Walrasian*: it takes all prices and the full endowment as arguments and returns the state's purchases across all commodities simultaneously. This is a methodological choice that deserves comment, because the dominant tradition in applied economics—and in most empirical work in international relations—is Marshallian.

Friedman (1949) offers the sharpest defense of the Marshallian approach. His reconstruction of Marshall's demand curve holds real income constant (via compensating price variations) and isolates the commodity in question from all but its closest substitutes. The justification is methodological: the analyst "will inevitably separate commodities that are closely related to the one immediately under study from commodities that are more distantly related," because "the forces affecting demand in any problem and the forces affecting supply will yield two lists that contain few items in common" (469). On this view, the economist's job is to identify a handful of related markets, hold everything else constant, and study the local mechanism in isolation.

This is sound advice for a world in which the interdependencies across distant markets are negligible—where the demand for tea has little to do with the price of steel. It is less sound in a world where states fight over trade flows. The battle function $s_i(m, x, \omega)$ creates precisely the kind of interdependence that the Marshallian approach assumes away. When conflict targets net exports, the demand for oil cannot be analyzed in isolation from the price of grain, because both enter the spoils function; the deployment that determines how much grain a state can seize also determines how much oil it diverts from the market. The "other things" that the Marshallian holds constant are exactly the things that move.

There is also a substantive reason, beyond the formal coupling, why the Walrasian framework is necessary here. General equilibrium highlights *resources*. In the study of intranational violence, resources play a decisive—often the decisive—role (Collier and Hoeffler, 2004; Humphreys, 2005). Some resources create conflict while others do not; the difference lies in their lootability, their position in the global price system, and their role in the opportunity costs of militarization (Le Billon, 2013). A general equilibrium approach forces the analyst to specify which commodities exist, how they substitute, and how conflict redirects them. The Marshallian approach, by isolating a single market, permits the analyst to set these questions aside. But they are the questions that matter: whether a state fights over oil depends on what oil is worth, which depends on the prices of everything oil substitutes for and competes with, which depend in turn on who is fighting over what. The resource problem of war is inherently general.

There is, finally, a deeper formal connection. Friedman's Marshallian demand curve, because it compensates for income effects, is itself a kind of price-dependent demand function: the income at each point on the curve is a function of the price, chosen to keep real income constant. Balasko (2009, Chapter 6) shows that the equilibrium theory extends cleanly to such functions. In our setting, the price-dependence of f_i arises not from income compensation but from conflict: the state's effective preferences over market purchases depend on prices because the spoils of conflict and the costs of militarization do. The mathematical structure is the same—a demand function whose arguments include the price vector in a non-standard way—but the economic content is different. Friedman's demand curve is price-dependent because the analyst chooses to hold real income constant; ours is price-dependent because the world is dangerous.

2.1.6 Militarization and Opportunity Costs

State i's level of militarization is $m_i \in M_i \subseteq \mathbb{R}$. The militarization profile is $m = (m_1, \ldots, m_I) \in M := \prod_i M_i$. Militarization imposes an opportunity cost on the endowment through the *opportunity cost function*

$$c_i : M_i \times S \times \mathbb{R}^L_{++} \to \mathbb{R}^L,$$

so that the resulting diminished endowment is $\omega_i - c_i(m_i, p, \omega_i)$.

The function c_i is a derived object. The state is a producer of force (we develop this at length in Carroll, 2026): it possesses a technology τ_i mapping resource inputs to military output and a cost function κ_i measuring the burden of extraction. The state's production problem is

$$\min_{e_i \in \mathbb{R}^L} \kappa_i(e_i) \quad \text{subject to} \quad \tau_i(e_i) = m_i,$$

and the solution defines a policy function $\pi_i : M_i \to \mathbb{R}^L$ mapping force targets to optimal mobilization plans. The opportunity cost vector is $c_i(m_i, p, \omega_i) = \pi_i(m_i; p, \omega_i)$. When κ_i depends on prices—as with the linear cost $\kappa_i(e) = p \cdot e$ or the price-augmented cost $\kappa_i(e) = p \cdot e + \gamma \|e\|^2$—the optimal plan changes with p, and c_i inherits this dependence.

The dependence on ω_i is equally natural. A state rich in oil but poor in labor faces a different mobilization calculus than a state with the reverse profile, even at the same prices and force target. What the state *has* shapes what it costs to extract: an agrarian state mobilizing labor pays in foregone harvests; an industrial state mobilizing capital pays in foregone output.

2.2 Assumption (Opportunity costs)

For each i, the opportunity cost function c_i is smooth in all arguments, satisfies $c_i(0, p, \omega_i) = 0$ (zero deployment costs nothing), and is coercive in m_i ($\|c_i(m_i, p, \omega_i)\| \to \infty$ as $|m_i| \to \infty$, uniformly in (p, ω_i) on compacts). The diminished endowment respects the consumption space:

$$\omega_i^\ell - c_i^\ell(m_i, p, \omega_i) > b_i^\ell \quad \textit{for all feasible } m_i,\ \ell.$$

2.1.7 Battle Functions

The *battle function* determines how conflict shifts resources among states.

2.3 Assumption (Conflict technology)

For each i, the battle function

$$s_i : M \times X \times \Omega_0 \to \mathbb{R}^L$$

is smooth. We write $s_i(m_i, m_{-i}, x, \omega)$, where x is the allocation and ω is the endowment profile.

Total consumption respects the consumption space:

$$x_i^\ell + s_i^\ell(m, x, \omega) > b_i^\ell \quad \textit{for all feasible } (m, x, \omega),\ \ell.$$

The arguments of s_i each carry substantive meaning.

- **Dependence on** m: how much force each state projects. This is the standard channel.

- **Dependence on** x: what has been produced and traded. If two states are fighting over oil, the oil available to seize depends on how much has been consumed, exported, or stockpiled—all equilibrium outcomes.

- **Dependence on ω**: what resources exist in the first place. You can only loot what is there. The joint dependence on x and ω allows the battle function to operate on *net flows*: if s_i depends on $(\omega_j - x_j)$ for other states j, then conflict targets what is being shipped rather than what is being kept.

2.4 Example (Lootable net exports)
Consider the two-state battle function

$$s_i^\ell(m_i, m_{-i}, x, \omega) = \frac{\lambda^\ell \cdot (\omega_{-i}^\ell - x_{-i}^\ell) \cdot (m_i - m_{-i})}{1 + m_i + m_{-i}},$$

where $\lambda^\ell \geqslant 0$ is a lootability parameter. The spoils are proportional to the other state's net exports $(\omega_{-i}^\ell - x_{-i}^\ell)$: you loot what the other side is shipping, not what it keeps at home. When the other state exports more of good ℓ, there is more to intercept; when it imports, there is nothing to seize along this channel.

2.2 The State's Problem under Threat

State i observes prices p and chooses $(x_i, m_i) \in X_i \times M_i$ to solve

$$\max_{\substack{x_i \in X_i,\, m_i \in M_i \\ p \cdot x_i \leqslant p \cdot (\omega_i - c_i(m_i, p, \omega_i))}} u_i\big(x_i + s_i(m_i, m_{-i}, x, \omega)\big). \tag{SPT$_i$}$$

The state maximizes the utility of *total consumption*—what she buys at market plus what she gains or loses through conflict—subject to the budget constraint that market purchases not exceed the diminished endowment. This is the *state's problem under threat*.

Two features distinguish (SPT$_i$) from a standard consumer problem. First, the budget set depends on the deployment m_i through the opportunity cost c_i. Second, the objective depends on m_i and on the full allocation x through the battle function s_i. The state's consumption and militarization decisions are therefore jointly determined: what you buy at market depends on what you expect to gain or lose in conflict, and what you gain or lose in conflict depends on what everyone has bought.

2.3 The Solution

2.3.1 First-Order Conditions

Let $x_i' := x_i + s_i(m_i, m_{-i}, x, \omega) \in X_i$ denote total consumption. The derivative of the objective with respect to x_i picks up the chain rule through s_i:

$$\frac{\partial}{\partial x_i} u_i(x_i') = Du_i(x_i') \cdot (I + D_{x_i} s_i),$$

where I is the $L \times L$ identity and $D_{x_i} s_i$ is the Jacobian of the battle function with respect to x_i. The matrix $(I + D_{x_i} s_i)$ is the *effective return* on a market purchase: buying one more unit of good ℓ yields not just one unit of consumption but also an induced change in spoils.

A necessary condition for (x_i^*, m_i^*) to solve (SPT$_i$) is that there exists a multiplier $\lambda^* > 0$ such that:

$$Du_i(x_i') \cdot (I + D_{x_i} s_i) = \lambda^* p, \tag{FOC 1}$$

$$Du_i(x_i') \cdot D_{m_i} s_i = \lambda^* p \cdot D_{m_i} c_i(m_i^*, p, \omega_i), \tag{FOC 2}$$

$$p \cdot \left[\omega_i - c_i(m_i^*, p, \omega_i) - x_i^*\right] = 0. \tag{FOC 3}$$

Equation (FOC 1) is the modified tangency condition: the marginal rate of substitution, adjusted by the effective return on market purchases, equals the price ratio. Equation (FOC 2) equates the marginal utility of force to its marginal budgetary cost. Equation (FOC 3) is the budget constraint, holding with equality by monotonicity.

2.3.2 The Coupling

From (FOC 1):

$$Du_i(x_i') = \lambda^* p \cdot (I + D_{x_i} s_i)^{-1}.$$

Substituting into (FOC 2):

$$p \cdot (I + D_{x_i} s_i)^{-1} \cdot D_{m_i} s_i = p \cdot D_{m_i} c_i(m_i^*, p, \omega_i). \tag{FOC 2*}$$

The left side depends on x_i through the matrix $(I + D_{x_i} s_i)^{-1}$. The militarization decision cannot be separated from the consumption decision: the optimal deployment depends on what the state is buying, and what the state buys depends on its deployment. This is the intended coupling.

2.5 Remark (When the coupling vanishes)

If $D_{x_i}s_i = O$—the battle function does not depend on the allocation—then $(I + D_{x_i}s_i)^{-1} = I$, and (FOC 2) reduces to $p \cdot D_{m_i}s_i = p \cdot D_{m_i}c_i$, which determines m_i^* independently of x_i^*. The problem decomposes: first choose the deployment, then solve a standard consumer problem. This is the separable special case.*

2.3.3 Regularity

The first-order system (FOC 1)–(FOC 3) comprises $L + 2$ equations in the $L + 2$ unknowns $(x_i^1, \ldots, x_i^L, \lambda, m_i)$. Write its Jacobian as

$$J = \begin{bmatrix} A & B \\ E & C \end{bmatrix},$$

where A is the $(L+1) \times (L+1)$ block from (FOC 1) and (FOC 3) differentiated with respect to (x_i, λ), C is the scalar from (FOC 2) differentiated with respect to m_i, and B and E are the cross blocks. We now develop conditions under which J is nonsingular.

2.6 Assumption (Conflict feedback)

For each i and all (m, x, ω) in the relevant domain, the matrix $R := I + D_{x_i}s_i(m, x, \omega)$ is invertible.

A sufficient condition is $\|D_{x_i}s_i\| < 1$ in the operator norm: the spoils respond to the allocation, but less than one-for-one. This rules out situations in which buying a unit of a good at market causes you to lose more than a unit through conflict—which would undermine the logic of exchange.

2.7 Assumption (Diminishing returns to militarization)

For each i, the net return to deployment

$$m_i \longmapsto p \cdot (s_i(m_i, m_{-i}, x, \omega) - c_i(m_i, p, \omega_i))$$

is strictly concave in m_i at any solution to (SPT$_i$).

This says that each additional unit of militarization yields diminishing marginal returns, taking into account both the spoils it generates and the resources it consumes. It is the direct analogue of the standard second-order condition for the militarization problem.

2.8 Lemma (Jacobian invertibility)

Under Assumptions 2.1, 2.6, and 2.7, the Jacobian J is nonsingular at any interior solution to (SPT$_i$).

Proof. Write $R = I + D_{x_i}s_i$, let $H = D^2u_i(x_i')$ denote the Hessian of u_i at total consumption, and let $d = D_{m_i}s_i$ denote the marginal spoils vector.

The block A contains the bordered Hessian of u_i evaluated at the *effective prices* $q = R^{-\top}p$. Since R is invertible (Assumption 2.6) and $p \gg 0$, the effective prices satisfy $q \gg 0$ whenever R is sufficiently close to I (guaranteed by $\|D_{x_i}s_i\| < 1$). Strict quasiconcavity of u_i (Assumption 2.1(2)) ensures that the bordered Hessian of u_i at any strictly positive prices is nonsingular with sign $(-1)^L$. A change of variables by the invertible matrix R preserves nonsingularity. Therefore $\det(A) \neq 0$.

Since A is invertible, the determinant of the full Jacobian factors via the Schur complement:

$$\det(J) = \det(A) \cdot \left(C - E\,A^{-1}\,B\right).$$

The Schur complement $\sigma := C - E\,A^{-1}B$ is a scalar. We show it is nonzero by decomposing it into a leading term and a perturbation.

In the separable case ($D_{x_i}s_i = O$), the cross blocks E and B encode only the direct coupling through D^2u_i between the consumption and militarization first-order conditions. The Schur complement reduces to

$$\sigma_0 = p \cdot \left(D^2_{m_im_i}s_i - D^2_{m_im_i}c_i\right) + d^\top H\,d - (d^\top H\,p)(p^\top H\,p)^{-1}(p^\top H\,d),$$

where $d = D_{m_i}s_i$ and the last two terms arise from the bordered-Hessian structure of A^{-1}. The first term is strictly negative by Assumption 2.7. The remaining terms constitute a negative-semidefinite quadratic form (the Schur complement of H in the bordered Hessian, projected onto d). Thus $\sigma_0 < 0$.

In the coupled case, write $\sigma = \sigma_0 + \eta$, where η collects all terms involving $D_{x_i}s_i$ and its derivatives. Since η is continuous in $D_{x_i}s_i$ and vanishes when $D_{x_i}s_i = O$, there exists a threshold $\varepsilon_0 > 0$ (depending on σ_0, the curvature of u_i, and the derivatives of s_i and c_i at the solution) such that $|\eta| < |\sigma_0|$ whenever $\|D_{x_i}s_i\| < \varepsilon_0$. Under the conflict feedback bound $\|D_{x_i}s_i\| < 1$, and provided $\varepsilon_0 \geq 1$ (which holds when the diminishing returns in Assumption 2.7 are strict relative to the conflict feedback), the Schur complement satisfies $\sigma < 0$. ∎

2.9 Remark

When s_i is affine in x_i (so that $D^2_{x_i x_i} s_i = 0$), as in Example 2.4, the perturbation η takes a closed form involving only $D_{x_i} s_i$, H, d, and the bordered-Hessian inverse, and the threshold ε_0 can be computed explicitly. The condition $\varepsilon_0 \geqslant 1$ then reduces to a concrete inequality relating the curvature of u_i, the diminishing returns rate, and the magnitude of the conflict feedback.

2.3.4 Uniqueness

Invertibility of the Jacobian gives smoothness but not uniqueness. For that we need one further condition.

2.10 Assumption (Joint quasiconcavity)

For each i, the composite objective $(x_i, m_i) \mapsto u_i(x_i + s_i(m_i, m_{-i}, x, \omega))$ is jointly strictly quasiconcave in (x_i, m_i) on the budget set.

2.11 Remark

When s_i is affine in (x_i, m_i), Assumption 2.10 follows from the strict quasiconcavity of u_i alone: quasiconcavity is preserved under composition with affine maps. For non-affine s_i, the assumption imposes that the nonlinearity of the battle function does not introduce non-convexities into the objective.

2.3.5 Main Results

With the regularity conditions in hand, we can state the two main results on the state's problem.

2.12 Proposition (Existence)

Under Assumptions 2.1, 2.2, 2.3, and 2.6, the state's problem under threat (SPT$_i$) has a solution for any $(p, m_{-i}, \omega) \in S_N \times M_{-i} \times \Omega_0$.

Proof. Fix (p, m_{-i}, ω). Coerciveness of c_i (Assumption 2.2) ensures that sufficiently large $|m_i|$ renders the budget infeasible; this bounds m_i. Given bounded m_i, the budget constraint bounds $p \cdot x_i$ from above. Boundary blow-up of u_i (Assumption 2.1(4)) ensures that any maximizing sequence stays away from ∂X_i: if some component $x_i^\ell \to (b_i^\ell)^+$, utility diverges to $-\infty$, which cannot be optimal. Together these bounds yield a compact effective feasible set in the interior of X_i. The objective is continuous, so a maximum exists by Weierstrass. ■

2.13 Proposition (Uniqueness and smoothness)

Under Assumptions 2.1, 2.2, 2.3, 2.6, 2.7, and 2.10, the solution (x_i^*, m_i^*) to (SPT$_i$) is unique and varies smoothly in (p, m_{-i}, ω).

Proof. Uniqueness follows from joint strict quasiconcavity (Assumption 2.10): if (x_i, m_i) and (x_i', m_i') are two distinct solutions, their midpoint is feasible and strictly preferred, contradicting optimality of either. Smoothness follows from the implicit function theorem applied to the first-order system (FOC 1)–(FOC 3): nonsingularity of J (Lemma 2.8) is exactly the required hypothesis. ■

2.4 The Demand Function under Threat

Propositions 2.12 and 2.13 yield a *demand function under threat*:

$$f_i(p, m_{-i}, \omega) := x_i^*(p, m_{-i}, \omega),$$

the market purchases of a state that has optimally chosen both its consumption and its militarization. This function depends on prices through three channels: the budget constraint, the opportunity cost $c_i(m_i, p, \omega_i)$, and the battle function $s_i(m, x, \omega)$. It is a price-dependent demand function in the sense of Balasko (2009, Chapter 6).

2.14 Lemma (Walras' law under threat)

The demand function under threat satisfies

$$p \cdot f_i(p, m_{-i}, \omega) = p \cdot (\omega_i - c_i(m_i^*, p, \omega_i)).$$

Proof. Monotonicity of u_i ensures the budget constraint binds at any solution. The binding constraint (FOC 3) is the stated identity. ∎

The demand function under threat is smooth (Proposition 2.13) and satisfies Walras' law (Lemma 2.14). It also satisfies a properness condition that prevents equilibria from escaping to infinity.

2.15 Lemma (Desirability)
Consider a sequence $(p^n, m^n_{-i}, \omega^n) \to (p^0, m^0_{-i}, \omega^0)$ where $(p^0, m^0_{-i}, \omega^0)$ lies on the boundary of the domain (some price $p^{0,\ell} = 0$, or some endowment component $\omega_i^{0,\ell} = b_i^\ell$). Then $\|f_i(p^n, m^n_{-i}, \omega^n)\| \to \infty$.

Proof. Suppose $p^{0,\ell} \to 0$ for some ℓ. The budget constraint $p \cdot x_i \leq p \cdot (\omega_i - c_i)$ permits the state to purchase up to $(p \cdot (\omega_i - c_i))/p^\ell$ units of good ℓ (holding all other purchases at zero). As $p^\ell \to 0$, this upper bound diverges. Monotonicity of u_i (Assumption 2.1(3)) ensures the state exploits this: any bounded demand in good ℓ could be improved by purchasing more, so $f_i^\ell(p^n, m^n_{-i}, \omega^n) \to \infty$.

Now suppose all prices remain bounded away from zero but $\omega_i^{0,\ell} \to (b_i^\ell)^+$ for some ℓ. At the optimum, total consumption $x_i' = x_i + s_i$ lies in the interior of X_i (the existence proof, Proposition 2.12, established this). Boundary blow-up of u_i (Assumption 2.1(4)) provides a quantitative lower bound: for any target utility level $\overline{u}$, there exists $\varepsilon > 0$ such that $x_i'^\ell > b_i^\ell + \varepsilon$ at any solution with $u_i(x_i') \geq \overline{u}$. Since the spoils s_i are bounded on compact subsets of the parameter space, the market purchase $x_i^\ell = x_i'^\ell - s_i^\ell$ is bounded below. But Walras' law (Lemma 2.14) requires $p \cdot x_i = p \cdot (\omega_i - c_i)$, and as $\omega_i^\ell \to (b_i^\ell)^+$ the right side collapses in good ℓ. To maintain both the lower bound on x_i^ℓ and the budget identity, the state must shift purchases toward good ℓ and away from others—but the budget is shrinking. Contradiction with boundedness of demand: at least one component of f_i must diverge. ∎

Smoothness, Walras' law, and desirability are the three properties required by the price-dependent equilibrium manifold theory of Balasko (2009, Chapter 6). That theory establishes that the equilibrium manifold of an economy with price-dependent demand retains all of its key structural properties—smooth submanifold, diffeomorphism with Euclidean space, smooth and proper natural projection—provided total resources are variable. Since the endowment ω is our main parameter, this condition is satisfied. We exploit this connection in Chapter 3.

2.16 Remark (What is not a diffeomorphism)
In the standard theory of smooth economies, the individual demand function $f_i : S \times \mathbb{R}_{++} \to X_i$ is itself a diffeomorphism: given a consumption bundle, one can recover the unique price-wealth pair that generated it (Balasko, 2011, Proposition 3.23). This property fails for price-dependent demand (Balasko, 2009, 131). In our setting, the demand function under threat depends on prices through both the budget constraint and the battle function, so the same bundle x_i can arise from different prices via different conflict-adjusted effective returns. The inverse map does not exist.

What is *a diffeomorphism is the* extended *demand function, which separates the two roles of the price system: the effective price $q = R^{-\top} p$ entering through the modified tangency condition (FOC 1), and the budget price p entering through the constraint (FOC 3). The extended demand, which keeps these arguments distinct, retains the diffeomorphism property (Balasko, 2009, Proposition 6.2.1). The equilibrium manifold is diffeomorphic to Euclidean space regardless: that result requires only smoothness, Walras' law, and desirability of the strict demand, not the diffeomorphism property.*

Alongside the demand function, the state's problem also determines an optimal deployment:

$$\mu_i(p, m_{-i}, \omega) := m_i^*(p, m_{-i}, \omega),$$

which is smooth by the same argument. The equilibrium concept in the next chapter will require that these deployment functions are mutually consistent.

2.5 The Space of Behaviors under Threat

The pair (f_i, μ_i) encodes everything about how State i responds to the competitive environment: given prices, the deployments of others, and the distribution of resources, it tells us what the state buys and how much force it projects. Different states produce different pairs, because they differ in preferences, costs, and conflict technologies. We now ask a topological question about the collection of all such pairs: what is the shape of the space of possible behaviors under threat?

2.5.1 The Parameter and Behavior Spaces

An admissible *state character* is a triple $(u_i, c_i, s_i) \in \mathcal{U}_i \times \mathcal{C}_i \times \mathcal{S}_i$, where $\mathcal{U}_i$ is the set of utility functions satisfying Assumption 2.1, $\mathcal{C}_i$ is the set of opportunity cost functions satisfying Assumption 2.2, and $\mathcal{S}_i$ is the set of battle functions satisfying Assumptions 2.3 and 2.6. We equip each factor with the smooth (compact-open C^∞) topology and the product $\mathcal{A}_i := \mathcal{U}_i \times \mathcal{C}_i \times \mathcal{S}_i$ with the product topology.

The *solution map* sends each state character to the corresponding demand-deployment pair:

$$\Phi_i : \mathcal{A}_i \to C^\infty(S_N \times M_{-i} \times \Omega_0,\ X_i \times M_i), \qquad \Phi_i(u_i, c_i, s_i) = (f_i, \mu_i).$$

The image $\mathcal{F}_i := \Phi_i(\mathcal{A}_i)$ is the *space of behaviors under threat*: the collection of all demand-deployment functions that can arise from some admissible state character.

2.5.2 Contractibility

The central topological result is the following.

2.17 Theorem

The space $\mathcal{F}_i$ is contractible.

The proof proceeds in two stages. We first retract $\mathcal{F}_i$ onto the subspace of peaceful behaviors, then retract that subspace to a point.

Stage 1: Dialing down conflict. For $t \in [0,1]$ and $(u_i, c_i, s_i) \in \mathcal{A}_i$, define the scaled character $(u_i, c_i, t \cdot s_i)$. This remains admissible: the scaled battle function $t \cdot s_i$ is smooth, the spoils conditions are preserved (scaling toward zero only shrinks the spoils), and the conflict feedback condition (Assumption 2.6) holds since

$$\|D_{x_i}(t \cdot s_i)\| = t \cdot \|D_{x_i} s_i\| \leqslant \|D_{x_i} s_i\|.$$

The domain constraint $x_i + t \cdot s_i(m, x, \omega) \in X_i$ is also preserved: if $x_i + s_i \in X_i$, then $x_i + t \cdot s_i$ lies between x_i (which is in X_i) and $x_i + s_i$ (which is in X_i), and X_i is convex.

By Proposition 2.13, the solution varies smoothly in the battle function, so the induced map

$$H_1 : \mathcal{A}_i \times [0,1] \to \mathcal{F}_i, \qquad H_1((u_i, c_i, s_i),\ t) = \Phi_i(u_i, c_i, t \cdot s_i)$$

is continuous. At $t = 1$, $H_1(\cdot, 1) = \Phi_i$. At $t = 0$, $H_1(\cdot, 0) = \Phi_i(u_i, c_i, 0)$: the behavior arising from zero conflict.

When $s_i = 0$, conflict confers no benefit and militarization is pure cost, so the optimal deployment is $m_i^* = 0$ (by centeredness of c_i: zero deployment costs nothing). The demand function reduces to standard demand at full wealth $p \cdot \omega_i$. The image at $t = 0$ is therefore the subspace

$$\mathcal{F}_i^0 := \{\Phi_i(u_i, c_i, 0) \mid (u_i, c_i) \in \mathcal{U}_i \times \mathcal{C}_i\},$$

which consists of pairs $(f_i^0, 0)$ where f_i^0 is the standard demand function associated with u_i and $\mu_i \equiv 0$. Since H_1 acts as the identity on $\mathcal{F}_i^0$ (scaling zero by t gives zero), this is a deformation retraction of $\mathcal{F}_i$ onto $\mathcal{F}_i^0$.

Stage 2: Contracting the peaceful subspace. Elements of $\mathcal{F}_i^0$ are indexed by $u_i \in \mathcal{U}_i$ alone (since $s_i = 0$ and $m_i^* = 0$, the cost function plays no role). The class $\mathcal{U}_i$ of smooth, strictly quasiconcave, monotone utilities with boundary blow-up on X_i is *not* convex: convex combinations of quasiconcave functions need not be quasiconcave.

We appeal to a classical result in smooth preference theory. Since X_i is open and convex, every $u_i \in \mathcal{U}_i$ admits a smooth, strictly concave representative $\tilde{u}_i : X_i \to \mathbb{R}$ that generates the same demand function: there exists a smooth, strictly increasing $\phi_i : \mathbb{R} \to \mathbb{R}$ such that $\tilde{u}_i = \phi_i \circ u_i$ is strictly concave on X_i (Mas-Colell, 1985, Proposition 2.6.4). The concave representative inherits monotonicity and boundary blow-up from u_i.

Denote by $\tilde{\mathcal{U}}_i$ the class of smooth, strictly concave, monotone functions on X_i with boundary blow-up. This class *is* convex: if $\tilde{u}$ and $\tilde{v}$ are both strictly concave, monotone, and blow up at ∂X_i, then so does $(1-t)\tilde{u} + t\tilde{v}$ for any $t \in [0, 1]$. Fix a distinguished $\tilde{u}_0 \in \tilde{\mathcal{U}}_i$. The homotopy

$$H_2 : \tilde{\mathcal{U}}_i \times [0,1] \to \tilde{\mathcal{U}}_i, \qquad H_2(\tilde{u}_i,\, t) = (1-t)\,\tilde{u}_i + t\,\tilde{u}_0$$

is continuous, with $H_2(\cdot, 0) = \mathrm{id}$ and $H_2(\cdot, 1) \equiv \tilde{u}_0$. Since the demand function depends only on ordinal preferences—and each $\tilde{u}_i$ generates the same demand as the original u_i—the induced map on $\mathcal{F}_i^0$ contracts it to the single demand function of $\tilde{u}_0$.

Conclusion. Composing the two deformation retractions, $\mathcal{F}_i$ deformation retracts onto $\mathcal{F}_i^0$, which in turn contracts to a point. Hence $\mathcal{F}_i$ is contractible.

2.18 Remark (Geometric mean alternative)

The concave representation theorem provides the cleanest route to Stage 2, but there is a structurally illuminating alternative that mirrors the geometric-mean homotopy used to establish contractibility of production technology spaces.

If utilities are normalized to be strictly positive (by adding a constant or working with $1 + u_i$), and if we restrict to the subclass of log-concave *utilities—those for which $\log u_i$ is concave—then the geometric mean*

$$u_t(x) := u_0(x)^{1-t} \cdot u_1(x)^t$$

preserves log-concavity, since $\log u_t = (1-t)\log u_0 + t\log u_1$ is a convex combination of concave functions. Log-concavity implies quasiconcavity, so u_t remains admissible. This is the direct analogue of the production technology homotopy (Carroll, 2026, Chapter 4) $\tau_t = (1+\tau_0)^{1-t}(1+\tau_1)^t - 1$, which preserves log-concavity through geometric interpolation. The parallel suggests a deeper structural unity between the supply side (force production) and the demand side (exchange under threat) of the state's behavior.

2.19 Remark (What contractibility means)

Contractibility says that the space of possible behaviors under threat has no topological obstructions: any state's demand-and-deployment behavior can be continuously deformed into any other's. There are no isolated clusters of behaviors, no holes in the space that would prevent smooth transitions, no topological invariants that distinguish one family of behaviors from another.

This is the demand-side analogue of the supply-side result that the space of policy functions—the set of all possible ways a state can convert resources into force—is contractible (Carroll, 2026). Together, the two results say that the full behavioral profile of a state—how it produces force and *how it trades and fights—lives in a topologically trivial space. Any state can, in principle, be continuously transformed into any other, on both sides of the ledger.*

The implications for institutions are direct: if the space of behaviors has nontrivial topology, then some institutional transitions—smooth paths from one equilibrium behavior to another—would be topologically obstructed. Contractibility removes this concern. We return to this point in Chapter 8.

2.6 Summary

The state's problem under threat is the atom of the framework. A state chooses how much to buy at market and how much force to project, jointly, subject to a budget that reflects the opportunity cost of militarization. What it ultimately consumes is the sum of its market purchases and the spoils (or losses) of conflict—and the spoils themselves depend on what everyone has bought. This coupling between exchange and conflict is the core of the model.

The chapter has established the following.

1. The problem has a solution (Proposition 2.12), which is unique and smooth in the parameters (Proposition 2.13). Smoothness is derived, not assumed: it follows from the bordered Hessian of u_i at effective prices and the diminishing returns to militarization, via the Schur complement of the coupled Jacobian (Lemma 2.8).

2. The solution defines a demand function under threat f_i and a deployment function μ_i, both smooth. The demand function satisfies Walras' law (Lemma 2.14) and desirability (Lemma 2.15). These are the properties required by the price-dependent equilibrium manifold theory of Balasko (2009, Chapter 6).

3. The demand function is price-dependent in a non-standard way—not by analytic choice (as in Friedman's Marshallian construction) but because the world is dangerous. The strict demand function is not a diffeomorphism, but the extended demand function (which separates effective and budget prices) is (Remark 2.16).

4. The space of all possible demand-deployment behaviors is contractible (Theorem 2.17). Any state's behavior under threat can be continuously deformed into any other's, with no topological obstructions.

The next chapter defines the dual-competitive equilibrium—a set of prices and deployments at which all markets clear and all states are in mutual best response—and uses the properties developed here to establish its existence and structure.

Chapter 3

Dual-Competitive Equilibrium

The previous chapter developed the state's problem under threat and established that it has a unique, smooth solution: a demand function f_i and a deployment function μ_i. We recall the notation: I states, L goods, prices $p \in S_N \subset \mathbb{R}^L_{++}$ (numeraire-normalized), deployments $m = (m_1, \dots, m_I) \in M$, consumption spaces $X_i = \prod_\ell (b_i^\ell, \infty)$, and endowment profiles $\omega \in \Omega_0 := \mathbb{R}^{IL}_{++}$. This chapter assembles the individual solutions into a system. We define the dual-competitive economy as a collection of demand-deployment pairs (§3.1), state the equilibrium concept (§3.2), solve the Nash layer to reduce the system to a price-dependent exchange model (§3.3), and then develop the structure of the equilibrium manifold—its local and global geometry, the natural projection, and the existence and finiteness of equilibria (§3.5–§3.7).

3.1 The Dual-Competitive Economy

3.1.1 Demand-Deployment Pairs

The state's problem under threat produces, for each state i, a pair of smooth functions:

$$\begin{aligned} f_i &: S_N \times M_{-i} \times \Omega_0 \to X_i, && \text{(demand under threat)} \\ \mu_i &: S_N \times M_{-i} \times \Omega_0 \to M_i. && \text{(optimal deployment)} \end{aligned}$$

The demand function f_i records what state i buys at market; the deployment function μ_i records how much force it projects. Both depend on prices, on the deployments of others, and on the endowment profile. Together, the pair (f_i, μ_i) encodes state i's complete behavioral response to the competitive environment.

The pair (f_i, μ_i) is derived from a state character (u_i, c_i, s_i)—preferences, costs, and a battle function—via the state's problem (Chapter 2). But for the purposes of this chapter, we take the pairs as given and work with their properties directly. The relevant properties, all established in Chapter 2, are:

- **Smoothness**: f_i and μ_i are C^∞ (Proposition 2.13).
- **Walras' law**: $p \cdot f_i(p, m_{-i}, \omega) = p \cdot (\omega_i - c_i(m_i^*, p, \omega_i))$ (Lemma 2.14).
- **Desirability**: demand diverges as prices or endowments approach the boundary (Lemma 2.15).

3.1.2 The Economy

We can now define the object of study.

3.1 Definition (Dual-competitive economy)

A dual-competitive economy *is a list*

$$\mathcal{E} = \left((f_i, \mu_i, c_i)_{i=1}^{I},\ \omega\right)$$

consisting of:

1. *for each state i, a demand-deployment pair (f_i, μ_i) satisfying smoothness, Walras' law, and desirability, together with an opportunity cost function c_i satisfying Assumption 2.2; and*
2. *an endowment profile $\omega \in \Omega_0$.*

This parallels the construction in Balasko (2011), where an exchange economy is a list of individual demand functions together with endowments. The difference is that our "demand functions" are demand-deployment pairs, and they depend on prices in a non-standard way (through the battle function as well as the budget constraint).

We think of the behavioral data $(f_i, \mu_i, c_i)_{i=1}^{I}$ as *fixed*—they describe the states' characters—and the endowment profile ω as the *variable parameter*. The central question is: for a given ω, does there exist a set of prices and deployments at which all states' behaviors are mutually consistent?

3.1.3 Multiple Agents and the Dimension Count

There are I states, each with L goods. The economy has:

- $L - 1$ free prices (one is the numéraire),
- I deployment levels (one per state),
- IL endowment components (the parameter space).

The equilibrium conditions comprise $L - 1$ market-clearing equations (one is redundant by Walras' law) and I Nash conditions (one per state). The total number of equations is $L - 1 + I$, matching the number of free variables (p, m). The parameter space has dimension IL. When the equilibrium conditions are satisfied, the system has

$$\dim(p, m, \omega) - \dim(\text{equations}) = (L - 1 + I + IL) - (L - 1 + I) = IL$$

degrees of freedom. This is the dimension of the equilibrium manifold.

The count reveals the role of multiple agents. With $I = 1$, there is no trade and no strategic interaction; the lone state simply consumes its diminished endowment. With $I = 2$, there are $L - 1$ price equations and 2 Nash conditions, and the parameter space has dimension $2L$. As I grows, the Nash layer grows linearly (one equation per state), while the market-clearing layer remains fixed at $L - 1$ equations. The strategic complexity of the system scales with I, not with L.

3.2 The Equilibrium Concept

We are now in a position to state the central behavioral hypothesis of the framework. A dual-competitive equilibrium asks for two things at once: that markets clear and that no state regrets its deployment.

3.2 Definition (Dual-competitive equilibrium)

Given a dual-competitive economy $\mathcal{E}$, a tuple $(p, m) \in S_N \times M$ is a dual-competitive equilibrium *(DCE) if:*

1. ***Market clearing.*** *Aggregate demand equals aggregate diminished supply:*

$$\sum_i f_i(p, m_{-i}, \omega) = \sum_i (\omega_i - c_i(m_i, p, \omega_i)).$$

2. ***Mutual best response.*** *Each state's deployment is optimal given the others':*

$$m_i = \mu_i(p, m_{-i}, \omega) \quad \text{for all } i.$$

Condition (1) is Walrasian: at the prevailing prices, the total quantity of each good demanded across all states equals the total available after militarization costs have been paid. Condition (2) is Nash: no state can improve its payoff by unilaterally changing its deployment. The two conditions are linked by the state's problem under threat, in which f_i and μ_i are jointly determined.

The term "dual-competitive" reflects the two modes of competition. Exchange operates through the price system: states compete for goods by bidding against one another in markets. Conflict operates through the deployment profile: states compete for spoils by projecting force. A DCE is a configuration in which both forms of competition have reached a rest point simultaneously.

3.2.1 The Equilibrium Equation

By Walras' law, one of the L market-clearing equations is redundant. Dropping the L-th (numéraire) equation, define the *reduced excess demand*:

$$\overline{d}(p, m, \omega) := \sum_i f_i(p, m_{-i}, \omega) - \sum_i (\omega_i - c_i(m_i, p, \omega_i)) \in \mathbb{R}^{L-1},$$

retaining only the first $L - 1$ components. Define the *Nash residual*:

$$r_i(p, m, \omega) := m_i - \mu_i(p, m_{-i}, \omega), \qquad r := (r_1, \dots, r_I) \in \mathbb{R}^I.$$

The *equilibrium function* is

$$e(p, m, \omega) := (\overline{d}(p, m, \omega),\ r(p, m, \omega)) \in \mathbb{R}^{L-1+I}.$$

A tuple (p, m, ω) is a DCE if and only if $e(p, m, \omega) = 0$. Since f_i, μ_i, and c_i are all smooth, the equilibrium function e is smooth. The equilibrium set is

$$C := e^{-1}(0) \subseteq S_N \times M \times \Omega_0.$$

3.3 The Nash Layer

For fixed (p, ω), the Nash condition $m = \mu(p, m, \omega)$ is a fixed-point problem in M. We show that it has a unique solution, which allows us to eliminate m from the equilibrium equation entirely.

3.3.1 Deployment Contraction

The deployment best-response $\mu_i(p, m_{-i}, \omega)$ depends on the other states' deployments through the battle function. The sensitivity of state i's optimal deployment to a change in state j's deployment is governed by the cross-partial $D_{m_j}\mu_i$.

3.3 Lemma (Deployment contraction)

Under the conflict feedback bound (Assumption 2.6) and diminishing returns to militarization (Assumption 2.7), the best-response mapping $m \mapsto \mu(p, m, \omega)$ is a contraction in m, uniformly in (p, ω) on compacts.

Proof. Fix i and consider how μ_i responds to a change in m_j, $j \neq i$. The optimal deployment μ_i is defined implicitly by the coupled first-order condition (FOC 2*). Differentiating with respect to m_j via the implicit function theorem and solving for $D_{m_j}\mu_i$:

$$D_{m_j}\mu_i = -\frac{p \cdot (I + D_{x_i}s_i)^{-1} \cdot D_{m_j}D_{m_i}s_i}{p \cdot (I + D_{x_i}s_i)^{-1} \cdot D^2_{m_im_i}s_i - p \cdot D^2_{m_im_i}c_i}.$$

The denominator is bounded away from zero: it is the second derivative of the net return to deployment (the Schur complement σ_0 from Lemma 2.8), which is strictly negative by Assumption 2.7. The numerator is bounded by

$$\|p\| \cdot \frac{\|D_{m_j}D_{m_i}s_i\|}{1 - \|D_{x_i}s_i\|},$$

using the Neumann bound on the inverse $(I + D_{x_i}s_i)^{-1}$.

Summing across $j \neq i$, the off-diagonal row norm of $D_m\mu$ is

$$\sum_{j \neq i} |D_{m_j}\mu_i| \leq \frac{\|p\| \cdot \sum_{j \neq i}\|D_{m_j}D_{m_i}s_i\|}{(1 - \|D_{x_i}s_i\|) \cdot |\sigma_0|}.$$

The conflict feedback bound ensures the denominator is large (small $\|D_{x_i}s_i\|$ and strict $|\sigma_0|$); the cross-derivatives $D_{m_j}D_{m_i}s_i$ measure how much one state's spoils respond to another's deployment. When these cross-effects are small relative to the own diminishing returns, the row norm is less than 1, and the mapping is a contraction. ∎

3.4 Remark

The contraction condition has a clean economic reading: each state's optimal deployment is not too sensitive to what the others are doing. This holds when the battle function's cross-effects (how state j's deployment affects state i's spoils) are small relative to the own diminishing returns (how fast state i's net return to deployment declines). It is the multi-agent analogue of the standard uniqueness condition in contest models.

3.3.2 Nash Uniqueness and Smoothness

The contraction property delivers both existence and regularity of the Nash equilibrium.

3.5 Proposition (Nash uniqueness)

For each $(p, \omega) \in S_N \times \Omega_0$, there exists a unique deployment profile $m^*(p, \omega) \in M$ satisfying $m^* = \mu(p, m^*, \omega)$. The function m^* is smooth in (p, ω).

Proof. The contraction mapping theorem gives existence and uniqueness of the fixed point $m^*(p, \omega)$. Smoothness follows from the implicit function theorem applied to $r(p, m, \omega) = m - \mu(p, m, \omega) = 0$: the Jacobian $D_m r = I - D_m \mu$ is invertible because μ is a contraction (all eigenvalues of $D_m \mu$ have modulus less than 1, so $I - D_m \mu$ has all eigenvalues bounded away from zero). ∎

3.4 The Reduced Economy

Proposition 3.5 allows us to eliminate deployments from the equilibrium equation. Substituting the Nash equilibrium $m^*(p, \omega)$ into the market-clearing condition gives the *reduced excess demand*:

$$\widehat{d}(p, \omega) := \overline{d}(p,\ m^*(p, \omega),\ \omega).$$

A price-endowment pair (p, ω) is part of a DCE if and only if $\widehat{d}(p, \omega) = 0$. The deployment profile is then recovered as $m = m^*(p, \omega)$.

3.6 Lemma (Properties of the reduced excess demand)
The function $\widehat{d} : S_N \times \Omega_0 \to \mathbb{R}^{L-1}$ *is smooth, satisfies Walras' law, and satisfies desirability.*

Proof. Smoothness: $\widehat{d}$ is a composition of smooth functions. Walras' law: the unreduced excess demand satisfies $p \cdot d(p, m, \omega) = 0$ by aggregating the individual Walras' law (Lemma 2.14) across states; substituting $m = m^*(p, \omega)$ preserves the identity. Desirability: if $p^\ell \to 0$, individual demand in good ℓ diverges (Lemma 2.15), so aggregate excess demand diverges; the argument for $\omega_i^\ell \to (b_i^\ell)^+$ is analogous. ■

The reduced economy is now a price-dependent exchange model in precisely the sense of Balasko (2009, Chapter 6): I agents, L goods, a smooth excess demand function satisfying Walras' law and desirability, with the endowment profile ω as the parameter. The entire machinery of that chapter applies. What follows records the consequences.

3.5 The Equilibrium Manifold

Define the *equilibrium manifold*

$$E := \{(p, \omega) \in S_N \times \Omega_0 \mid \widehat{d}(p, \omega) = 0\}.$$

This is the set of all price-endowment pairs at which a DCE obtains.

3.5.1 Local Structure

The reduced excess demand $\widehat{d}$ is a smooth map from $S_N \times \Omega_0$ to $\mathbb{R}^{L-1}$. The equilibrium manifold is its zero set.

3.7 Proposition (Smooth submanifold)
E is a smooth submanifold of $S_N \times \Omega_0$ of dimension IL.

Proof. We apply the preimage theorem. It suffices to show that 0 is a regular value of $\widehat{d}$—that is, the Jacobian $D_{(p,\omega)}\widehat{d}$ has full rank $L-1$ at every $(p,\omega) \in E$.

Consider the partial Jacobian $D_\omega \widehat{d}$. The reduced excess demand is $\widehat{d}(p,\omega) = \sum_i f_i(p, m^*_{-i}, \omega) - \sum_i(\omega_i - c_i(m^*_i, p, \omega_i))$. Differentiating with respect to ω_j (the endowment of a single state j) and evaluating at a no-trade equilibrium (where $f_j = \omega_j - c_j$):

$$D_{\omega_j}\widehat{d}\big|_{\text{no-trade}} = D_{\omega_j} f_j - I_L = D_{w_j} f_j \cdot D_{\omega_j} w_j - I_L,$$

where $w_j = p \cdot (\omega_j - c_j)$ is diminished wealth. By Walras' law, $D_{w_j} f_j$ points in the direction of the price-normalized gradient of u_j, which has full rank. Since j can be chosen freely among the I states, the columns of $D_\omega \widehat{d}$ span $\mathbb{R}^{L-1}$. At non-no-trade equilibria, the argument generalizes by continuity (the rank condition is open). Hence 0 is a regular value, and $E = \widehat{d}^{-1}(0)$ is a smooth submanifold of codimension $L-1$ in the ambient space $S_N \times \Omega_0$ of dimension $L - 1 + IL$. Thus $\dim E = IL$. ∎

3.5.2 Global Structure

The global structure of E is determined by its fiber decomposition.

3.8 Definition (No-trade equilibrium)

A DCE (p, m^*, ω) is a no-trade equilibrium *if each state consumes its diminished endowment at market:*

$$f_i(p, m^*_{-i}, \omega) = \omega_i - c_i(m^*_i, p, \omega_i) \quad \textit{for all } i.$$

At a no-trade equilibrium, no goods change hands through exchange. States may still fight: spoils may be nonzero and deployments positive. The absence of trade does not imply the absence of conflict—it means only that each state's optimal market purchase happens to coincide with what it has left after paying the costs of militarization.

For each $(p, w) \in S_N \times \mathbb{R}^I_{++}$, where $w_i = p \cdot (\omega_i - c_i(m^*_i, p, \omega_i))$ is state i's diminished wealth, the *fiber* $F(p, w)$ consists of all endowment profiles ω such that $(p, \omega) \in E$ and each state has wealth w_i.

3.9 Proposition (Fiber structure)

Each fiber $F(p, w)$ is a nonempty linear manifold of dimension $(L-1)(I-1)$. Every fiber contains exactly one no-trade equilibrium.

Proof. Fix $(p, w) \in S_N \times \mathbb{R}^I_{++}$. The fiber $F(p, w)$ is the set of $\omega \in \Omega_0$ such that $\widehat{d}(p, \omega) = 0$ and $p \cdot (\omega_i - c_i(m_i^*, p, \omega_i)) = w_i$ for all i. The wealth constraints are I linear equations in IL endowment components, defining an affine subspace of dimension $IL - I = I(L-1)$. The market-clearing conditions $\widehat{d} = 0$ impose $L - 1$ additional equations on this subspace, and one is redundant (by Walras' law, market clearing in $L - 1$ goods implies clearing in the L-th). But since one state's endowment is determined by the others' through the aggregate resource constraint $\sum_i \omega_i = \sum_i f_i$, the effective number of free parameters within the fiber is $I(L-1) - (L-1) = (I-1)(L-1)$.

The fiber is a translate of a linear subspace (the constraints are affine in ω once (p, w) is fixed), so it is a linear manifold of dimension $(I-1)(L-1)$. It is nonempty because the no-trade allocation $\omega_i = f_i(p, m^*_{-i}, \omega) + c_i(m_i^*, p, \omega_i)$ with $p \cdot \omega_i = w_i + p \cdot c_i$ satisfies all constraints. This no-trade allocation is the unique element of $F(p, w) \cap T$ (since at no-trade, ω_i is determined by (p, w_i) and the demand function). ∎

The no-trade equilibria form a smooth submanifold $T \subset E$ diffeomorphic to $S_N \times \mathbb{R}^I_{++}$. The manifold T is the geometric spine of E: every equilibrium lies in a unique fiber, and every fiber is pinned to a unique element of T.

3.10 Proposition (Diffeomorphism)
E is diffeomorphic to $\mathbb{R}^{IL}$.

Proof. Define the map $\Phi : E \to S_N \times \mathbb{R}^I_{++} \times \mathbb{R}^{(L-1)(I-1)}$ by

$$\Phi(p, \omega) = (p,\, w,\, \overline{\omega}_1, \ldots, \overline{\omega}_{I-1}),$$

where $w_i = p \cdot (\omega_i - c_i(m_i^*, p, \omega_i))$ is state i's diminished wealth and $\overline{\omega}_i$ denotes the first $L - 1$ components of ω_i (the L-th is determined by the wealth constraint $p \cdot \omega_i = w_i + p \cdot c_i$). The state I's endowment is determined by the aggregate resource constraint $\omega_I = \sum_i f_i - \sum_{j<I} \omega_j + c_I$.

The map Φ is smooth (composition of smooth functions), and one can verify directly that it is bijective: given $(p, w, \overline{\omega}_1, \ldots, \overline{\omega}_{I-1})$, each ω_i is recovered from $(p, w_i, \overline{\omega}_i)$, and ω_I from the resource constraint. The inverse is smooth by the same argument. Hence Φ is a diffeomorphism. The target space is $S_N \times \mathbb{R}^I_{++} \times \mathbb{R}^{(L-1)(I-1)} \cong \mathbb{R}^{L-1} \times \mathbb{R}^I \times \mathbb{R}^{(L-1)(I-1)} = \mathbb{R}^{IL}$. ∎

3.11 Definition (Peaceful equilibrium)
A DCE (p, m^*, ω) *is* peaceful *if conflict is entirely absent:* $s_i(m^*, x^*, \omega) = 0$ *and* $m_i^* = 0$ *for all* i.

A peaceful equilibrium is a no-trade equilibrium in which the deployment is also zero (by centeredness of c_i: zero deployment costs nothing, and without spoils there is no reason to deploy). Consumption equals the full endowment: $f_i = \omega_i$. This is a standard Walrasian equilibrium of the underlying exchange economy, undistorted by conflict. We return to peaceful equilibria in Chapter 5.

3.5.3 Topological Consequences

Since E is diffeomorphic to $\mathbb{R}^{IL}$, it inherits all the topological properties of Euclidean space:

3.12 Corollary
The equilibrium manifold E is connected, pathconnected, simply connected, and contractible.

These properties have institutional significance. Pathconnectedness means any two equilibria can be connected by a smooth path through equilibria—policies can be adjusted continuously from one rest point to another. Simple connectedness means any two such paths can be continuously deformed into each other—there are no topological obstructions to modifying a policy trajectory on the fly. Contractibility means the entire equilibrium manifold can be continuously deformed to a single point—there is no global topological structure to navigate. We develop these implications in Chapter 8.

3.6 The Natural Projection

The *natural projection* is the map

$$\pi : E \to \Omega_0, \qquad \pi(p, \omega) = \omega,$$

which sends each equilibrium to its underlying endowment profile. The fiber $\pi^{-1}(\omega)$ is the set of all equilibrium prices for a given distribution of resources.

The natural projection is the fundamental object for comparative statics: it tells us how the set of equilibria changes as endowments vary.

3.13 Proposition (Smooth and proper)

The natural projection $\pi : E \to \Omega_0$ is smooth and proper.

Proof. Smoothness: π is the restriction of the smooth coordinate projection $(p, \omega) \mapsto \omega$ to the smooth submanifold E.

Properness: let $K \subset \Omega_0$ be compact. We show $\pi^{-1}(K)$ is compact. Take any sequence $(p^n, \omega^n) \in \pi^{-1}(K)$. Since $\omega^n \in K$ and K is compact, ω^n has a convergent subsequence with limit $\omega^0 \in K \subset \Omega_0$. It remains to show that p^n also has a convergent subsequence with limit in S_N (i.e., prices do not escape to the boundary).

Market clearing gives $\sum_i f_i(p^n, m_{-i}^{*n}, \omega^n) = \sum_i(\omega_i^n - c_i(m_i^{*n}, p^n, \omega_i^n))$. The right side is bounded (since $\omega^n \in K$ and c_i is continuous with $c_i(0, \cdot, \cdot) = 0$, and m_i^{*n} is bounded because c_i is coercive in m_i). Hence $\sum_i f_i$ is bounded. Since each f_i takes values in X_i (which is bounded below), the individual demands $f_i(p^n, m_{-i}^{*n}, \omega^n)$ are bounded above as well.

Now suppose p^n has a subsequence converging to a boundary point $p^0 \in \partial S_N$ (some price $p^{0,\ell} = 0$). Desirability (Lemma 2.15) implies $\|f_i(p^n, m_{-i}^{*n}, \omega^n)\| \to \infty$, contradicting the boundedness just established. Hence every convergent subsequence of p^n has limit in the interior S_N, and $\pi^{-1}(K)$ is sequentially compact, hence compact. ∎

Properness is what makes the degree-theory argument for existence work. It also ensures that the number of equilibria cannot jump discontinuously as endowments vary (Proposition 3.16 below).

3.7 Existence and Finiteness

3.7.1 Existence

We now prove the main result.

3.14 Theorem (Existence of DCE)

A dual-competitive equilibrium exists for every $\omega \in \Omega_0$.

Proof. The argument uses the modulo-2 degree of the natural projection.

For $t \in [0, 1]$, define the scaled battle function $s_i^t := t \cdot s_i$. As in Stage 1 of the contractibility proof (Theorem 2.17), the scaled conflict technology satisfies all admissibility conditions for every t: the conflict feedback bound is preserved ($\|D_{x_i}(t \cdot s_i)\| \leq \|D_{x_i} s_i\|$), the domain constraints hold (by convexity of X_i), and the Nash contraction condition is maintained (the cross-derivatives of s_i scale by t, tightening the contraction).

The corresponding reduced excess demand $\widehat{d}^t(p, \omega)$ is smooth in (p, ω, t), satisfies Walras' law, and satisfies desirability—all uniformly in t, because scaling the battle function does not affect the boundary behavior of demand. The equilibrium manifold E^t and natural projection $\pi^t : E^t \to \Omega_0$ are therefore well-defined for every t.

At $t = 0$, conflict is absent. The reduced excess demand $\widehat{d}^0$ is the aggregate excess demand of a standard exchange economy with smooth, strictly quasiconcave, monotone preferences satisfying boundary blow-up. This is the classical smooth economy of Balasko (2011). Its natural projection π^0 has modulo-2 degree 1 (Balasko, 2009, Proposition 4.6.3).

The family $\{\pi^t\}_{t \in [0,1]}$ constitutes a proper homotopy from π^0 to $\pi^1 = \pi$: each π^t is smooth and proper (the desirability argument is uniform in t), and π^t varies smoothly in t. Since the modulo-2 degree is a homotopy invariant of proper maps,

$$\deg_2(\pi) = \deg_2(\pi^0) = 1.$$

A proper map of positive degree is surjective. Hence for every $\omega \in \Omega_0$, there exists $p \in S_N$ such that $(p, \omega) \in E$. Together with the Nash equilibrium $m^*(p, \omega)$ from Proposition 3.5, the tuple $(p, m^*(p, \omega), \omega)$ is a dual-competitive equilibrium. ∎

3.15 Remark (The homotopy as a thought experiment)

The proof scales conflict continuously from zero to its actual level. At $t = 0$, the world is peaceful and the standard existence theorem applies. As t increases, conflict is gradually "turned on," and the equilibrium manifold deforms continuously. The degree of the natural projection—a topological invariant counting, in a parity-adjusted sense, how many equilibria lie above each endowment profile—cannot change during this continuous deformation. Since it starts at 1, it remains 1, and existence is preserved.

This is more than a proof technique. It provides a conceptual picture: the equilibria of the dual-competitive economy are continuous deformations of the equilibria of the peaceful economy. Conflict distorts the price system and reallocates resources, but it does not create or destroy equilibria (in the modulo-2 sense).

3.7.2 Finiteness

Existence tells us the equilibrium set is nonempty. Properness constrains its size.

3.16 Proposition (Generic finiteness)

There exists a closed set $\Sigma \subset \Omega_0$ of Lebesgue measure zero such that for every $\omega \in \Omega_0 \setminus \Sigma$, the set of equilibrium prices $\pi^{-1}(\omega)$ is finite, and the number of equilibria is odd.

Proof. The natural projection π is smooth and proper (Proposition 3.13) with modulo-2 degree 1. By Sard's theorem, the set Σ of critical values of π is closed and has Lebesgue measure zero. At any regular value $\omega \in \Omega_0 \setminus \Sigma$, the preimage $\pi^{-1}(\omega)$ is a discrete set; properness makes it compact, hence finite. The modulo-2 count is 1, so the number of equilibria is odd. ■

3.17 Remark (Uniqueness)

The generic finiteness result says nothing about uniqueness: the number of equilibria is generically odd, but it need not be 1. Multiple equilibria are the rule, not the exception, in general equilibrium models with more than two goods and two agents (Mas-Colell, Whinston and Green, 1995). In our setting, multiplicity is compounded by the strategic interaction: different deployment profiles can support different market-clearing prices, and vice versa. The question of when the DCE is unique—and whether uniqueness can be guaranteed by restrictions on preferences, costs, or battle functions—is left for future work.

3.8 Summary

This chapter has established the following.

1. A dual-competitive economy is a list of demand-deployment pairs (one per state) together with an endowment profile (Definition 3.1). A dual-competitive equilibrium requires Walrasian market clearing and Nash mutual best response (Definition 3.2).

2. The Nash layer has a unique, smooth solution: for each price-endowment pair, there is a unique deployment profile in mutual best response, derived from the conflict feedback bound and diminishing returns to militarization (Proposition 3.5).

3. The reduced equilibrium—obtained by substituting the Nash solution into the market-clearing condition—is a price-dependent exchange model in the sense of Balasko (2009, Chapter 6) (Lemma 3.6).

4. The equilibrium manifold E is a smooth submanifold of $S_N \times \Omega_0$, diffeomorphic to $\mathbb{R}^{IL}$ (Propositions 3.7 and 3.10). It is connected, pathconnected, simply connected, and contractible (Corollary 3.12).

5. The natural projection $\pi : E \to \Omega_0$ is smooth and proper (Proposition 3.13).

6. A dual-competitive equilibrium *exists* for every endowment profile (Theorem 3.14). The proof constructs a proper homotopy from a peaceful economy to the dual-competitive economy by continuously scaling up the battle function, and uses the homotopy invariance of the modulo-2 degree. For generic endowments, the number of equilibria is finite and odd (Proposition 3.16).

7. Two classes of distinguished equilibria are defined. No-trade equilibria form the geometric spine of the equilibrium manifold. Peaceful equilibria—a subclass in which conflict is entirely absent—are the bridge to the normative theory.

The next chapter develops the positive theory: the local behavior of the equilibrium correspondence, stability, comparative statics, and the index theorem.

Chapter 4

The Positive Theory

The previous chapter established the structural facts about the equilibrium manifold E: it is a smooth submanifold of $S_N \times \Omega_0$, diffeomorphic to $\mathbb{R}^{IL}$, fibered over endowments by a smooth and proper natural projection π. These are global, topological facts. They say what the set of equilibria *is*. The present chapter asks how equilibria *behave*: how do they respond to changes in the world's resource distribution, and what does the distribution of their local properties look like?

The central object is the *price Jacobian*—the derivative of the reduced excess demand with respect to prices at an equilibrium. Its invertibility distinguishes well-behaved equilibria from degenerate ones; its structure encodes how exchange and conflict interact in determining the sensitivity of the price system. We introduce this object and decompose it into exchange and conflict components in §4.1, use it to classify international systems as regular or critical in §4.2, show that the regular case is generic in §4.3, prove that regular equilibria track parameters smoothly in §4.4, derive the comparative statics formula in §4.5, and establish the index theorem in §4.6.

4.1 The Price Jacobian

The equilibrium condition, after the Nash layer has been resolved, is $\widehat{d}(p, \omega) = 0$, where $\widehat{d} : S_N \times \Omega_0 \to \mathbb{R}^{L-1}$ is the reduced excess demand of §3.4. The behavior of the equilibrium correspondence near any solution depends fundamentally on how $\widehat{d}$ responds to changes in prices. We define the *price Jacobian*

$$J(p,\omega) := D_p\widehat{d}(p,\omega) \in \mathbb{R}^{(L-1)\times(L-1)},$$

the matrix of first derivatives of the reduced excess demand with respect to the $L-1$ free prices, evaluated at (p,ω).

Because $\widehat{d}$ is the excess demand of the reduced economy—obtained by substituting the Nash equilibrium deployment $m^*(p,\omega)$ into the aggregate excess demand—the price Jacobian J inherits effects from both exchange and conflict. Recall that

$$\widehat{d}(p,\omega) = \overline{d}(p,\, m^*(p,\omega),\, \omega),$$

where $\overline{d}(p,m,\omega) = \sum_i f_i(p,m_{-i},\omega) - \sum_i(\omega_i - c_i(m_i,p,\omega_i))$ is the unreduced excess demand. Differentiating by the chain rule:

$$J = D_p\overline{d}\big|_{m^*} + D_m\overline{d}\big|_{m^*} \cdot D_p m^*.$$

The two summands have distinct economic content. Define

$$J^{\text{ex}} := D_p\overline{d}\big|_{m^*} = \sum_i D_p f_i\big|_{m^*} + \sum_i D_p c_i\big|_{m^*},$$
$$J^{\text{conf}} := D_m\overline{d}\big|_{m^*} \cdot D_p m^*(p,\omega),$$

so that $J = J^{\text{ex}} + J^{\text{conf}}$.

The first term, J^{ex}, is the price response of excess demand with deployments held fixed at their Nash values. It captures the direct effects of price changes on demand (through the budget constraint and substitution effects) and on supply (through the Shephard's lemma derivative of opportunity costs). It is the analog of the aggregate Slutsky matrix in a standard exchange economy, evaluated at the conflict-adjusted equilibrium allocation.

The second term, J^{conf}, is new to the dual-competitive setting. It captures the indirect effects of price changes that flow through the conflict layer: a change in prices shifts the Nash equilibrium deployment profile by $D_p m^*$—which is given by differentiating the fixed-point condition $m^* = \mu(p,m^*,\omega)$ to obtain $D_p m^* = (I - D_m\mu)^{-1} D_p\mu$, invertible because μ is a contraction (Lemma 3.3)—and these deployment shifts alter excess demand by $D_m\overline{d}$. The matrix J^{conf} is therefore the product of two objects: how sensitive equilibrium deployments are to prices ($D_p m^*$), and how sensitive excess demand is to deployments ($D_m\overline{d}$).

4.1 Remark

The conflict term J^{conf} vanishes in two limiting cases. First, if the battle function does not depend on prices, so $D_p\mu_i = 0$ and hence $D_p m^* = 0$, then prices do not alter the strategic balance and $J^{\mathrm{conf}} = 0$. Second, in the peaceful economy ($s_i \equiv 0$), we have $m^* = 0$ and $D_p\mu = 0$, again giving $J^{\mathrm{conf}} = 0$ and $J = J^{\mathrm{ex}}$. In the peaceful case, the price Jacobian reduces to the standard exchange-economy Jacobian. In general, however, $J^{\mathrm{conf}} \neq 0$, and the conflict layer distorts the price sensitivity of the system.

The decomposition $J = J^{\mathrm{ex}} + J^{\mathrm{conf}}$ is the workhorse of the comparative statics analysis in §4.5. For now the key point is structural: J is a well-defined smooth function of (p, ω), since $\hat{d}$ is smooth (Lemma 3.6), and $J(p, \omega)$ varies smoothly across the equilibrium manifold.

4.2 Regular International Systems

Not all equilibria are created equal. At some, the price Jacobian J is invertible; at others, it is not. The distinction is fundamental, and it is worth understanding why before making it precise.

A singular J at (p^*, ω^*) means that some direction in price space leaves the reduced excess demand unchanged at first order. Geometrically, the equilibrium manifold E meets the fiber $\pi^{-1}(\omega^*)$ tangentially rather than transversally: the natural projection fails to be a local diffeomorphism at (p^*, ω^*). This has two consequences. Analytically, the implicit function theorem cannot be applied—there is no smooth local selection of equilibrium prices as a function of endowments near (p^*, ω^*). Economically, the world is at a knife edge: an arbitrarily small redistribution of resources may cause a discrete jump in equilibrium prices and deployments rather than the smooth adjustment one would expect.

Regular equilibria are precisely those at which this pathology is absent.

4.2 Definition (Regular DCE)

A dual-competitive equilibrium $(p^*, \omega^*) \in E$ is regular *if the price Jacobian $J(p^*, \omega^*)$ is nonsingular. It is* critical *otherwise.*

Equivalently, (p^*, ω^*) is regular if and only if the natural projection $\pi : E \to \Omega_0$ is a local diffeomorphism at (p^*, ω^*)—that is, the tangent map $T_{(p^*,\omega^*)}\pi :$

$T_{(p^*,\omega^*)}E \to T_{\omega^*}\Omega_0$ is an isomorphism. To verify the equivalence: a tangent vector $(dp, d\omega) \in T_{(p^*,\omega^*)}E$ must satisfy $J\,dp + D_\omega \widehat{d}\,d\omega = 0$, so the projection $(dp, d\omega) \mapsto d\omega$ is an isomorphism if and only if dp is uniquely determined by $d\omega$—which holds exactly when J is invertible.

4.3 Definition (Regular and critical endowments)

An endowment profile $\omega \in \Omega_0$ is regular *if every equilibrium above it is regular; it is* critical *otherwise. The* critical set *is*

$$\Sigma := \{\omega \in \Omega_0 : \omega \textit{ is a critical endowment}\}.$$

The critical set Σ collects all endowment profiles at which the equilibrium correspondence can fail to behave smoothly. Outside Σ, every equilibrium varies smoothly with the underlying distribution of resources. We turn to the question of how large Σ can be.

4.3 Genericity of Regular Systems

The critical set Σ is the collection of "bad" endowment profiles—those at which some equilibrium is degenerate. A priori, Σ could be a large and complicated set. The following result shows it is negligible.

4.4 Proposition (Generic regularity)

The critical set Σ is a closed subset of Ω_0 of Lebesgue measure zero. The complement $\Omega_0 \setminus \Sigma$ is open, dense, and of full measure.

Proof. The natural projection $\pi : E \to \Omega_0$ is smooth and proper (Proposition 3.13). The critical set Σ is the image under π of the set of critical points of π—those $(p^*, \omega^*) \in E$ at which $T\pi$ is not surjective, equivalently at which $J(p^*, \omega^*)$ is singular. By Sard's theorem applied to the smooth, proper map π, the set of critical values of π has Lebesgue measure zero. Hence Σ has measure zero.

For closedness: the set of critical points of π is closed in E, since it is the preimage of $\{0\}$ under the continuous function $(p, \omega) \mapsto \det J(p, \omega)$. Because π is proper, it maps closed sets in E to closed sets in Ω_0. Hence Σ is closed.

Since Σ is a closed set of measure zero in the connected open set Ω_0, its complement is open and has full measure. Denseness follows from openness and full measure in the connected space $\Omega_0 \cong \mathbb{R}^{IL}_{++}$. ∎

The proposition has a strong reading: the world is almost surely regular. Draw an endowment profile at random from any distribution that is absolutely continuous with respect to Lebesgue measure; the resulting international system is regular with probability one. Critical systems exist, but they are negligible.

4.5 Remark (The critical set as bifurcation locus)

Though negligible in measure, the critical set Σ is not uninteresting. It is the bifurcation locus *of the equilibrium correspondence: the set of endowment profiles at which equilibria are born or destroyed as parameters change. As ω crosses Σ transversally, the number of equilibria changes — always by an even number, since the modulo-2 count is preserved at the value 1 (Theorem 3.14). Locally and generically, Σ has the structure of a fold: two equilibria approach each other as $\omega \to \Sigma$, collide at Σ, and disappear on the other side. These are the tipping points of the international system—configurations at which a small shift in the world's resource distribution can force a discrete reorganization of prices and deployments, not because the world has changed dramatically, but because the system was already near a critical threshold.*

4.4 Stability

Generic regularity tells us that critical systems are rare. But what does regularity actually *buy*? The answer is stability in a precise sense: at a regular equilibrium, the entire equilibrium—prices and deployments alike—varies smoothly as endowments change.

4.4.1 Smooth selections

We formalize this in the following proposition. The claim is not just that prices move smoothly—it is that the entire equilibrium tuple does. The deployment selection $\widehat{m}$ is recovered from the price selection via the Nash layer, and smoothness propagates through that layer because m^* is itself smooth in (p, ω) (Proposition 3.5).

4.6 Proposition (Stability)

Let $(p^*, \omega^*) \in E$ be a regular dual-competitive equilibrium. Then there exist open neighborhoods $U_\omega \ni \omega^*$ in Ω_0 and $U_p \ni p^*$ in S_N, and unique smooth functions

$$\widehat{p} : U_\omega \to U_p \qquad \text{and} \qquad \widehat{m} : U_\omega \to M,$$

such that:

1. $\widehat{p}(\omega^*) = p^*$ *and* $\widehat{m}(\omega^*) = m^*(p^*, \omega^*)$*;*

2. *for each* $\omega \in U_\omega$*, the pair* $(\widehat{p}(\omega), \widehat{m}(\omega))$ *is the unique dual-competitive equilibrium in* $U_p \times M$*; and*

3. *both selections are* C^∞*.*

Proof. We apply the implicit function theorem to $\widehat{d}(p, \omega) = 0$ at (p^*, ω^*). Two conditions are required. First, (p^*, ω^*) must be a zero of $\widehat{d}$: this holds because (p^*, ω^*) is an equilibrium. Second, the partial Jacobian with respect to p must be nonsingular at (p^*, ω^*): this is $D_p\widehat{d}(p^*, \omega^*) = J(p^*, \omega^*)$, which is nonsingular by the regularity assumption.

The implicit function theorem yields open neighborhoods U_p and U_ω and a unique smooth function $\widehat{p} : U_\omega \to U_p$ satisfying $\widehat{d}(\widehat{p}(\omega), \omega) = 0$ for all $\omega \in U_\omega$, with $\widehat{p}(\omega^*) = p^*$. This establishes item (3) for the price selection.

The deployment selection is then $\widehat{m}(\omega) := m^*(\widehat{p}(\omega), \omega)$. Since m^* is smooth (Proposition 3.5) and $\widehat{p}$ is smooth, $\widehat{m}$ is smooth as a composition of smooth functions, giving item (3) for deployments. Item (1) holds by construction. For item (2): $(\widehat{p}(\omega), \widehat{m}(\omega))$ satisfies market clearing (since $\widehat{d}(\widehat{p}(\omega), \omega) = 0$) and Nash best response (since $\widehat{m}(\omega) = m^*(\widehat{p}(\omega), \omega)$, the unique Nash deployment at price $\widehat{p}(\omega)$ from Proposition 3.5). Uniqueness of the price selection in U_p follows from the uniqueness clause of the implicit function theorem; uniqueness of the deployment selection then follows from the global uniqueness of m^*. ■

The proposition says that at a regular equilibrium, the equilibrium *tracks* the parameter ω smoothly. A small shock to the world's resource distribution—a harvest failure that reduces a state's food endowment, a resource discovery that augments another's oil endowment—yields small, well-determined changes in both commodity prices and military deployments. The equilibrium does not jump. It moves.

It is worth isolating the smoothness claim on its own, since it is the form in which we will use the stability result downstream.

4.7 Corollary
At a regular DCE, the equilibrium prices $\widehat{p}(\omega)$ and deployments $\widehat{m}(\omega)$ are jointly smooth in ω. In particular, the full equilibrium is a smooth function of the underlying distribution of resources.

This is not merely a mathematical observation. The deployment selection $\widehat{m}(\omega)$ records how military postures respond to economic changes: a shift in endowments that moves prices will, through the Nash layer, also move deployments. Stability says this response is smooth and bounded, not catastrophic. It is the formal content of the intuition that small causes have small effects in a well-behaved international system.

4.4.2 Instability at the critical set

The stability result depends entirely on regularity. At a critical DCE—one where $J(p^*, \omega^*)$ is singular—the implicit function theorem cannot be applied, and the smooth selection generally does not exist. The equilibrium correspondence may fail to be a function near (p^*, ω^*): there may be multiple equilibria near p^* for some perturbations of ω^*, or equilibria may disappear entirely.

The critical set Σ is therefore the set of endowment profiles at which the world can be tipped: at which a small redistribution of resources may force a discrete jump in the system's equilibrium configuration. While Σ is negligible in measure (Proposition 4.4), it is the locus at which the positive theory's guarantees break down and the international system's behavior becomes structurally sensitive to the distribution of resources.

4.5 Comparative Statics

Stability guarantees that the equilibrium moves smoothly with endowments; comparative statics asks in which direction it moves and by how much. At a regular DCE, the stability result furnishes smooth selections $\widehat{p}(\omega)$ and $\widehat{m}(\omega)$. The question becomes: what are the derivatives $D_\omega \widehat{p}$ and $D_\omega \widehat{m}$?

4.5.1 The comparative statics formula

We begin with prices. Differentiate the identity $\widehat{d}(\widehat{p}(\omega), \omega) \equiv 0$ with respect to ω and apply the chain rule:

$$J(p^*, \omega^*) \cdot D_\omega \widehat{p} + D_\omega \widehat{d}(p^*, \omega^*) = 0.$$

Since (p^*, ω^*) is regular, J is invertible, and the equation can be solved.

4.8 Theorem (Comparative statics)
At a regular dual-competitive equilibrium (p^*, ω^*), the derivatives of the equilibrium price and deployment selections are

$$\begin{aligned} D_\omega \widehat{p} &= -J(p^*, \omega^*)^{-1} \cdot D_\omega \widehat{d}(p^*, \omega^*), \\ D_\omega \widehat{m} &= D_p m^*(p^*, \omega^*) \cdot D_\omega \widehat{p} + D_\omega m^*(p^*, \omega^*). \end{aligned}$$

Proof. The price formula follows from solving the differentiated identity as above. For deployments: differentiate $\widehat{m}(\omega) = m^*(\widehat{p}(\omega), \omega)$ with respect to ω and apply the chain rule to get the stated formula. ∎

The formula has the standard structure of comparative statics in smooth systems: the price response is the product of an inverse sensitivity matrix $-J^{-1}$ and a direct impact vector $D_\omega \widehat{d}$. What distinguishes the dual-competitive formula from its Walrasian analogue is the content of both factors. We examine each in turn.

4.5.2 The endowment impact and the conflict channel

The direct impact $D_\omega \widehat{d}$ measures how the equilibrium condition shifts when endowments change while prices are held fixed. Because $\widehat{d}$ is the reduced excess demand—with Nash deployments already substituted in—this derivative accounts for endowment changes not only through their direct effects on wealth and supply, but also through any shifts they induce in the equilibrium deployment profile.

To see the structure concretely, consider a perturbation to a single state j's endowment, holding all other endowments fixed. Differentiating $\widehat{d}(p, \omega) = \sum_i f_i(p, m^*_{-i}(p, \omega), \omega_i) - \sum_i (\omega_i - c_i(m^*_i(p, \omega), p, \omega_i))$ with respect to ω_j, and using $D_{\omega_j} m^* = (I - D_m \mu)^{-1} D_{\omega_j} \mu$ for the Nash response (from differentiating the fixed-point condition):

$$D_{\omega_j} \hat{d} = \underbrace{D_{\omega_j} f_j \big|_{m^*}}_{\text{(i) wealth}} + \underbrace{\sum_i D_{m_{-i}} f_i \cdot D_{\omega_j} m^*_{-i} + \sum_i D_{m_i} c_i \cdot D_{\omega_j} m^*_i}_{\text{(ii) Nash adjustment}} + \underbrace{D_{\omega_j} c_j \big|_{m^*} - I_{L-1}}_{\text{(iii) supply}} . \tag{4.1}$$

The three terms carry distinct interpretations.

1. The *wealth effect*, $D_{\omega_j} f_j|_{m^*}$: a larger endowment raises state j's wealth at any given prices, shifting its demand upward. This is the standard Walrasian income effect, present regardless of conflict.

2. The *Nash adjustment*, $\sum_i D_{m_{-i}} f_i \cdot D_{\omega_j} m^*_{-i} + \sum_i D_{m_i} c_i \cdot D_{\omega_j} m^*_i$: the endowment change alters the distribution of resources, which shifts the strategic calculus of all states. Their Nash equilibrium deployments respond by $D_{\omega_j} m^*$: some states arm more, others arm less, depending on how the endowment change affects the relative returns to deployment. These deployment shifts feed back into both demand (states respond to the altered threat environment through their consumption decisions) and supply (deployment levels determine opportunity costs through c_i). This term is unique to the dual-competitive setting; it vanishes in a peaceful economy.

3. The *supply effect*, $D_{\omega_j} c_j|_{m^*} - I_{L-1}$: the endowment increase raises state j's supply directly ($-I_{L-1}$), partially offset by any change in how much of the endowment is consumed by military activity ($D_{\omega_j} c_j$). The net supply response depends on whether the additional resources are also military inputs.

The Nash adjustment term is the heart of the matter. Its presence means that even a change in endowments that, holding deployments fixed, would have a simple and predictable effect on prices can have a more complex effect once the strategic response of all states is accounted for.

4.9 Remark (The Walrasian benchmark)

In a peaceful economy ($m^* = 0$, $s_i \equiv 0$), the Nash adjustment vanishes identically ($D_{\omega_j} m^* = 0$ when $D_{\omega_j} \mu_i = 0$), opportunity costs vanish ($c_i = 0$), and the formula *reduces to the standard Walrasian comparative statics:*

$$D_\omega \widehat{p}^{\text{Walras}} = -(D_p \widehat{d}^{\text{Walras}})^{-1}(D_{\omega_j} f_j - I_{L-1}),$$

with only the wealth and supply effects present. The dual-competitive formula adds the Nash adjustment to the endowment impact and replaces $D_p \widehat{d}^{\text{Walras}}$ *with* $J = J^{\text{ex}} + J^{\text{conf}}$ *in the denominator. The deviation of* $D_\omega \widehat{p}^{\text{DCE}}$ *from* $D_\omega \widehat{p}^{\text{Walras}}$ *is entirely a product of the conflict layer.*

The implication is striking: two otherwise identical international systems, one peaceful and one dual-competitive, will not only have different equilibrium prices but will respond *differently* to the same exogenous shock. When state j discovers oil, the Walrasian formula predicts a price response driven by wealth and supply effects. The dual-competitive formula adds a second channel: the discovery changes the strategic calculus of all states—oil-rich states become more attractive targets, militarization levels adjust across the system—and this strategic adjustment feeds back into commodity markets. The dual-competitive framework does not merely impose a constant offset on Walrasian prices; it alters the *structure* of how prices respond to the world.

4.5.3 Comparative statics of deployments

The deployment comparative statics follow from the price formula together with the Nash layer. From Theorem 4.8:

$$D_\omega \widehat{m} = D_p m^* \cdot D_\omega \widehat{p} + D_\omega m^*.$$

The first term is the indirect effect: price changes induced by the endowment shock alter Nash deployments through $D_p m^*$. The second term is the direct effect: the endowment shock changes the strategic calculus holding prices fixed. Both terms are smooth functions of (p^*, ω^*), and together they give a complete picture of how the military balance responds to any parameter change.

The deployment comparative statics have immediate substantive content. An increase in state j's endowment may raise or lower deployments, depending on whether the endowment change makes state j a more attractive target (raising others' incentives to arm), makes state j richer and thus better able to arm (raising j's own deployment), or shifts prices in ways that alter the opportunity costs of militarization throughout the system. The formula organizes these competing forces into a single expression, whose sign and magnitude can be evaluated parametrically in the applications of Chapter 6.

4.6 The Index

The finiteness result of Chapter 3 establishes that the number of equilibria is generically finite and odd. This is already a strong statement, but it leaves open questions of structure: are all odd numbers realized? Do equilibria of different "local character" alternate in a systematic way? The price Jacobian allows us to go further. To each regular equilibrium, we can assign an index that measures its local orientation relative to the natural projection. The sum of these indices over all equilibria is a fixed invariant of the system, and the invariant turns out to be 1—with rich consequences.

4.6.1 The index at a regular equilibrium

At a regular DCE $(p^*, \omega^*) \in E$, the natural projection π is a local diffeomorphism. Since $E \cong \mathbb{R}^{IL}$ (Proposition 3.10), E carries a natural orientation induced by the global diffeomorphism $\Phi : E \to \mathbb{R}^{IL}$. With $\Omega_0 = \mathbb{R}^{IL}_{++}$ oriented by the standard Euclidean orientation, the local diffeomorphism $T_{(p^*,\omega^*)}\pi$ has a well-defined sign.

4.10 Definition (Index)
The index *of a regular DCE* $(p^*, \omega^*) \in E$ *is*

$$\operatorname{ind}(p^*, \omega^*) := (-1)^{L-1} \operatorname{sign} \det J(p^*, \omega^*) \in \{+1, -1\}.$$

The factor $(-1)^{L-1}$ is an orientation correction: it ensures that the index is $+1$ at the unique equilibrium of the peaceful economy, where $J = J^{\text{ex}}$ is the standard exchange-economy price Jacobian (which satisfies $(-1)^{L-1} \det J^{\text{ex}} > 0$ at any regular equilibrium of a smooth economy, by the classical index result of Balasko (2009)).

An equilibrium of index $+1$ is *stable*: near it, the natural projection is orientation-preserving, and the equilibrium price responds to parameter perturbations in the locally expected direction. An equilibrium of index -1 is *unstable*: the natural projection reverses orientation, and the equilibrium is surrounded, locally, by configurations with more and fewer equilibria. In the familiar two-state, two-good case ($I = L = 2$), a positive-index equilibrium is one at which demand exceeds supply below the equilibrium price and supply exceeds demand above it—the standard partial-equilibrium stability condition. The index generalizes this to the multi-state, multi-good, dual-competitive setting.

4.6.2 The index theorem

The sum of indices over all equilibria above a regular endowment is a global invariant, and we can compute it exactly.

4.11 Theorem (Index theorem)

For every regular endowment $\omega \in \Omega_0 \setminus \Sigma$,

$$\sum_{\substack{(p^*,\omega')\in E \\ \omega'=\omega}} \operatorname{ind}(p^*, \omega) = 1.$$

Proof. Since $E \cong \mathbb{R}^{IL}$ is orientable and $\pi : E \to \Omega_0$ is smooth and proper (Proposition 3.13), the integer degree $\deg(\pi)$ of π with respect to the orientations on E and Ω_0 is well-defined. For any regular value $\omega \in \Omega_0 \setminus \Sigma$, the integer degree equals the signed count of preimages, which by Definition 4.10 is exactly $\sum \operatorname{ind}(p^*, \omega)$. It therefore suffices to show $\deg(\pi) = 1$.

Consider the homotopy $\pi^t : E^t \to \Omega_0$ used in the existence proof (Theorem 3.14), where the battle function is scaled as $s_i^t = t \cdot s_i$ for $t \in [0, 1]$. As shown there, each π^t is smooth and proper (desirability is uniform in t), so the homotopy is a proper homotopy. The orientation on E^t is induced by the diffeomorphism $\Phi^t : E^t \to \mathbb{R}^{IL}$; since the fiber structure of E^t is preserved under scaling (the no-trade spine and fiber decomposition of Proposition 3.9 apply for every t), Φ^t varies smoothly in t and the orientation is constant along the homotopy. The integer degree is therefore constant in t.

At $t = 0$, the economy is a standard smooth exchange economy. By Balasko (2009, Proposition 4.6.3), its natural projection has integer degree 1 (with the orientation convention above). Hence $\deg(\pi^t) = 1$ for all $t \in [0, 1]$, and in particular $\deg(\pi) = \deg(\pi^1) = 1$. ∎

The proof is a straightforward consequence of the homotopy constructed in Chapter 3, now used to track not just the modulo-2 count but the full integer degree. The orientation argument is what upgrades the modulo-2 result to an integer one: because the diffeomorphism $E^t \cong \mathbb{R}^{IL}$ is stable under the scaling homotopy, no orientation reversal occurs, and the degree is anchored at 1 throughout.

From the theorem, a short algebraic argument extracts the precise structure of the equilibrium set at any regular endowment.

4.12 Corollary (Index consequences)

For every regular $\omega \in \Omega_0 \setminus \Sigma$*:*

1. *The number of stable equilibria (index* $+1$*) exceeds the number of unstable equilibria (index* -1*) by exactly* 1.
2. *At least one stable equilibrium exists.*
3. *The number of equilibria is odd (recovering Proposition 3.16).*

Proof. Let n_+ and n_- denote the number of positive- and negative-index equilibria respectively. The index theorem gives $n_+ - n_- = 1$, which is (1). Item (2) follows since $n_- \geqslant 0$ implies $n_+ = 1 + n_- \geqslant 1$. For (3): $n_+ + n_- = (n_+ - n_-) + 2n_- = 1 + 2n_-$, which is odd. ■

The corollary upgrades the finiteness result from Chapter 3 in two ways. First, it provides the exact algebraic structure of the equilibrium set: stable and unstable equilibria alternate, with stable ones always ahead by exactly one. Second, and more substantially, it guarantees the existence of a *stable* equilibrium, not merely any equilibrium. The existence theorem showed that at least one equilibrium exists for every endowment profile; the index theorem shows that at least one of these equilibria is locally well-behaved—that is, the price system responds smoothly and in the expected direction to small parameter changes near it.

4.13 Remark (Multiplicity and the index)

The corollary says nothing about uniqueness, and uniqueness is not the generic case in multi-state, multi-good general equilibrium models (Mas-Colell, Whinston and Green, 1995). When $n_- \geqslant 1$*, there are multiple equilibria, and they come in pairs of one stable and one unstable. Economically, multiple equilibria represent distinct self-consistent configurations of the international system: different price vectors and military postures, all supported by the same underlying distribution of resources and state characters. Which equilibrium the world finds itself at depends on history and coordination, not on the parameters alone. The index theorem constrains this multiplicity but does not resolve it; the applications in Chapter 6 will examine cases with unique equilibria, where these questions are set aside.*

4.7 Summary

This chapter has developed the positive theory of dual-competitive equilibrium: how equilibria behave as a function of the endowment parameter.

1. The *price Jacobian* $J = D_p\widehat{d}$ is the key local object. It decomposes as $J = J^{\mathrm{ex}} + J^{\mathrm{conf}}$: an exchange component that captures the direct price response of excess demand at fixed deployments, and a conflict component that captures the indirect response flowing through changes in the Nash deployment profile (§4.1).

2. A DCE is *regular* if J is nonsingular; the *critical set* Σ collects all endowment profiles at which some equilibrium is degenerate. Regularity is equivalent to the natural projection being a local diffeomorphism at the equilibrium (§4.2).

3. Regular equilibria are *generic*: Σ is closed with Lebesgue measure zero, and regular endowments form an open dense set of full measure (Proposition 4.4). The critical set is the bifurcation locus of the equilibrium correspondence, where tipping-point dynamics can occur.

4. At a regular DCE, the equilibrium *tracks parameters smoothly*: unique smooth price and deployment selections $\widehat{p}(\omega)$ and $\widehat{m}(\omega)$ exist in a neighborhood of ω^*, and small endowment perturbations produce small, well-determined changes in both (Proposition 4.6).

5. The *comparative statics formula* is $D_\omega\widehat{p} = -J^{-1}D_\omega\widehat{d}$. The endowment impact $D_\omega\widehat{d}$ decomposes into a Walrasian wealth effect, a supply effect, and a Nash adjustment unique to the dual-competitive setting (Theorem 4.8, (4.1)). The Nash adjustment is the general equilibrium footprint of conflict: endowment changes alter the strategic balance, shift deployments, and feed back into commodity markets. The dual-competitive formula departs from its Walrasian counterpart in both the inverse matrix and the endowment impact.

6. The *index theorem* establishes that the sum of indices over all equilibria above any regular endowment equals 1 (Theorem 4.11). Stable equilibria outnumber unstable ones by exactly 1; at least one stable equilibrium always exists; and the total count is odd (Corollary 4.12). The proof

anchors the integer degree at 1 via the homotopy to the peaceful economy and the stability of the orientation under scaling.

The next chapter turns to the normative theory: given that dual-competitive equilibria exist, behave well, and respond systematically to parameter changes, what do they imply for welfare?

Chapter 5

The Normative Theory

The preceding chapter characterized the behavior of dual-competitive equilibria: their local structure, stability, comparative statics, and the global count index. The present chapter turns to evaluation: are these equilibria *good*? The answer has a sharp structure.

Every DCE is efficient in a qualified sense—efficient conditional on the prevailing military deployments. This is a weakened First Welfare Theorem: the exchange mechanism, given the strategic environment it faces, resolves the allocation problem without waste. However, the qualification is not cheap. Military deployments are themselves chosen by Nash logic, and this generates a *conflict externality*: each state's militarization imposes resource costs and strategic distortions on others that no price signal corrects. DCEs are therefore generically *not* globally Pareto optimal—there exist feasible alternatives, achievable by collectively demilitarizing, that would make everyone better off.

The exception is striking. When warfare is zero- or negative-sum—when conflict cannot create aggregate resources, only destroy or redistribute them—any peaceful DCE is globally efficient. In that case the invisible hand of the market, operating on a world at peace, delivers a first-best outcome. The failure of the Second Welfare Theorem then sharpens the normative picture: even when an efficient allocation exists, it is generally not supportable as a DCE via redistribution of endowments alone, because the Nash condition on military deployment is not correctable by prices.

This chapter develops each of these results with precision, using the smooth structure of the Pareto set—following Balasko (2009, Appendix A)—to give each claim its sharpest geometric form. Section 5.1 develops the Pareto-set toolkit. Section 5.2 establishes conditional efficiency and the welfare weight representation. Section 5.3 characterizes the conflict externality and the efficiency gap. Section 5.4 proves the efficiency of peaceful equilibria under zero-sum warfare. Section 5.5 discusses the failure of the Second Welfare Theorem. Section 5.6 summarizes.

5.1 The Structure of Pareto Optima

Before studying the normative properties of equilibria, we need a precise account of the Pareto optima available to the system. This section develops the requisite toolkit, drawing on Balasko (2009, Appendix A) and adapting its notation to our setting.

Let $r \in X := \prod_i X_i$ denote a vector of *total resources* available for consumption. The *resource-feasible set* at r is

$$\Omega(r) := \left\{ x = (x_1, \ldots, x_I) \in X \mid \sum_i x_i = r \right\}.$$

An allocation $x \in \Omega(r)$ is *Pareto optimal* if no $y \in \Omega(r)$ satisfies $u_i(y_i) \geq u_i(x_i)$ for all i with strict inequality for some j. Let $\mathcal{P}(r)$ denote the set of Pareto optima.

The *utility possibility set* at r tracks what utility vectors are achievable. Define the open and closed versions:

$$\begin{aligned} \mathcal{U}(r) &:= \{ (u_1(x_1), \ldots, u_I(x_I)) \in \mathbb{R}^I \mid \textstyle\sum_i x_i < r,\ x \in X \}, \\ \overline{\mathcal{U}}(r) &:= \mathrm{cl}(\mathcal{U}(r)), \end{aligned}$$

where the strict inequality $\sum_i x_i < r$ holds componentwise (incomplete resource utilization). The *utility possibility frontier* is $\partial\overline{\mathcal{U}}(r) := \overline{\mathcal{U}}(r) \setminus \mathcal{U}(r)$.

The following structural facts are established in Balasko (2009, Appendix A) under Assumption 2.1.

5.1 Proposition (Structure of Pareto optima)

For each $r \in X$:

1. (Manifold) *$\mathcal{P}(r)$ is a smooth submanifold of $\Omega(r)$ diffeomorphic to $\mathbb{R}^{I-1}$.*
2. (Frontier) *The utility map $x \mapsto (u_1(x_1), \ldots, u_I(x_I))$ restricts to a diffeomorphism from $\mathcal{P}(r)$ onto $\partial\overline{\mathcal{U}}(r)$.*

3. (Boundary condition) *A utility vector u belongs to $\partial\overline{\mathcal{U}}(r)$ if and only if $(u + \mathbb{R}^I_+) \cap \overline{\mathcal{U}}(r) = \{u\}$: no coordinate can be increased while remaining feasible.*

4. (Diffeomorphism) *$\mathcal{U}(r)$ is diffeomorphic to $\mathbb{R}^I$.*

Proof. Part (1) is Balasko 2009, Proposition A.6.2; part (2) is Proposition A.8.1; part (3) is Proposition A.1.4; part (4) is Proposition A.5.6, proved by induction using $\mathcal{U}_{k+1}(r) \cong \mathcal{U}_k(r) \times (0, \infty)$ (Proposition A.5.3) and $\mathcal{U}_1(r) \cong \mathbb{R}$. ■

Part (2) says that the utility possibility frontier and the Pareto set are in smooth one-to-one correspondence: every point on the frontier corresponds to exactly one Pareto optimal allocation, and vice versa. Part (3) characterizes the frontier geometrically: you are on the frontier precisely when you cannot improve any state's utility without making the bundle infeasible.

Two parameterizations of $\mathcal{P}(r)$ will play important roles. The first is free.

5.2 Proposition (Parameterization by utility levels)

The smooth map $R_{I-1}(r, \cdot) : \mathcal{U}_{I-1}(r) \to \Omega(r)$ of Balasko (2009, Proposition A.4.1) associates to each utility vector $(u_1, \ldots, u_{I-1}) \in \mathcal{U}_{I-1}(r)$ a unique Pareto optimal allocation $R_{I-1}(r, u_1, \ldots, u_{I-1}) \in \mathcal{P}(r)$. This defines a diffeomorphism between $\mathcal{U}_{I-1}(r)$ and $\mathcal{P}(r)$.

The Pareto set is therefore parameterized by the utility levels of any $I - 1$ consumers: specifying $(u_1, \ldots, u_{I-1})$ pins down a unique Pareto optimal allocation, with state I's utility determined as the highest achievable given the others' levels and the resource constraint.

The second parameterization uses welfare weights and requires stronger preferences.

5.3 Assumption (Strict concavity)

For each i, $u_i : X_i \to \mathbb{R}$ is strictly smoothly concave: the Hessian matrix $D^2 u_i(x_i)$ is negative definite for all $x_i \in X_i$.

Strict concavity strengthens the strict quasiconcavity in Assumption 2.1. It is needed for the welfare analysis because it ensures that the social welfare function $\sum_i \lambda_i u_i$ is globally concave in aggregate allocations—a property that quasi-concavity does not deliver. It also implies the strict convexity of the utility possibility set.

5.4 Proposition (Parameterization by welfare weights)
Under Assumption 5.3, the maximization problem

$$\max_{x\in\Omega(\leqslant r)}\left[\sum_{i=1}^{I-1}\lambda_i\, u_i(x_i)+u_I(x_I)\right]$$

has a unique solution $T_{I-1}(r,\lambda)\in\mathcal{P}(r)$ for each $\lambda\in\mathbb{R}^{I-1}_{++}$. The map $\lambda\mapsto T_{I-1}(r,\lambda)$ is a diffeomorphism from $\mathbb{R}^{I-1}_{++}$ onto $\mathcal{P}(r)$.

Proof. Balasko (2009, Propositions A.9.2 and A.9.4). Strict concavity makes the restricted Lagrangean $\mathcal{L}(x,\lambda)=\sum_i\lambda_i u_i(x_i)+u_I(x_I)$ strictly concave in x, giving a unique maximizer. Smoothness and bijectivity of $\lambda\mapsto T_{I-1}(r,\lambda)$ follow from the nonvanishing Jacobian of the first-order conditions (Balasko Section A.3.2). ■

5.5 Proposition (Strict convexity of the utility possibility set)
Under Assumption 5.3, $\overline{\mathcal{U}}(r)$ is strictly convex.

Proof. Balasko (2009, Proposition A.9.5). For distinct $u,u'\in\overline{\mathcal{U}}(r)$ achieved by x,x' with $\sum x_i,\sum x_i'\leqslant r$, the midpoint $x''=(x+x')/2$ satisfies $\sum x_i''\leqslant r$ and $u_i(x_i'')\geqslant(u_i(x_i)+u_i(x_i'))/2$ by concavity, with strict inequality for at least one i (since $u\neq u'$). So $(u+u')/2\in\mathcal{U}(r)$, the interior of $\overline{\mathcal{U}}(r)$. ■

One further fact is crucial for the efficiency analysis: the utility possibility set grows strictly when resources increase.

5.6 Lemma (Monotone utility possibility)
If $r'<r$ componentwise, then $\overline{\mathcal{U}}(r')\subset\mathcal{U}(r)$. In particular, $\partial\overline{\mathcal{U}}(r')$ lies in the interior of $\overline{\mathcal{U}}(r)$.

Proof. Let $u\in\overline{\mathcal{U}}(r')$. There exists $x\in X$ with $\sum_i x_i\leqslant r'$ and $u_i(x_i)\geqslant u_i$ for all i. Since $r'<r$ componentwise, $\sum_i x_i\leqslant r'<r$ strictly, so the allocation x uses strictly less than r in at least one good. By smooth monotonicity (Assumption 2.1(3)), increase x_j slightly in that good for some j; this raises $u_j(x_j)$ strictly while keeping $\sum x_i<r$. The resulting utility vector dominates u in at least one coordinate, so $u\in\mathcal{U}(r)$ (the open interior). ■

Geometrically: more resources strictly expand the feasible utility set, and the frontier at lower resources lies strictly inside the frontier at higher resources. This embedding is the key to measuring the welfare cost of conflict.

5.2 Conditional Efficiency

With the geometric toolkit in place, we can state and prove the normative properties of DCE. The first result is the analogue of the First Welfare Theorem.

5.2.1 Total Consumption and the Conditional Frontier

At a DCE (p^*, x^*, m^*), the military deployments m^* reduce the endowments available for market exchange through the opportunity costs $c_i(m_i^*, p^*, \omega_i)$, while the battle outcomes $s_i(m^*, x^*, \omega)$ shift what states actually consume away from what they purchased. To analyze efficiency, the right object is not market purchases x_i but the total consumption that results.

5.7 Definition (Total consumption and effective resources)
At a DCE (p^*, x^*, m^*) at endowment profile ω:

1. *The* total consumption *of state i is $x_i'^* := x_i^* + s_i(m^*, x^*, \omega) \in X_i$.*
2. *The* effective endowment *of state i is $\omega_i'^* := \omega_i - c_i(m_i^*, p^*, \omega_i) + s_i(m^*, x^*, \omega)$.*
3. *The* effective resources *are $r^* := \sum_i x_i'^* = \sum_i \omega_i'^*$.*

Market clearing gives $\sum_i x_i^* = \sum_i(\omega_i - c_i(m_i^*, p^*, \omega_i))$, so

$$r^* = \sum_i \omega_i - \sum_i c_i(m_i^*, p^*, \omega_i) + \sum_i s_i(m^*, x^*, \omega).$$

Under zero-sum warfare ($\sum_i s_i = 0$): $r^* = \omega - \sum_i c_i(m^*) \leqslant \omega$, with equality if and only if $m^* = 0$. At a peaceful DCE, $c_i(0, p^*, \omega_i) = 0$ and $s_i(0, x^*, \omega) = 0$ (by Assumption 2.2 and the natural normalization that zero deployment yields zero transfers), giving $r^* = \omega$ and $x_i'^* = x_i^*$.

5.2.2 The First Welfare Theorem for DCE

The key question is whether the total consumption allocation $x'^* = (x_1'^*, \dots, x_I'^*)$ is Pareto optimal given the resources r^*. Conditional efficiency says yes—no reallocation of market purchases, holding the military and battle outcomes fixed, can Pareto-improve on the equilibrium.

5.8 Proposition (Conditional efficiency)
At any DCE (p^*, x^*, m^*), we have $x'^* \in \mathcal{P}(r^*)$.

Proof. Fix m^* and the equilibrium battle outcomes $s_i^* := s_i(m^*, x^*, \omega)$. We argue by contradiction. Suppose there exists a market-purchase profile $y = (y_1, \dots, y_I)$ in the *conditionally feasible set*

$$CF(m^*) := \left\{ y \in X \mid \sum_i y_i \leq \sum_i (\omega_i - c_i(m_i^*, p^*, \omega_i)) \right\}$$

such that $u_i(y_i + s_i^*) \geq u_i(x_i'^*)$ for all i with strict inequality for some j.

At the DCE, x_i^* is the unique solution to state i's problem (SPT$_i$) with $m_{-i} = m_{-i}^*$ and battle outcomes fixed at s_i^*: it maximizes $u_i(z + s_i^*)$ subject to $p^* \cdot z \leq w_i^* := p^* \cdot (\omega_i - c_i(m_i^*, p^*, \omega_i))$.

By uniqueness and the Weak Axiom of Revealed Preference: if $y_i \neq x_i^*$ and $u_i(y_i + s_i^*) \geq u_i(x_i^* + s_i^*)$, then y_i is not affordable, so $p^* \cdot y_i > w_i^*$. For state j with strict preference, this is immediate. Therefore

$$p^* \cdot y_i \geq w_i^* \text{ for all } i, \qquad p^* \cdot y_j > w_j^*.$$

Summing over i and applying Walras's Law ($\sum_i w_i^* = p^* \cdot \sum_i (\omega_i - c_i(m_i^*))$):

$$p^* \cdot \sum_i y_i > p^* \cdot \sum_i (\omega_i - c_i(m_i^*, p^*, \omega_i)).$$

■

Since $p^* \gg 0$, this forces $\sum_i y_i^\ell > \sum_i (\omega_i^\ell - c_i^\ell(m_i^*, p^*, \omega_i))$ for some good ℓ, contradicting $y \in CF(m^*)$.

The proof is the standard First Welfare Theorem argument, applied to the economy in which each state i maximizes $u_i(\cdot + s_i^*)$ over its budget set. Revealed preference does the rest. The market mechanism achieves conditional efficiency without any coordination device beyond prices.

The qualification "conditional" is substantive. Conditional efficiency excludes from comparison any outcome achievable by changing the military posture: it compares only alternative market-purchase profiles at fixed (m^*, s_i^*). Two states locked in a costly arms race may still be conditionally efficient if, given their forces, they trade rationally. The waste is in the arming itself—a waste that prices cannot correct. The richer question of whether the full outcome (x'^*, m^*) is efficient in a global sense is taken up in Section 5.3.

5.2.3 The Welfare Weight Representation

Conditional efficiency establishes that $x'^* \in \mathcal{P}(r^*)$. Proposition 5.4 then provides a dual representation: every Pareto optimal allocation has a unique set of welfare weights. Applied to x'^*:

5.9 Corollary (Welfare weight representation)

Under Assumption 5.3, at any DCE (p^*, x^*, m^*), there exist unique weights $\lambda^* \in \mathbb{R}^{I-1}_{++}$ such that

$$x'^* = T_{I-1}(r^*, \lambda^*).$$

That is, x'^* is the unique solution to the social welfare problem

$$\max_{x \in \Omega(\leq r^*)} \left[\sum_{i=1}^{I-1} \lambda_i^* \, u_i(x_i) + u_I(x_I) \right].$$

The weights λ_i^* are inversely proportional to the Lagrange multipliers μ_i^* on state i's budget constraint at the DCE.

Proof. Proposition 5.8 gives $x'^* \in \mathcal{P}(r^*)$. Proposition 5.4 gives the unique λ^* with $T_{I-1}(r^*, \lambda^*) = x'^*$. The relationship $\lambda_i^* \propto 1/\mu_i^*$ follows from the first-order conditions of the social welfare problem and the budget-constraint multipliers of the individual SPTs, which share the same optimality conditions at x'^*. ∎

The welfare weight representation says: every DCE behaves as though a social planner were maximizing a weighted utilitarian social welfare function over the effective resource base r^*. The planner is not real—the allocation is decentralized—but the equilibrium is *consistent with* such a planner, for the particular weights λ^* determined by the initial endowments, military costs, and battle outcomes.

The weights encode the distributional implications of the equilibrium. State i receives weight $\lambda_i^* \propto 1/\mu_i^*$: a high weight when its marginal utility of income is low (it is resource-rich) and a low weight when income is scarce. This is the competitive logic of the invisible hand: the market does not equalize welfare but produces welfare weights that reflect the competitive distribution of resources.

Moreover, the weights $\lambda^*(\omega)$ vary smoothly with the endowment profile at regular equilibria, because $p^*(\omega)$ and $m^*(\omega)$ are smooth selections by the stability results of Chapter 4. The normative content of the equilibrium—the distributional preference it embeds—is therefore itself a smooth function of the primitive resource distribution.

5.3 The Conflict Externality and Global Inefficiency

Conditional efficiency is a real but limited result. To evaluate the DCE more fully, we must compare it against outcomes achievable by varying military deployments as well. The Nash equilibrium over m is not a social optimum, because each state's militarization creates external effects—on others' battle outcomes and on the aggregate resource base—that market prices do not discipline.

5.3.1 The Social Planner's Conditions

An outcome (y, m') is *globally feasible* if $\sum_i y_i \leqslant \sum_i \omega_i - \sum_i c_i(m'_i, p, \omega_i) + \sum_i s_i(m', y, \omega)$ for some price vector p. Total consumption at (y, m') is $y'_i := y_i + s_i(m', y, \omega)$. A DCE outcome (x'^*, m^*) is *globally Pareto optimal* if no globally feasible (y, m') with y'_i satisfies $u_i(y'_i) \geqslant u_i(x'^*_i)$ for all i with at least one strict inequality.

A social planner choosing (y, m') to maximize $\sum_i \alpha_i u_i(y'_i)$ subject to global feasibility would require, at an interior optimum, the following first-order condition for state i's deployment:

$$\sum_{j=1}^{I} \alpha_j Du_j(y'_j) \cdot D_{m_i} s_j(m', y, \omega) = \alpha_i \mu'_i p \cdot D_{m_i} c_i(m'_i, p, \omega_i), \tag{5.1}$$

where $\mu_i^{\prime}$ is state i's shadow income. Compare this to the Nash condition (FOC 2) at the DCE, which requires only:

$$Du_i(x_i^{\prime *}) \cdot D_{m_i} s_i(m^*, x^*, \omega) \;=\; \mu_i^* \, p^* \cdot D_{m_i} c_i(m_i^*, p^*, \omega_i). \tag{5.2}$$

The Nash condition (5.2) equates state i's own marginal benefit of arming against its own marginal cost. The social planner's condition (5.1) adds the cross-effects: $\sum_{j \neq i} \alpha_j \, Du_j(y_j^{\prime}) \cdot D_{m_i} s_j$, the welfare consequences that state i's militarization imposes on all other states through their battle outcomes. This discrepancy is the *conflict externality*.

5.3.2 Generic Inefficiency

The conflict externality vanishes at a Nash equilibrium only if the cross-effects happen to cancel—a condition that imposes I restrictions on the battle function Jacobians and holds generically for no $m^* \neq 0$.

5.10 Proposition (Generic global inefficiency)

Generically over the space of battle functions satisfying Assumption 2.3, no DCE with $m^* \neq 0$ is globally Pareto optimal.

Proof. At a globally Pareto optimal DCE with multipliers $\alpha_i = \lambda_i^*$ (from Corollary 5.9), conditions (5.1) and (5.2) must hold simultaneously. Subtracting, the conflict externality must vanish:

$$\sum_{j \neq i} \lambda_j^* \, Du_j(x_j^{\prime *}) \cdot D_{m_i} s_j(m^*, x^*, \omega) = 0 \qquad \text{for all } i.$$

This is a system of I scalar equations in the derivatives $D_{m_i} s_j$ ($j \neq i$) of the battle functions, evaluated at (m^*, x^*, ω). For generic smooth s satisfying Assumption 2.3, these equations are not satisfied simultaneously at any $m^* \neq 0$: the constraints lie on a set of measure zero in the space of battle function profiles. ■

The inefficiency has two structural sources. The *resource externality* arises because $c_i(m_i^*) > 0$ for $m_i^* \neq 0$ (Assumption 2.2): military spending diverts resources from exchange, shrinking the feasible consumption set in a way that

Nash logic does not internalize. The *security externality* arises because $D_{m_i} s_j \neq 0$ for $j \neq i$: state i's deployment shifts other states' battle outcomes, affecting their welfare in ways that are invisible to the Nash equilibrium.

Together, these externalities imply that DCEs generically over-militarize relative to the social optimum. The Nash equilibrium arms because arming is individually rational; but the resulting arms level is collectively wasteful. This is the international-political analogue of the prisoners' dilemma at the level of aggregate welfare.

5.3.3 The Efficiency Gap

Proposition 5.1 gives us the tools to measure the welfare cost of conflict precisely. Following Balasko (2009, Section A.5.1), define the *value function*

$$\rho(r, u_{-I}) := \max\left\{u_I(x_I) \mid (x_1, \ldots, x_I) \in \Omega(\leqslant r),\ u_i(x_i) \geqslant u_{-I,i} \text{ for } i < I\right\}$$

for $(r, u_{-I}) \in X \times \mathcal{U}_{I-1}(r)$. By Proposition A.5.1 of Balasko (2009), ρ is a smooth function; by strict monotonicity of u_I, it is strictly increasing in r for fixed u_{-I}. The Pareto frontier $\partial\overline{\mathcal{U}}(r)$ is the graph of $\rho(\cdot, r)$ over $\mathcal{U}_{I-1}(r)$.

At any DCE, let $u'^*_{-I} := (u_1(x'^*_1), \ldots, u_{I-1}(x'^*_{I-1}))$ be the utility levels of the first $I-1$ states. Since $x'^* \in \mathcal{P}(r^*)$, we have $u'^*_{-I} \in \mathcal{U}_{I-1}(r^*)$ and state I achieves $u^*_I = \rho(r^*, u'^*_{-I})$. The hypothetical utility available for state I at the full Pareto frontier—holding other states' utilities fixed and supplying the full endowment ω—is $\rho(\omega, u'^*_{-I})$. Define the *efficiency gap*

$$\Delta^* := \rho(\omega, u'^*_{-I}) - \rho(r^*, u'^*_{-I}). \tag{5.3}$$

5.11 Proposition (Efficiency gap)

Under zero-sum warfare:

1. $\Delta^* > 0$ *whenever* $m^* \neq 0$.
2. $\Delta^*(\omega)$ *is a smooth function of* ω *at regular equilibria.*
3. $\Delta^* \to 0$ *as* $\sum_i \|c_i(m^*_i, p^*, \omega_i)\| \to 0$.

Proof. Under zero-sum warfare, $r^* = \omega - \sum_i c_i(m^*)$. When $m^* \neq 0$, Assumption 2.2 gives $\sum_i c_i(m^*) > 0$, so $r^* < \omega$ componentwise in at least one good. Since $\rho(r, u_{-I})$ is strictly increasing in r (more resources allow a higher payoff for state I holding others fixed), $\rho(\omega, u'^*_{-I}) > \rho(r^*, u'^*_{-I})$, giving part (1). Part (2) follows from smoothness of ρ (Balasko Proposition A.5.1) and smoothness of the equilibrium selection $r^*(\omega) = \omega - \sum_i c_i(m^*(\omega))$ at regular equilibria (Chapter 4). Part (3) is continuity of ρ in r. ∎

The efficiency gap Δ^* measures, in utility units for state I, how much welfare is lost to conflict. It is zero only at peaceful equilibria, positive wherever there is militarization, and increasing in total military expenditure. Part (2) connects the normative and positive theories: the welfare cost of conflict is a smooth function of the resource distribution, so comparative statics on ω (Chapter 4) directly yield comparative statics on Δ^*.

5.4 Peaceful Equilibria

The conflict externality identified in the preceding section is not inevitable. Under conditions on the battle technology that prevent conflict from creating aggregate resources, peaceful equilibria escape the externality entirely and achieve the global Pareto frontier.

5.4.1 Zero- and Negative-Sum Warfare

The relevant restriction is on the aggregate effect of conflict: whether the battle technology is capable of creating resources at the social level, or merely redistributes what already exists.

5.12 Definition (Zero- and negative-sum warfare)
Warfare is zero-sum *if* $\sum_i s_i^\ell(m, x, \omega) = 0$ *for all goods* ℓ*, deployment profiles* m*, allocations* x*, and endowments* ω*. Warfare is* negative-sum *if* $\sum_i s_i^\ell(m, x, \omega) \leq 0$ *for all* ℓ*,* m*,* x*, and* ω*.*

Zero-sum warfare means that any resources gained by one state through conflict are lost by others; the aggregate is unchanged. Negative-sum warfare means that conflict destroys aggregate resources as well—the weapons, lives, and infrastructure consumed in fighting are gone from both sides. Both conditions say that conflict is not a productive technology at the social level: it cannot expand the total resource pie, only reshape (zero-sum) or shrink (negative-sum) it.

The "lootable net exports" battle function of Example 2.4 is zero-sum in the two-state case: $s_A^\ell + s_B^\ell = \lambda^\ell(\omega_B^\ell - x_B^\ell)(m_A - m_B)/(1 + m_A + m_B) + \lambda^\ell(\omega_A^\ell - x_A^\ell)(m_B - m_A)/(1 + m_A + m_B)$, which vanishes if $(\omega_A^\ell - x_A^\ell) + (\omega_B^\ell - x_B^\ell) = 0$ (net exports sum to zero) or if one interprets the formula as the state capturing the other's net exports (what one gains, the other loses). More generally, any battle function whose spoils represent pure redistribution—what one state seizes, another loses—satisfies the zero-sum condition.

5.4.2 Conditional Efficiency of Peaceful DCE on the Full Frontier

A *peaceful DCE* has $m^* = 0$. By Assumption 2.2, $c_i(0, p^*, \omega_i) = 0$; and the normalization $s_i(0, x^*, \omega) = 0$ (zero deployment yields zero transfers) gives $x_i'^* = x_i^*$ and $r^* = \omega$. Proposition 5.8 then immediately implies:

5.13 Corollary (Peaceful DCE on the full Pareto frontier)
At any peaceful DCE $(p^*, x^*, 0)$, the equilibrium allocation satisfies $x^* \in \mathcal{P}(\omega)$.

The market mechanism, operating on the undistorted endowment profile, places the equilibrium on the boundary of the maximal utility possibility set $\partial\overline{\mathcal{U}}(\omega)$. This is the full Pareto frontier—the highest utility levels simultaneously achievable across all states given the economy's resources.

5.4.3 Global Pareto Optimality

Corollary 5.13 establishes conditional efficiency at the full resource level. The key question is whether a peaceful DCE is also superior to all *non-peaceful* alternatives—whether changing the deployment profile could improve outcomes. Under zero- or negative-sum warfare, the answer is no.

5.14 Proposition (Peaceful efficiency)
Suppose warfare is zero- or negative-sum. Then any peaceful DCE $(p^*, x^*, 0)$ is globally Pareto optimal.

Proof. By Corollary 5.13, $x^* \in \mathcal{P}(\omega)$, so $u(x^*) \in \partial\overline{\mathcal{U}}(\omega)$ by Proposition 5.1(2).

Suppose for contradiction that (y, m') is globally feasible and $u_i(y_i') \geqslant u_i(x_i^*)$ for all i, with strict inequality for some j, where $y_i' = y_i + s_i(m', y, \omega)$.

Case 1: $m' = 0$. Then $y_i' = y_i$ and $y \in CF(0) = \{z \in X : \sum_i z_i \leqslant \omega\}$. But $x^* \in \mathcal{P}(\omega)$ means no reallocation in $CF(0)$ Pareto dominates x^* (conditional efficiency at $m^* = 0$), contradiction.

Case 2: $m' \neq 0$. Global feasibility gives $\sum_i y_i \leqslant \omega - \sum_i c_i(m_i')$, so

$$\begin{aligned}\sum_i y_i' &= \sum_i y_i + \sum_i s_i(m', y, \omega) \\ &\leqslant (\omega - \sum_i c_i(m_i')) + \sum_i s_i(m', y, \omega).\end{aligned}$$

Under zero-sum warfare ($\sum_i s_i = 0$): $\sum_i y_i' \leqslant \omega - \sum_i c_i(m')$. Under negative-sum warfare ($\sum_i s_i \leqslant 0$): $\sum_i y_i' \leqslant \omega - \sum_i c_i(m')$ as well. Since $m' \neq 0$ implies $\sum_i c_i(m_i') > 0$ in at least one component (Assumption 2.2), we have $\sum_i y_i' < \omega$ componentwise in at least one good. Therefore

$$u(y') \in \overline{\mathcal{U}}\,(\omega - \textstyle\sum_i c_i(m')) \subset \mathcal{U}(\omega),$$ ■

where the inclusion is Lemma 5.6.

Now $u(x^*) \in \partial\overline{\mathcal{U}}(\omega)$ and $u(y') \in \mathcal{U}(\omega) = \text{int}(\overline{\mathcal{U}}(\omega))$. Suppose $u(y') \geqslant u(x^*)$ componentwise. Then $u(y') \in (u(x^*) + \mathbb{R}^I_+) \cap \overline{\mathcal{U}}(\omega)$, so the boundary condition (Proposition 5.1(3)) gives $u(y') = u(x^*)$. But $u(y')$ is in the interior and $u(x^*)$ is on the boundary—two disjoint sets—a contradiction.

In both cases, no globally feasible outcome Pareto dominates x^*.

The proof exploits the Balasko utility possibility set in an essential way. The geometric content is:

1. x^* sits on the boundary $\partial\overline{\mathcal{U}}(\omega)$: the utility vector $u(x^*)$ is maximal; no feasible reallocation can improve any component.

2. Any alternative with $m' \neq 0$ under zero/negative-sum warfare produces total consumption in a strictly smaller utility possibility set $\overline{\mathcal{U}}(\omega - \sum c_i)$, which lies in the *interior* of $\overline{\mathcal{U}}(\omega)$.

3. Interior and boundary are disjoint: no point in the interior can weakly Pareto dominate a point on the boundary.

The zero/negative-sum condition is what makes step (2) work. Without it, an aggressive state could in principle seize enough from others to push aggregate resources beyond ω, creating new space for a Pareto improvement. Under zero/negative-sum warfare, conflict cannot expand aggregate resources, so paying the military cost $\sum_i c_i(m') > 0$ strictly shrinks the feasible utility set.

5.15 Remark (Comparison with the bargaining model)

In the bargaining model of war (Fearon, 1995), it is assumed as a premise that war is inefficient: any allocation achievable by fighting can also be achieved by agreement. Proposition 5.14 establishes the efficiency of peace as a theorem, under explicit conditions on the conflict technology. The premise of the bargaining model is the conclusion of our analysis. This not only provides a foundation for the standard assumption but also identifies exactly what it requires: that warfare be zero- or negative-sum and that a peaceful DCE exist. Neither condition is automatic, and the applications of Chapter 6 will show how each can fail.

5.5 The Failure of the Second Welfare Theorem

The First Welfare Theorem finds its standard companion in the Second: any Pareto optimal allocation can be supported as a competitive equilibrium if initial endowments are appropriately redistributed. In a standard exchange economy, this means that efficiency and distribution can be separated: the market handles allocation, and distribution is handled by lump-sum transfers. In the dual-competitive setting, this separation fails.

5.16 Proposition (Failure of the Second Welfare Theorem)

Generically over battle functions and opportunity cost functions satisfying Assumptions 2.3–2.2, a globally Pareto optimal allocation $(x^*, 0) \in \mathcal{P}(\omega) \times \{0\}$ *is not supportable as a peaceful DCE at any endowment profile* ω' *obtained by redistributing* ω.

Proof (Proof sketch). For $(x^*, 0)$ to be a peaceful DCE at ω', two conditions must hold simultaneously:

Exchange condition. There exist prices p and endowments ω' with $\sum_i \omega_i' = \omega$ such that $x_i^* = f_i(p, p \cdot \omega_i')$ for all i (each state demands x_i^* at prices p from endowment ω_i'). Setting $\omega_i' = x_i^*$ trivially satisfies this for any p that rationalizes the allocation, and such p exist by the classical Second Welfare Theorem for pure exchange economies.

Nash condition. $m_i = 0$ must be a best response for each state given $m_{-i} = 0$ and prices p. By (FOC 2), this requires:

$$Du_i(x_i^*) \cdot D_{m_i} s_i(0, x^*, \omega') \;\leq\; \mu_i^* \, p \cdot D_{m_i} c_i(0, p, \omega_i') \tag{5.4}$$

■

for all i. The left-hand side is the marginal utility of increasing m_i from zero (marginal gain from arming); the right-hand side is its marginal budgetary cost.

The exchange condition can be satisfied by choosing ω'. The Nash condition, however, depends on $D_{m_i} s_i$ and $D_{m_i} c_i$—the marginal conflict and cost derivatives at zero deployment—which are properties of the battle and cost technologies, not of the endowment distribution. In particular, endowment redistribution changes the level of consumption x_i^* and the income $p \cdot \omega_i'$, but it cannot alter the *shape* of s_i and c_i as functions of m_i. For generic (s, c), there exist allocations x^* and prices p for which the marginal gain from arming exceeds the marginal cost in (5.4) for some i, regardless of ω_i'. At such allocations, $m_i = 0$ is not a Nash best response for any wealth level, and no endowment redistribution can support $(x^*, 0)$ as a DCE.

The failure of the Second Welfare Theorem reflects a fundamental incompleteness of the price mechanism in the dual-competitive setting. Prices coordinate the exchange of goods; they do not coordinate the use of force. The Nash equilibrium over m is governed by the conflict technology—by the marginal returns to militarization, as encoded in $D_{m_i} s_i$ and $D_{m_i} c_i$—and this logic is not correctable by price signals or endowment transfers.

To support an efficient peaceful outcome as a DCE, a society would need not just the right wealth distribution but the right conflict technology: battle functions for which peace is individually optimal at the efficient allocation. This is an institutional and technological problem, not merely a distributional one. It is this observation that motivates the study of international institutions in

Chapter 8: institutions are mechanisms that alter the effective conflict technology, making peace individually rational where the underlying s and c would otherwise make conflict tempting.

5.6 Summary

The normative theory of dual-competitive equilibria organizes itself around two findings.

The first is conditional efficiency. Every DCE places total consumption on the conditional Pareto frontier $\mathcal{P}(r^*)$, where $r^* = \omega - \sum_i c_i(m^*) + \sum_i s_i(m^*, x^*, \omega)$ are the effective resources after accounting for military costs and battle transfers. Under strict concavity, the equilibrium has a unique welfare weight representation $x^{I*} = T_{I-1}(r^*, \lambda^*)$, with weights λ^* determined by the competitive distribution of income. The market mechanism resolves the exchange problem efficiently—conditional on whatever strategic environment it faces.

The second finding is global inefficiency. For any $m^* \neq 0$, the conflict externality drives a wedge between the Nash optimum and the social optimum: each state's militarization imposes resource and security externalities that prices cannot correct. The efficiency gap $\Delta^* = \rho(\omega, u_{-I}^{I*}) - \rho(r^*, u_{-I}^{I*})$ measures this cost in utility terms; it is positive, smooth in ω, and grows with total military expenditure. Peaceful equilibria under zero- or negative-sum warfare are the sole exception: by eliminating both the resource externality ($c_i = 0$) and the security externality ($s_i = 0$), they place the equilibrium on the global frontier $\mathcal{P}(\omega)$.

The Second Welfare Theorem fails in this setting because the Nash condition on military deployment is governed by conflict technology, not endowment distribution. No redistribution of initial resources can make peace individually rational if the battle function renders arming profitable. Addressing the conflict externality requires changing the strategic environment itself—the function of institutions, to which Chapter 8 turns.

Chapter 6

Applications

The foregoing chapters establish the dual-competitive equilibrium as a mathematically rigorous foundation for the study of international politics. The positive theory (Chapter 4) characterizes the equilibrium manifold and its comparative statics. The normative theory (Chapter 5) establishes conditional efficiency, the conflict externality, and the structural failure of the Second Welfare Theorem. These results are general, but they are abstract. The present chapter derives three substantive consequences of the framework that speak to central debates in international political economy.

No functional-form restrictions beyond the model's axioms are imposed. No numerical calibration is required. Each result is derived from the structure of the dual-competitive equilibrium directly, by identifying which features of the equilibrium determine a specific observable phenomenon. The three applications concern: (i) the liberal peace hypothesis; (ii) the reliability of GDP-based measures of military effort; and (iii) the informational content of current trade flows as a predictor of conflict.

6.1 Opening Up Economies Need Not Make Peace

The liberal peace hypothesis holds that trade promotes peace. States that exchange goods face high opportunity costs of conflict; the more they trade, the higher those costs; hence trade-intensive economies should be more peaceful. A large empirical literature documents a negative correlation between bilateral trade flows and conflict (Polachek, 1980; Findlay and O'Rourke, 2007). The analytic literature has incorporated trade into bargaining and contest models

(Garfinkel, Skaperdas and Syropoulos, 2015; Bó and Bó, 2011).

The present framework introduces a distinction the dyadic evidence cannot capture. In the DCE, market-clearing prices aggregate all states' endowments and preferences simultaneously. A bilateral trade relationship between states i and j cannot be studied in isolation; prices are determined globally, and a third party's entry into the economy can alter i's conflict incentive against j even without engaging either of them directly. The following result makes this precise.

6.1.1 The Peace Condition

Consider a dual-competitive economy with I states and L goods. At a corner equilibrium $(p^*, 0)$ (peace), each state solves an unconstrained Walrasian demand problem. State i has no incentive to deviate to positive deployment if and only if the marginal gain from militarization does not exceed its marginal cost. By the envelope theorem applied to state i's problem, and using the battle and cost structures of Chapter 2, this condition is

$$\sum_{\ell=1}^{L} p^{*\ell}\left(\lambda^{\ell} - \kappa_i^{\ell}\right) \leq 0, \qquad \forall\, i, \tag{6.1}$$

where $\lambda^{\ell} \geq 0$ is the lootability of good ℓ and $\kappa_i^{\ell} \geq 0$ is i's marginal cost of deploying against good ℓ. The condition is intuitive: peace holds for state i whenever the price-weighted sum of lootability-minus-cost is non-positive across all goods.

Condition (6.1) depends on the equilibrium price vector p^*, which in turn depends on the entire endowment profile ω. This is the crucial point. Even if the parameters λ^{ℓ} and κ_i^{ℓ} are fixed, changes in ω—including changes originating with states other than i and j—can move p^* and thereby shift (6.1) from satisfied to violated. The conflict between i and j is not a bilateral phenomenon; it is an equilibrium outcome of a global system.

6.1.2 A New Entrant Can Break the Peace

We now show that introducing a new state into a peaceful economy can destroy that peace. The mechanism is pure price pressure: the new state's demand for traded goods reshapes the global price vector, potentially pushing (6.1) into positive territory for some original state.

6.1 Definition (Inert entrant)
A state W is an inert entrant *to the economy* $(\{1,\ldots,I\},\{1,\ldots,L\})$ *if it is characterized by:*

1. *a large endowment of a new good $L{+}1$, which is non-lootable ($\lambda^{L+1}=0$) and not held by the original states;*
2. *strictly positive preferences over goods $1,\ldots,L$; and*
3. *opportunity costs high enough that W never attacks ($\kappa_W^\ell \geq \lambda^\ell$ for all $\ell \leq L$).*

The extended economy has $I{+}1$ states and $L{+}1$ goods.

An inert entrant is, by construction, peaceful: it never militarizes. Its only effect is through trade. Because it holds a large endowment of good $L{+}1$ that the original states desire, and because it competes with them for goods $1,\ldots,L$, its entry changes the equilibrium price vector.

6.2 Proposition (Systemic conflict from trade openness)
Let $\omega_c \in \Omega_0$ be an endowment profile at which the I-state, L-good economy has a peaceful DCE. Suppose that for some state i there exists a good ℓ^* with $\kappa_i^{\ell^*} < \lambda^{\ell^*}$ (good ℓ^* is profitably lootable for i at sufficiently high prices). Then there exists an inert entrant W such that in the extended economy, state i strictly prefers positive deployment.

Proof. Let W have preferences $\alpha_W^{\ell^*} > 0$ over good ℓ^* and endowment $\omega_W^{L+1} = R$ of good $L{+}1$. As $R \to \infty$, W's wealth $p^{L+1}R \to \infty$, and its demand for good ℓ^* grows without bound. By the market-clearing condition and the desirability property (Lemma 2.15), the equilibrium price $p^{*\ell^*}$ must rise without bound as R increases: the original states' aggregate supply of good ℓ^* is bounded, so unbounded demand forces the price up.

Formally, market clearing for good ℓ^* requires $p^{*\ell^*} \to \infty$ as $R \to \infty$. The peace condition (6.1) for state i involves the term $p^{*\ell^*}(\lambda^{\ell^*} - \kappa_i^{\ell^*})$, which is positive by assumption and grows without bound. For R large enough, this term dominates all others in (6.1), making the sum strictly positive and the peace condition violated. State i then strictly prefers to deploy. ∎

6.3 Remark (Dyadic measurement misses systemic causation)
Proposition 6.2 does not require W to trade with or interact directly with either i or any potential adversary. W merely participates in the same market. The empirical correlation between bilateral trade flows and conflict, estimated at the dyadic level, cannot detect this channel: dyadic flows between i and j may be unchanged by W's entry, while the conflict incentive between them changes entirely because the global price of a contested good has moved.

6.4 Remark (Peaceful dyads in a violent world)
The condition $\kappa_i^{\ell^*} < \lambda^{\ell^*}$ is compatible with the small-complementarities assumption (Assumption 2.3), which only requires $\kappa_i^{\ell^*} > 2\lambda^{\ell^*}/27$. For $\kappa_i^{\ell^*} \in (2\lambda^{\ell^*}/27, \lambda^{\ell^*})$, the economy is conditionally peaceful (the original equilibrium satisfies (6.1)) but vulnerable to systemic conflict induction. The original states may never have contested anything before W's arrival, yet find themselves at war once W bids up the price of a lootable good.

6.2 A Pathology in a Standard Measure of Militarization

Military deployments are not directly observable as a single scalar. Empirical researchers aggregate multi-commodity resource commitments into a scalar measure of military effort. The most widely used is the fraction of gross domestic product devoted to military expenditure, compiled by the Stockholm International Peace Research Institute (Tian et al., 2018). The present framework allows us to examine whether this statistic is a reliable proxy for the underlying deployment in a general-equilibrium setting. It is not.

6.2.1 The GDP-Share Statistic

Since supply is exogenous, state i's total market wealth is $p^* \cdot \omega_i$. The value of the resources devoted to militarization is $p^* \cdot c_i(m_i^*, p^*)$. Define the GDP-share statistic

$$\mu_i = \frac{p^* \cdot c_i(m_i^*, p^*)}{p^* \cdot \omega_i} \times 100. \tag{6.2}$$

This is dimensionless, expressed as a percentage, directly comparable across states and across time, and zero at peace. Its failure modes are structural.

By the diffeomorphism $E \cong \mathbb{R}^{IL}$ (Proposition 3.10), the equilibrium map $\omega \mapsto (p^*, m^*)$ is a smooth bijection. The map is informationally complete: the full equilibrium is uniquely recoverable from the endowment profile. The statistic μ_i is a further projection—a single real number derived from the full equilibrium. It discards the structure that the diffeomorphism preserves. In particular, μ_i is not a coordinate of the diffeomorphism; there is no reason it should be monotone in any equilibrium variable, or single-valued as a function of any component of m^*. The following two propositions establish that it is generically neither.

6.2.2 Non-Monotonicity and Non-Single-Valuedness

Two distinct failure modes arise. The first is a failure of monotonicity: as a state's opportunity cost parameter increases, its deployment falls, yet the GDP-share statistic may rise, because the equilibrium price response raises the cost-weighted value of the remaining deployment faster than the deployment reduction lowers it. The second is a failure of injectivity: a given deployment level by one state can correspond to two different values of the GDP-share statistic, because two distinct equilibrium price vectors can support the same Nash behavior with different cost valuations. Both failures arise from the general-equilibrium determination of prices, which is suppressed by the statistic but is essential to the equilibrium.

6.5 Proposition (Non-monotone μ)
In any economy with price-dependent opportunity costs ($\partial c_i^\ell / \partial p^{-\ell} \neq 0$ for some i, ℓ), there generically exist smooth parameter paths along which m_i^* is strictly decreasing while μ_i is strictly increasing.

Proof. Fix a smooth path $\kappa_i^1(t)$ increasing in t, holding all other parameters fixed. Since $\partial^2 V_i / \partial m_i \partial \kappa_i^1 < 0$ (higher cost depresses deployment), $m_i^*(t)$ is strictly decreasing. Differentiate μ_i along the path:

$$\frac{d\mu_i}{dt} = \frac{1}{p^* \cdot \omega_i} \left[\frac{\partial(p^* \cdot c_i)}{\partial \kappa_i^1} + \frac{\partial(p^* \cdot c_i)}{\partial m_i^*} \cdot \frac{dm_i^*}{dt} + D_{p^*}(p^* \cdot c_i) \cdot \frac{dp^*}{dt} \right].$$

The first term, $\partial(p^* \cdot c_i)/\partial\kappa_i^1$, is strictly positive: raising the cost parameter directly raises the cost-weighted expenditure, holding m_i^* and p^* fixed. For small κ_i^1 (near zero), the deployment m_i^* is large and this direct effect is proportionally large, while $dm_i^*/dt < 0$ remains bounded away from $-\infty$. Hence the first term dominates at small κ_i^1, giving $d\mu_i/dt > 0$ despite $dm_i^*/dt < 0$. As $\kappa_i^1 \to \infty$, $m_i^* \to 0$ and both the first and third terms vanish; the second term also vanishes since m_i^* has nothing further to lose. Hence $d\mu_i/dt \to 0^-$ for large κ_i^1. By the intermediate value theorem, $d\mu_i/dt$ changes sign on the path, establishing non-monotonicity of μ_i in m_i^*. ∎

6.6 Proposition (Non-single-valued μ)

In any economy with price-dependent opportunity costs ($\partial c_j^\ell/\partial p^{-\ell} \neq 0$ for some $j \neq i$, ℓ), there generically exist two distinct equilibria with the same deployment m_j^* but different values of μ_j.

Proof. The Nash equilibrium deployment m_j^* is uniquely determined by the Nash first-order condition at any fixed price p (Proposition 3.5). This condition pins down m_j^* as a function of p; it does not require p to be unique given m_j^*. With price-dependent costs, the Nash condition is $\partial V_j/\partial m_j = 0$, an equation that depends on p both through the budget constraint and through $c_j(m_j, p)$. The set of (m_j, p) pairs satisfying this condition is a smooth submanifold of $\mathbb{R} \times S_N$ of dimension $L - 1$; points on this submanifold with the same m_j value but different p generally exist when $L > 1$. Since $\mu_j = p \cdot c_j(m_j^*, p)/(p \cdot \omega_j)$ depends on p through the cost function (via $\partial c_j/\partial p^{-\ell} \neq 0$), two such points yield different values of μ_j generically. ∎

6.2.3 The Root Cause

Both pathologies share a common origin. The non-monotonicity of μ_i in m_i^* arises because p^* is not monotone in κ_i^1: prices equilibrate the entire market, not just the relationship between A's deployment and A's wealth. More deployment withdraws goods from the market, raising prices, which feeds back into the cost of deployment. The μ statistic registers this feedback as if it were an increase in the thing it is supposed to measure.

The non-single-valuedness of μ_j on m_j^* arises because the price space has more degrees of freedom than the scalar deployment. Two distinct equilibrium price vectors can support the same Nash deployment with different cost valuations. This is not a knife-edge phenomenon; it is generic whenever $L > 1$ and costs depend on prices.

The underlying lesson is about what the diffeomorphism $E \cong \mathbb{R}^{IL}$ does and does not guarantee. It guarantees that the space of equilibria is topologically simple and that the full equilibrium is smoothly and bijectively parameterized by endowments. It offers no guarantee about statistics derived from the full equilibrium by further aggregation. Any such statistic—μ being a canonical example—inherits the complexity of the price system it ignores.

6.3 Trade Expectations and Preventive Conflict

The previous two applications are static: the equilibrium is analyzed at a fixed endowment profile. The present application concerns how anticipated changes in the global economy affect current conflict incentives. Copeland (2015) argues that the relevant variable for conflict onset is not the current level of bilateral trade but a state's *expectation* of the future trade environment: a state that anticipates favorable future access will remain at peace; one that anticipates cutoff will fight, and will fight sooner rather than later.

The present framework formalizes and extends this claim. By the diffeomorphism $E \cong \mathbb{R}^{IL}$, any smooth path in parameter space Ω_0 induces a smooth path in the space of equilibria. A state can read off, from the smooth comparative statics of Chapter 4, the derivative of its equilibrium welfare along any anticipated parametric path. The central result is that a state whose welfare is declining along the anticipated path has a strictly higher current conflict incentive than a state at the same allocation facing a flat trajectory—even if both are currently at peace.

6.3.1 Trade Expectations

Consider a dual-competitive economy with endowment profile $\omega(t) = \omega_0 + t\Delta\omega$ for $t \in [0, T]$, where $\Delta\omega \in \Omega_0$ is the anticipated direction of parametric change. By the smooth comparative statics of section 4.5, this induces a smooth path of peaceful equilibrium utilities $\overline{V}_i(t)$ and, when an interior equilibrium exists, a smooth path of conflict utilities $V_i^*(t)$.

6.7 Definition (Trade expectation)
The trade expectation *of state i at $t = 0$ along the path $\omega(\cdot)$ is*

$$\tau_i := \left.\frac{d\overline{V}_i}{dt}\right|_{t=0} = D_\omega \overline{V}_i(\omega_0) \cdot \Delta\omega.$$

State i has negative *trade expectations if $\tau_i < 0$: the anticipated path delivers a declining peaceful equilibrium utility.*

The trade expectation is a directional derivative of the peaceful-equilibrium value function. By the smooth comparative statics formula (section 4.5), τ_i is computable from the current equilibrium Jacobian and the direction $\Delta\omega$. It does not require a dynamic model; it is a property of the static equilibrium family.

The *deviation gain* of state i at ω_0 is $G_i(\omega_0) = V_i^*(\omega_0) - \overline{V}_i(\omega_0)$, the utility gain from deviating unilaterally to positive deployment. Peace holds when $G_i(\omega_0) \leq 0$ for all i. The following result shows that negative trade expectations cause G_i to grow.

6.8 Proposition (Trade expectations and conflict incentive)
Let $\lambda^\ell > 0$ for some good ℓ, and suppose that at the interior DCE state i receives a positive contest transfer: $s_i^\ell(m^*) > 0$. Along any path with $\Delta\omega_i^\ell < 0$ (declining endowment of the contested good for state i):

$$\left.\frac{dG_i}{dt}\right|_{t=0} > 0.$$

Hence i's deviation gain is strictly increasing along the path even when $G_i(\omega_0) < 0$ (while i is currently at peace).

Proof. By definition, $dG_i/dt = dV_i^*/dt - d\overline{V}_i/dt$. The claim reduces to showing

$$\frac{\partial V_i^*}{\partial \omega_i^\ell} < \frac{\partial \overline{V}_i}{\partial \omega_i^\ell},$$

since then multiplying both sides by $\Delta\omega_i^\ell < 0$ and taking the difference gives $dG_i/dt > 0$.

Consider the regime near the conflict threshold where $m^* \to 0$ and $p^{**} \to p^*$ by continuity. In this regime, the envelope theorem applied to state i's optimization yields:

$$\frac{\partial \overline{V}_i}{\partial \omega_i^\ell} = \frac{p^{*\ell}}{p^* \cdot \omega_i},$$
$$\frac{\partial V_i^*}{\partial \omega_i^\ell} = \frac{p^{*\ell}}{p^* \cdot (\omega_i + s_i(m^*) - c_i(m^*))}.$$

At an interior Nash equilibrium, the first-order condition $\partial V_i^* / \partial m_i = 0$ implies that the marginal benefit of arming equals its marginal cost, so $p^* \cdot s_i(m^*) \geq p^* \cdot c_i(m^*)$ with strict inequality for small positive m^*. Hence the denominator of $\partial V_i^* / \partial \omega_i^\ell$ is strictly greater than that of $\partial \overline{V}_i / \partial \omega_i^\ell$, giving the desired inequality.■

6.3.2 The Mechanism

The logic of Proposition 6.8 is transparent. Under peace, state i has no alternative source of good ℓ: any reduction in its endowment translates fully into reduced welfare, with no buffer. Under conflict, i's access to good ℓ includes the contest transfer $s_i^\ell > 0$. A reduction in the traded endowment therefore hurts i less in the conflict equilibrium than in the peaceful one: the alternative source partially compensates. Consequently, the relative payoff to conflict over peace increases as the traded endowment declines, and the deviation gain G_i rises—even before peace breaks down.

6.9 Remark (Copeland's mechanism, formalized)

Proposition 6.8 restates the core claim of Copeland (2015) in the language of the model. Copeland argues that a state dependent on trade for a critical input, facing anticipated cutoff, will fight—and will fight sooner rather than later, while it is still strong enough to win. The proposition formalizes the "sooner rather than later" logic: $dG_i/dt > 0$ means the incentive to deviate is already increasing at $t = 0$, before peace has formally broken down. The correct predictor of conflict onset is $\tau_i = d\overline{V}_i/dt$, not the current level of bilateral trade.

6.10 Remark (Preventive conflict)
Suppose $G_i(\omega_0) < 0$ (peace at $t = 0$) but $G_i(\omega(T)) > 0$ (conflict at some future T, which follows by continuity from $dG_i/dt > 0$ and a sufficiently large horizon). State i then has a preventive *incentive: it prefers to arm at $t = 0$, exploiting its current endowment strength, rather than waiting until $t = T$, when it will fight from a depleted position. The static model does not represent this dynamic preference directly, but the comparative statics make the structure clear: the payoff to early conflict exceeds the payoff to late conflict whenever $\Delta\omega_i^\ell < 0$ and $\lambda^\ell > 0$.*

6.11 Remark (Trade flows as an insufficient statistic)
At $t = 0$, bilateral trade between i and its trading partners may be high: i is currently exchanging good ℓ at the prevailing market price, and the dyadic flows are substantial. A model conditioned on current flows alone—the standard operationalization of the liberal peace hypothesis—would predict low conflict probability. The trade expectation τ_i, by contrast, is strictly negative, and conflict is imminent. Like the μ statistic of the preceding section, current trade flows are a projection of the full equilibrium onto an observable scalar. The diffeomorphism $E \cong \mathbb{R}^{IL}$ implies that the full equilibrium is informationally complete; any scalar derived from it discards information. The relevant information discarded by trade flows is the trajectory: $\Delta\omega$ and the directional derivative τ_i. The discarded information is precisely what determines conflict onset.

6.4 Summary

The three applications share a common structure: a substantive question in international political economy is shown to require the full equilibrium machinery of the DCE, and a specific piece of that machinery delivers the answer.

The first application shows that the liberal peace hypothesis fails systemically. The peace condition (6.1) depends on the global price vector, not on bilateral trade flows. A new entrant that never attacks can destroy a peaceful bilateral equilibrium by bidding up the price of a contested good, and this effect is invisible to dyadic empirical analysis.

The second application shows that the GDP-share statistic is an unreliable proxy for deployment in a general-equilibrium setting. Its failures— non-monotonicity in the underlying deployment and non-single-valuedness across distinct equilibrium configurations—are both consequences of projecting the high-dimensional equilibrium onto a scalar. The diffeomorphism $E \cong \mathbb{R}^{IL}$ guarantees that the full equilibrium is well-behaved; it offers no guarantee about derived statistics that suppress the price dimension.

The third application shows that current trade flows are an insufficient statistic for conflict onset. The trade expectation τ_i—the directional derivative of i's equilibrium utility along the anticipated parameter path—is the correct predictor. A state with negative trade expectations has a strictly increasing deviation gain, and the preventive conflict logic follows directly: fight now, from strength, rather than later, from weakness. This is the Copeland trade-expectations mechanism, derived from first principles within the DCE framework.

Chapter 7

Scaling the Framework

The analytical chapters establish properties of the dual-competitive equilibrium that hold under the model's axioms alone: existence, smoothness, conditional efficiency, the conflict externality, and the trade expectations amplification. These are general results, valid for any configuration of states, goods, and preferences. The present chapter examines what those general properties look like in specific instances, and then extends the framework to incorporate time, uncertainty, and discounting.

Section 7.1 populates the model numerically and renders visible three phenomena that the analytical results imply but do not depict: the non-monotone relationship between endowment inequality and equilibrium militarization, the clustering of conflict around resource-rich states in a multi-state economy, and the rate at which the efficiency gap grows as the international system expands. Section 7.2 embeds the static DCE in an infinite-horizon stochastic environment. The trade expectations analysis of section 6.3 computed the derivative of the peaceful value function along an anticipated endowment path—a property of the static equilibrium family. The dynamic extension internalizes the temporal structure, derives the recursive analogue of that derivative, and shows that endowment persistence amplifies the conflict incentive by a factor that depends on the discount rate and the autocorrelation of the endowment process.

7.1 Numerical Equilibrium

The diffeomorphism $E \cong \mathbb{R}^{IL}$ (Corollary 3.12) makes numerical computation tractable in a precise sense: there is no multiplicity to navigate, the natural projection $\pi : E \to \Omega_0$ is a bijection, and the smooth Jacobian of the reduced excess demand ensures that Newton-type methods converge geometrically near the unique equilibrium. For any endowment profile ω, there is exactly one DCE, and it varies smoothly with ω. These properties mean that numerical illustrations are not merely examples; they are representative cross-sections of a smooth, globally well-behaved object.

The economies in this section share a common structure: two goods, a lootable numéraire ($\ell = 1$, $\lambda^1 > 0$, $p^{*2} = 1$) and a non-lootable traded good ($\ell = 2$, $\lambda^2 = 0$). Preferences satisfy Assumption 2.1 and are varied across illustrations. The structural parameters λ^1 and κ_i^1 are held in the range consistent with Assumption 2.3.

7.1.1 Computing the Dual-Competitive Equilibrium

The DCE is characterized by two nested conditions: market clearing at prices p^* and Nash optimality of the deployment profile m^*. These define a joint fixed-point problem in the pair (p, m). The alternating algorithm below exploits the separability of the two conditions.

1. *Initialization.* Set $m^{(0)} = 0$ (begin at peace) and choose a convergence tolerance $\varepsilon > 0$.

2. *Price step.* Given the current deployment profile $m^{(k)}$, find the Walrasian equilibrium price $p^{(k)}$ that clears markets with deployments fixed at $m^{(k)}$. Market clearing is a system of $L - 1$ equations in $L - 1$ unknowns; since the reduced excess demand $\widehat{d}(p, \omega)$ is smooth and the Jacobian J^{ex} is invertible at regular equilibria (Definition 4.2), Newton's method converges quadratically.

3. *Nash step.* Given $p^{(k)}$, compute each state i's best-response deployment $m_i^{(k+1)}$ by solving the first-order condition $\partial V_i^* / \partial m_i = 0$ at fixed $(p^{(k)}, m_{-i}^{(k)})$. Strict quasiconcavity of V_i^* in m_i (Assumption 2.2) ensures a unique solution.

4. *Convergence check.* If $\|m^{(k+1)} - m^{(k)}\| + \|p^{(k+1)} - p^{(k)}\| < \varepsilon$, stop and return $(p^{(k+1)}, m^{(k+1)})$. Otherwise, increment k and return to step 2.

Since $E \cong \mathbb{R}^{IL}$, the joint fixed point (p^*, m^*) is unique. Convergence of the alternating algorithm follows from the contraction properties of the Nash best-response map (Lemma 3.3) and the smoothness of the Walrasian price function. In practice, the algorithm converges in a small number of iterations for all configurations reported below.

7.1.2 Non-Monotone Deployment and Comparative Statics

A natural conjecture about equilibrium militarization is that it should be monotone in a state's lootable wealth: richer targets are more valuable, they arm more heavily, and their neighbors respond in kind. The equilibrium is subtler. Figure 7.1 traces the equilibrium deployments of two states, A and B, as state A's lootable endowment ω_A^1 increases from zero while B's endowment and all structural parameters are held fixed.

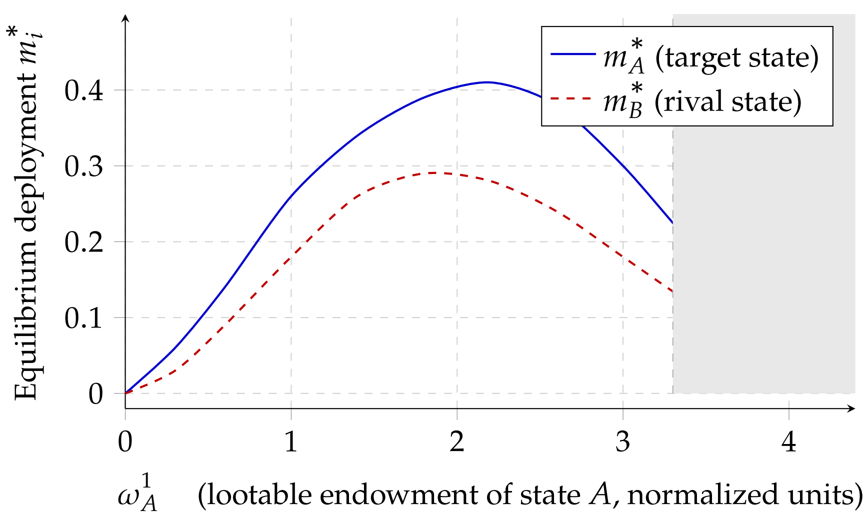

Figure 7.1: *Equilibrium deployments in a two-state, two-good economy as state A's lootable endowment ω_A^1 increases (state B's endowment and all structural parameters fixed). Both deployment schedules are hump-shaped: militarization rises with A's wealth, peaks at intermediate wealth, and declines toward zero as A approaches the inert-entrant zone of Definition 6.1 (shaded region).*

Both deployment schedules are non-monotone in ω_A^1. Three phases are visible.

Phase I (low ω_A^1): A holds little lootable wealth. The spoils from attacking A are small, and neither state deploys significantly. The peace condition (6.1) is comfortably satisfied for both.

Phase II (intermediate ω_A^1): A's growing lootable wealth makes it an increasingly valuable target. State B finds the peace condition increasingly strained; it begins deploying offensively, and A responds with defensive deployment. Both m_A^* and m_B^* rise. The peak of each schedule differs: B's deployment peaks earlier (when A is moderately wealthy), while A's peaks later (when A is wealthy enough to mount effective defense).

Phase III (high ω_A^1, shaded region): As A's wealth grows large, its demand in the market for the lootable good rises substantially. By the equilibrium price mechanism, p^{*1} rises. This has two effects. For state B, the rising price increases the opportunity cost of deploying (arms are drawn from the same resource pool as market purchases), eventually outweighing the gain from attacking A. For state A, the high price raises the implicit value of not deploying: peaceful trade now delivers large rents. Both forces push A toward the inert-entrant conditions of Definition 6.1, and both deployment schedules converge toward zero.

The figure also illustrates a feature that the analytical results establish but do not quantify: the width of the hump, and hence the range of endowments over which the system is armed, depends on the elasticity of the equilibrium price response. In economies where p^{*1} is highly responsive to endowment changes (elastic markets), the transition from conflict to inert peace is rapid; in economies with inelastic prices (no close substitutes for the lootable good), the transition is drawn out. The shape of the price response to ω_A^1 is itself a smooth function of ω—a consequence of Theorem 4.8—and can be computed directly from the comparative statics formula of section 4.5.

7.1.3 Conflict Clustering

The bilateral view of conflict—studying a pair of states in isolation—misses a feature that is structural to the general-equilibrium model: each pair's conflict incentive is determined by the same global price vector p^* that aggregates the endowments and demands of all states. A third party that holds lootable wealth changes p^{*1} and thereby alters the conflict incentive of every other dyad in the system, even those it never directly engages. Figure 7.2 illustrates this through a six-state economy.

Panel (a) shows the symmetric baseline. With equal endowments, the

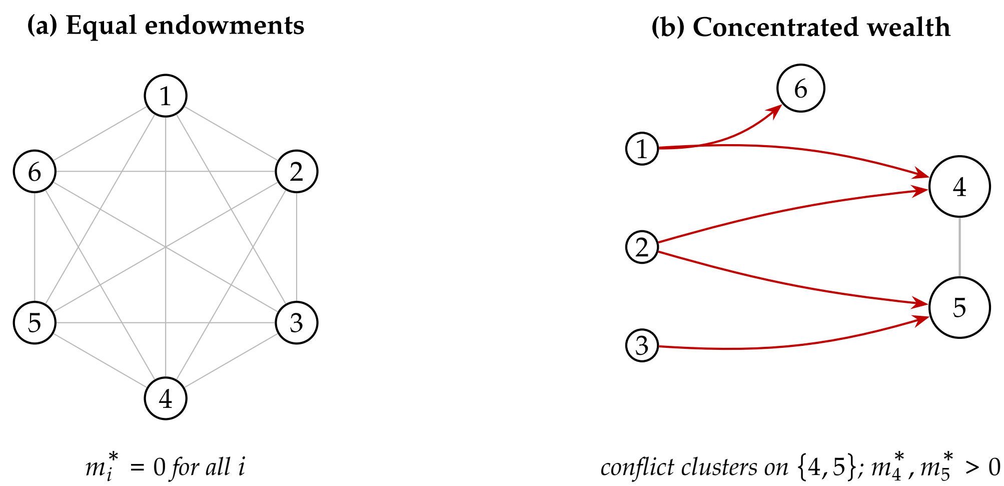

Figure 7.2: *Conflict networks in a six-state, two-good economy. Panel (a): symmetric endowments yield universal peace. Panel (b): states 4 and 5 hold concentrated lootable wealth (large circles); states 1–3 are resource-poor (small circles); state 6 is intermediate. Arrows indicate that the peace condition (6.1) is violated in that dyad. Rich-rich pairs remain at peace through mutual deterrence.*

Walrasian equilibrium price is also symmetric, and the peace condition (6.1) is satisfied for every state. No state finds it profitable to deviate to positive deployment.

Panel (b) introduces concentrated lootable wealth. States 4 and 5 each hold substantially larger endowments of good 1; state 6 holds a moderate surplus. As this concentration increases, three changes occur in sequence.

1. The Walrasian price p^{*1} of the lootable good rises, because all six states compete to purchase good 1 but the supply is increasingly concentrated.

2. For the poor states (1–3), the rising p^{*1} increases the prize from capturing the lootable good from a rich target, while the opportunity cost of deployment—paid in the numéraire—remains approximately constant. Eventually the peace condition flips: states 1–3 find positive deployment strictly profitable.

3. Between states 4 and 5, mutual deterrence prevails. Each holds enough resources to mount a sufficiently costly defense that attacking the other fails the peace condition. Their bilateral relationship remains peaceful,

sustained by the same equilibrium price p^{*1} that induces their weaker neighbors to arm.

The conflict network in panel (b) is *structurally* determined, not bilaterally. State 1's decision to attack state 4 is not a product of a bilateral grievance between 1 and 4 alone. It is a product of the price p^{*1} that states 5 and 6 (and their endowments) jointly determined. Removing state 5 from the economy would change p^{*1}, potentially resolving the conflict between 1 and 4—even though 5 is not party to that conflict. This is the systemic character of conflict documented in Proposition 6.2: the network structure of who fights whom is an equilibrium property of the full price system.

7.1.4 The Efficiency Gap at Scale

The efficiency gap Δ^*, defined in equation (5.3) and shown to be strictly positive at any armed equilibrium (Proposition 5.11), measures the aggregate welfare cost of militarization relative to the peaceful Pareto frontier. Figure 7.3 plots Δ^* as a fraction of total endowment value against the number of states I, holding the aggregate endowment and preference structure fixed as I increases (each new state receives an equal share of the total endowment).

Two features of the figure warrant comment. The first is monotonicity. Each additional state introduces new conflict relationships—new potential aggressors and new potential targets—and aggregate deployment $\sum_i m_i^*$ rises. More military resources are diverted from exchange, the feasible consumption set shrinks, and the efficiency gap between the peaceful frontier and the conflict equilibrium widens. This is consistent with Proposition 5.10: generically, armed equilibria are inefficient, and adding states generically adds to the aggregate inefficiency.

The second feature is concavity. Each new state holds a smaller per-state endowment share (since total endowment is fixed). The prizes in each bilateral conflict relationship are therefore smaller, and the marginal efficiency loss from the $(I + 1)$th state is less than that from the Ith. The system approaches a conflict-saturated limit in which adding further states changes Δ^* very little.

Two institutional implications follow from the concave shape. First, disarmament agreements generate the largest per-agreement efficiency gains in small systems, where each bilateral relationship involves high stakes. The marginal gain from transitioning from $I = 2$ to $I = 1$ (*i.e.*, eliminating the only conflict dyad) is the largest single gain available. Second, the tipping structure implies

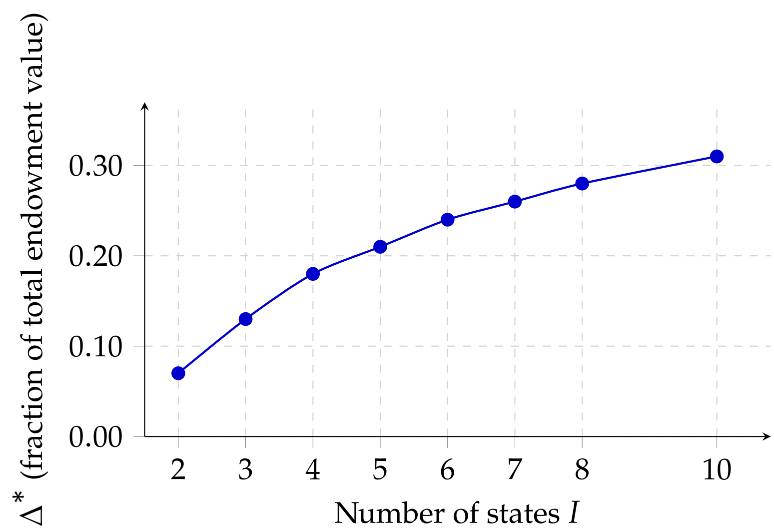

Figure 7.3: *Efficiency gap Δ^* (as a fraction of total endowment value) plotted against the number of states I in a symmetric economy (equal endowments, common preferences, two-good structure). The gap is strictly increasing and concave: each additional state contributes new conflict relationships, but with diminishing marginal impact as per-state endowment shares fall.*

that late-stage disarmament steps in large systems are cheap: in a ten-state system, removing the last conflict dyad yields only a modest efficiency improvement. A disarmament strategy that targets small high-stakes systems will outperform one that targets large low-stakes systems, holding the number of dyads addressed fixed.

7.2 A Dynamic Stochastic Extension

The static DCE takes the endowment profile ω as a given parameter and derives the equilibrium (p^*, m^*) for that parameter. This is appropriate for answering comparative statics questions—how does the equilibrium change when ω changes?—but it treats time implicitly. A state that anticipates future endowment changes faces a richer decision problem than one that takes the current endowment as permanent: it must weigh current military spending against future trade access, accounting for how much it discounts the future and how persistent the endowment process is.

This section develops a dynamic stochastic extension of the model by embedding the static DCE as the stage game of an infinite-horizon economy. The extension involves three new ingredients: an endowment process that evolves stochastically over time, a discount factor that makes future utility less valuable than current utility, and a recursive equilibrium concept that characterizes optimal behavior as a stationary function of the current state. All three are standard in the macroeconomic dynamics literature, and this section develops each from first principles before connecting the resulting structure to the results of Chapter 6.

7.2.1 The Dynamic Environment

Time and endowments. Time is discrete and infinite: $t = 0, 1, 2, \ldots$. At each period t, the world is described by an endowment profile $\omega_t = (\omega_{1,t}, \ldots, \omega_{I,t}) \in \Omega_0 := \mathbb{R}^{IL}_{++}$. The profile evolves according to a *Markov process*: ω_{t+1} is drawn from a conditional distribution $\Pi(\omega_t, \cdot)$ that depends on the current profile ω_t but not on the history prior to t. The function $\Pi : \Omega_0 \times \mathcal{B}(\Omega_0) \to [0, 1]$ is the *transition kernel*: $\Pi(\omega, A)$ is the probability that next period's endowment profile falls in the set A, given that this period's profile is ω.

Timing within a period. At each t:

1. The endowment profile ω_t is realized and observed by all states.
2. States simultaneously choose deployment levels $m_{i,t} \in M_i$.
3. Markets clear at prices p_t^*, and consumption $x_{i,t}$ is determined.
4. Period utilities $u_i(x_{i,t})$ are realized.

Steps 2–4 are exactly the static DCE of Chapter 3. The dynamic extension wraps this stage game in an infinite-horizon structure.

Preferences. Each state maximizes the expected discounted sum of period utilities:

$$\mathbb{E}_0\left[\sum_{t=0}^{\infty} \beta^t \, u_i(x_{i,t})\right],$$

where $\beta \in (0,1)$ is a common discount factor and $\mathbb{E}_0$ denotes expectation conditional on information available at $t = 0$. The discount factor captures the weight states place on the future relative to the present. When β is close to 1, states are patient: they weigh future outcomes nearly as highly as current ones. When β is close to 0, states are myopic: they care mainly about the present.

Stationarity. An important consequence of the Markov structure is that a state's optimal deployment at time t depends only on the current endowment profile ω_t, not on the entire history $(\omega_0, \ldots, \omega_{t-1})$. If the transition kernel Π does not vary with t, the decision problem is *stationary*: the same profile ω today will induce the same deployment decision regardless of how the system arrived at ω. This allows us to characterize the equilibrium as a collection of time-invariant functions of ω, called *policy functions*, rather than as an indexed sequence of period-by-period plans.

7.2.2 Recursive Dual-Competitive Equilibrium

A *recursive dual-competitive equilibrium* (RDCE) is an equilibrium of the infinite-horizon economy in which every state's decisions depend only on the current endowment profile. Concretely, it consists of two types of objects.

Policy functions describe how the economy behaves at any endowment profile ω: a price function $p^* : \Omega_0 \to S_N$, deployment functions $m_i^* : \Omega_0 \to M_i$ for each state i, and consumption functions $x_i^* : \Omega_0 \to X_i$ for each state i.

Value functions describe how much each state expects to receive, from the current period onward, when the endowment profile is ω: a DCE value $\mathcal{V}_i^* : \Omega_0 \to \mathbb{R}$ and a peaceful value $\overline{\mathcal{V}}_i : \Omega_0 \to \mathbb{R}$ for each state i.

7.1 Definition (Recursive dual-competitive equilibrium)
A recursive dual-competitive equilibrium *consists of policy functions* $(p^*, m_i^*, x_i^*)_{i=1}^I$ *and value functions* $(\mathcal{V}_i^*, \overline{\mathcal{V}}_i)_{i=1}^I$ *such that:*

1. (Static consistency) *For each* $\omega \in \Omega_0$*, the tuple* $(p^*(\omega), (m_i^*(\omega))_i, (x_i^*(\omega))_i)$ *is a dual-competitive equilibrium for the static economy with endowment profile* ω *in the sense of Definition 3.2.*

2. (Bellman equation) *For each state* i *and each* $\omega \in \Omega_0$*:*

$$\mathcal{V}_i^*(\omega) = u_i\left(x_i^*(\omega)\right) + \beta \int_{\Omega_0} \mathcal{V}_i^*(\omega')\,\Pi(\omega, d\omega'). \tag{7.1}$$

An analogous equation holds for $\overline{\mathcal{V}}_i$*, replacing the DCE stage payoff* $u_i(x_i^*(\omega))$ *with the peaceful DCE stage payoff* $u_i(\overline{x}_i^*(\omega))$*.*

The Bellman equation (7.1) has an interpretive structure worth dwelling on. The left side, $\mathcal{V}_i^*(\omega)$, is the total value that state i expects to receive from today onward when the endowment profile is ω: this is a number that captures the payoff from today's equilibrium *and* every future period, properly discounted. The right side decomposes this total into two parts. The first term, $u_i(x_i^*(\omega))$, is today's utility from the DCE allocation. The second term, $\beta \int \mathcal{V}_i^*(\omega')\Pi(\omega, d\omega')$, is the discounted expected continuation value: the probability-weighted average of the value function next period, multiplied by the discount factor β to bring it to today's units. The Bellman equation is therefore a *consistency condition*: the total value today must equal the current payoff plus the discounted expected total value tomorrow.

The elegant consequence of this structure is that an equilibrium of the infinite-horizon game is characterized by two functions $(\mathcal{V}_i^*, \overline{\mathcal{V}}_i)$ satisfying (7.1), rather than by an infinite sequence of decision rules indexed by time. Finding the RDCE reduces to finding the fixed points of two Bellman operators—an operation that is both analytically tractable and numerically implementable.

7.2.3 Existence and Uniqueness

The key tool for establishing the RDCE's existence is the Banach fixed-point theorem: an operator that is a *contraction*—that brings any two distinct inputs closer together by a fixed factor—has exactly one fixed point. We show that the Bellman operator is a contraction by verifying Blackwell's sufficient conditions.

7.2 Assumption (Endowment process)

The transition kernel Π *maps* $\Omega_0 \times \mathcal{B}(\Omega_0)$ *to* $[0,1]$ *and there exists a compact set* $K \subset \Omega_0$ *with* $\Pi(\omega, K) = 1$ *for all* $\omega \in K$ *(forward invariance). The map* $\omega \mapsto \int_K V(\omega')\,\Pi(\omega, d\omega')$ *is continuous for every bounded continuous* $V : K \to \mathbb{R}$*.*

Forward invariance of K ensures that starting inside K, the endowment process never leaves K. This is satisfied by processes that are ergodic on a bounded region of endowment space: random-walk-with-drift processes bounded away from zero and from above, or linear autoregressive processes with stationary invariant distributions.

7.3 Proposition (Existence and uniqueness of the RDCE)

Under Assumptions 2.1, 2.2, 2.3, and 7.2, for any $\beta \in (0,1)$*, there exists a unique RDCE. The value functions* $\mathcal{V}_i^*$ *and* $\overline{\mathcal{V}}_i$ *are bounded and continuous on* K*.*

Proof. Let $\mathcal{C}_b(K)$ denote the Banach space of bounded continuous functions from K to $\mathbb{R}$ under the supremum norm $\|\cdot\|_\infty$. By the static existence theorem (Theorem 3.14) and the smoothness of the DCE (Proposition 3.10), the stage payoff $\phi_i(\omega) := u_i(x_i^*(\omega))$ is bounded and continuous on K. Define the *Bellman operator* $T_i : \mathcal{C}_b(K) \to \mathcal{C}_b(K)$ by

$$(T_i V)(\omega) := \phi_i(\omega) + \beta \int_K V(\omega')\,\Pi(\omega, d\omega').$$

Continuity of $T_i V$ follows from Assumption 7.2. We verify Blackwell's two sufficient conditions for a contraction.

1. *(Monotonicity)* If $V \le W$ pointwise on K, then $(T_i V)(\omega) = \phi_i(\omega) + \beta \int V\,d\Pi \le \phi_i(\omega) + \beta \int W\,d\Pi = (T_i W)(\omega)$ for all ω.
2. *(Discounting)* For any constant $c > 0$, $T_i(V + c)(\omega) = \phi_i(\omega) + \beta \int (V + c)\,d\Pi = (T_i V)(\omega) + \beta c$, so $T_i(V + c) = T_i V + \beta c$ with $\beta < 1$.

By Blackwell's theorem, T_i is a contraction on $\mathcal{C}_b(K)$ with modulus β. By the Banach fixed-point theorem, T_i has a unique fixed point $\mathcal{V}_i^* \in \mathcal{C}_b(K)$. The same argument applies to $\overline{\mathcal{V}}_i$ with $\phi_i(\omega)$ replaced by $u_i(\overline{x}_i^*(\omega))$. The static consistency condition holds by construction. ∎

The fixed-point iteration $V^{(k+1)} = T_i V^{(k)}$ converges geometrically to $\mathcal{V}_i^*$ from any starting point $V^{(0)} \in \mathcal{C}_b(K)$, with $\|V^{(k)} - \mathcal{V}_i^*\|_\infty \le \beta^k \|V^{(0)} - \mathcal{V}_i^*\|_\infty$. The convergence rate is β: patient states (β near 1) require more iterations to converge; impatient states (β near 0) converge quickly.

7.2.4 Smooth Decision Rules

The static DCE is smooth: the map $\omega \mapsto (p^*(\omega), m^*(\omega), x^*(\omega))$ is C^∞ on Ω_0 (Corollaries 3.12 and 4.7). This smoothness of the stage payoff propagates to the dynamic value functions under a corresponding smoothness condition on the transition kernel.

7.4 Assumption (Smooth transitions)

The transition is generated by $\omega' = f(\omega, \varepsilon)$, where $f : \Omega_0 \times \mathcal{E} \to K$ is C^k in ω for μ-almost every ε, and ε is drawn from a fixed probability space $(\mathcal{E}, \mathcal{F}, \mu)$ independent of the current state.

A canonical example is the linear autoregressive process: $\omega_{t+1} = \overline{\omega} + \rho(\omega_t - \overline{\omega}) + \varepsilon_{t+1}$, with $\rho \in [0, 1)$ and ε_t drawn independently from a distribution supported on a compact set. Here $f(\omega, \varepsilon) = \overline{\omega} + \rho(\omega - \overline{\omega}) + \varepsilon$ is linear (and hence C^∞) in ω.

7.5 Proposition (Smooth RDCE value functions)

Under Assumptions 2.1, 2.2, 2.3, 7.2, and 7.4, the RDCE value functions $\mathcal{V}_i^*$ and $\overline{\mathcal{V}}_i$ are C^k on K.

Proof. Let $\mathcal{C}_b^k(K)$ denote the space of bounded C^k functions on K under the C^k norm. The stage payoff $\phi_i = u_i \circ x_i^*$ is C^k as a composition of smooth functions (Corollary 3.12). For any $V \in \mathcal{C}_b^k(K)$, the integral

$$\omega \mapsto \int_{\mathcal{E}} V\left(f(\omega, \varepsilon)\right) \mu(d\varepsilon)$$

is C^k by Leibniz's rule for differentiating under the integral sign, applicable because $f(\cdot, \varepsilon)$ is C^k uniformly in ε on the compact support of μ. Hence T_i maps $\mathcal{C}_b^k(K)$ into itself. Since $\mathcal{C}_b^k(K)$ is complete and T_i is a contraction with the same modulus β, the unique fixed point $\mathcal{V}_i^*$ belongs to $\mathcal{C}_b^k(K)$. ■

The smooth RDCE value functions imply that the policy functions $m_i^*(\omega)$ and $p^*(\omega)$ are also smooth functions of the endowment profile. States' deployment decisions, and the prices that clear markets, vary continuously and differentiably with the state of the world. There are no jumps, kinks, or non-differentiabilities in the RDCE policy functions. This is the dynamic analogue of the smooth comparative statics established in Theorem 4.8.

7.2.5 Dynamic Trade Expectations

The trade expectations result of section 6.3 computed the rate of change of state i's peaceful value function along an anticipated endowment path. That analysis was conducted in the static model: the "path" was a family of static economies parameterized by $t \in [0, T]$, and the trade expectation $\tau_i = d\overline{V}_i/dt|_{t=0}$ was a directional derivative of the *stage* peaceful utility. In the RDCE, the peaceful value function $\overline{\mathcal{V}}_i$ already accounts for all future periods, and the analogue of τ_i is the derivative of $\overline{\mathcal{V}}_i$ with respect to the current endowment profile.

7.6 Definition (Dynamic trade expectation)

The dynamic trade expectation *of state i at endowment profile ω along direction $\Delta\omega \in \mathbb{R}^{IL}$ is*

$$\tau_i^{\text{dyn}}(\omega) := D_\omega \overline{\mathcal{V}}_i(\omega) \cdot \Delta\omega.$$

State i has negative dynamic trade expectations if $\tau_i^{\text{dyn}}(\omega) < 0$.

The dynamic trade expectation is the gradient of the total peaceful value function dotted with the anticipated direction of change. Unlike the static τ_i, it accounts for the *entire discounted future path* of peaceful utility implied by the endowment process, not just the one-period derivative. The following proposition quantifies how τ_i^{dyn} relates to the static τ_i under the autoregressive process.

7.7 Proposition (Amplification of trade expectations)
Suppose the endowment process is autoregressive: $\omega_{t+1} = \overline{\omega} + \rho(\omega_t - \overline{\omega}) + \varepsilon_{t+1}$ *with* $\rho \in [0, 1)$ *and* ε_t *drawn independently with mean zero. Evaluated at the stationary mean* $\overline{\omega}$*, the dynamic and static trade expectations satisfy*

$$\tau_i^{\text{dyn}}(\overline{\omega}) = \frac{\tau_i(\overline{\omega})}{1 - \beta\rho}, \tag{7.2}$$

where $\tau_i(\overline{\omega}) = D_\omega \overline{V}_i(\overline{\omega}) \cdot \Delta\omega$ *is the static trade expectation evaluated at* $\overline{\omega}$*. Since* $\beta\rho > 0$*, the dynamic trade expectation is strictly larger in magnitude than the static one.*

Proof. Differentiate the peaceful Bellman equation with respect to ω at the stationary mean $\overline{\omega}$. For the autoregressive process, $\omega' = \overline{\omega} + \rho(\omega - \overline{\omega}) + \varepsilon$, so $D_\omega \omega' = \rho I_{IL}$ (the derivative of next period's endowment with respect to today's is ρ times the identity). Applying the chain rule to the integral:

$$D_\omega \overline{\mathcal{V}}_i(\omega) = D_\omega u_i\left(\overline{x}_i^*(\omega)\right) + \beta \int D_{\omega'} \overline{\mathcal{V}}_i(\omega') \cdot D_\omega \omega' \, \mu(d\varepsilon).$$

At the stationary mean $\overline{\omega}$, stationarity implies $\mathbb{E}[D_{\omega'} \overline{\mathcal{V}}_i(\omega')] = D_\omega \overline{\mathcal{V}}_i(\overline{\omega})$. Substituting $D_\omega \omega' = \rho I_{IL}$:

$$D_\omega \overline{\mathcal{V}}_i(\overline{\omega}) = D_\omega u_i\left(\overline{x}_i^*(\overline{\omega})\right) + \beta\rho \, D_\omega \overline{\mathcal{V}}_i(\overline{\omega}).$$

Collecting terms: $(1 - \beta\rho) D_\omega \overline{\mathcal{V}}_i(\overline{\omega}) = D_\omega u_i(\overline{x}_i^*(\overline{\omega}))$. The right side is $D_\omega \overline{V}_i(\overline{\omega})$—the gradient of the static peaceful value function, by the envelope theorem applied to the stage game. Hence $D_\omega \overline{\mathcal{V}}_i(\overline{\omega}) = D_\omega \overline{V}_i(\overline{\omega})/(1 - \beta\rho)$, and dotting both sides with $\Delta\omega$ gives (7.2). ∎

The amplification factor $1/(1 - \beta\rho)$ has a clean interpretation. Because the endowment process is persistent ($\rho > 0$), a current decline in the lootable endowment predicts future declines as well. The peaceful value function $\overline{\mathcal{V}}_i$ capitalizes not just the one-period loss but the entire discounted stream of expected future losses. The amplification is therefore greater when states are more patient (β large) or when endowments are more persistent (ρ large). In the polar cases: $\rho = 0$ (IID shocks) gives $\tau_i^{\text{dyn}} = \tau_i$ (no amplification, since today's endowment reveals nothing about the future); $\rho \to 1$ (unit root) drives

$\tau_i^{\text{dyn}} \to -\infty$ for any $\tau_i < 0$ (a permanent decline in access triggers unbounded conflict incentives in a patient state).

The same amplification applies to the deviation gain. Define the dynamic deviation gain as $\mathcal{G}_i := \mathcal{V}_i^* - \overline{\mathcal{V}}_i$. By the same argument applied to the DCE Bellman equation in parallel:

7.8 Corollary (Dynamic deviation gain)

Under the conditions of Proposition 7.7, the dynamic deviation gain satisfies

$$\left.\frac{d\mathcal{G}_i}{dt}\right|_{t=0} = \frac{1}{1-\beta\rho} \cdot \left.\frac{dG_i}{dt}\right|_{t=0}.$$

When $\tau_i(\overline{\omega}) < 0$ *and the conditions of Proposition 6.8 hold,* $dG_i/dt|_{t=0} > 0$ *and hence* $d\mathcal{G}_i/dt|_{t=0} > 0$*. The dynamic conflict incentive is strictly larger than the static one, and it grows faster.*

7.9 Remark (Copeland and persistence)

Proposition 6.8 formalized Copeland's argument that anticipated trade cutoff—not current trade flows—drives conflict onset. Proposition 7.7 adds a new dimension: the strength of the conflict incentive is not just a function of the direction $\Delta\omega$ *but of the product* $\beta\rho$*. A state that expects a persistent decline in its lootable endowment (high* ρ*), and that weights the future heavily (high* β*), faces a dramatically amplified incentive to fight preemptively. This provides a microfounded account of why resource-dependent states facing secular decline in their resource base are particularly prone to conflict: the dynamic capitalization of a persistent negative trade expectation compounds the static incentive by* $1/(1-\beta\rho)$*, which can be large when* β *and* ρ *are both close to one.*

7.10 Remark (Patience and conflict)

The amplification factor also implies a counterintuitive relationship between patience and conflict. In bargaining models, patient states are often less *prone to war, because they can wait for a negotiated settlement. In the present model, a patient state facing declining trade access is* more *prone to war, because it capitalizes a larger fraction of the anticipated stream of losses into the current conflict incentive. These effects operate through different channels and are not mutually exclusive: patience reduces the cost of waiting but increases the value of fighting now when the future looks bad. Which effect dominates depends on the degree of persistence* ρ *and on the specific bargaining environment that determines the outside option.*

7.2.6 Value Function Iteration

The existence proof of Proposition 7.3 is constructive: the sequence $\{T_i^k V^{(0)}\}_{k\geq 0}$ converges to $\mathcal{V}_i^*$ for any starting point $V^{(0)} \in \mathcal{C}_b(K)$. This is the basis of *value function iteration* (VFI), the standard numerical method for solving infinite-horizon dynamic programs.

To implement VFI:

1. Discretize the endowment space K to a finite grid $\{\omega^1, \ldots, \omega^N\} \subset K$ and represent the transition kernel Π as an $N \times N$ stochastic matrix $P_{nm} = \Pi(\omega^n, \{\omega^m\})$.

2. Initialize the value function on the grid: set $V_n^{(0)} = 0$ for all n (or use the myopic payoff $\phi_i(\omega^n)$ as a warm start).

3. Iterate: $V_n^{(k+1)} = \phi_i(\omega^n) + \beta \sum_m P_{nm} V_m^{(k)}$.

4. Stop when $\max_n |V_n^{(k+1)} - V_n^{(k)}| < \varepsilon$.

At each grid point ω^n, computing $\phi_i(\omega^n)$ requires solving the static DCE—a task handled by the alternating algorithm of section 7.1.1. The outer VFI loop then treats these static DCE payoffs as fixed parameters and finds the dynamic value function by iterating the Bellman operator.

The convergence rate of VFI is β, meaning that the error decreases by a factor of β at each iteration. For $\beta = 0.9$, roughly 100 iterations are required to reduce the error by a factor of 10^{-5}. For $\beta = 0.95$, roughly 200 iterations are required. In practice, *policy function iteration* (Howard improvement), which alternates between policy evaluation and policy improvement steps, can achieve the same accuracy in far fewer iterations and is often preferred for high-β economies.

7.3 Summary

This chapter has examined the dual-competitive equilibrium from two complementary angles not accessible to the analytical results alone.

The numerical section populated the model and demonstrated three findings. First, equilibrium militarization is non-monotone in endowment: conflict rises with wealth at low and intermediate wealth levels, then declines as the wealthiest states approach the inert-entrant conditions of Definition 6.1 (fig. 7.1). Second, conflict is structurally clustered—not dyadic. In a multi-state economy, conflict organizes around resource-rich targets and is shaped by the global price vector that all states jointly determine (fig. 7.2). Third, the efficiency gap Δ^* grows with the number of states but is concave: the marginal efficiency cost of each additional state diminishes as per-state endowment shares fall (fig. 7.3).

The dynamic section embedded the static DCE in an infinite-horizon stochastic environment and derived the recursive dual-competitive equilibrium (RDCE). Existence and uniqueness of the RDCE (Proposition 7.3) follow from the contraction structure of the Bellman operator—specifically, from Blackwell's monotonicity and discounting conditions. Smoothness of the RDCE value functions (Proposition 7.5) follows from the smooth stage payoffs delivered by the static DCE (Corollary 3.12) and smooth transitions. The central dynamic result (Proposition 7.7) shows that trade expectations are amplified in the RDCE by the factor $1/(1-\beta\rho)$: an anticipated endowment decline that reduces peaceful utility by one unit today reduces the dynamic peaceful value function by $1/(1-\beta\rho)$ units. The conflict incentive is correspondingly amplified, with the amplification increasing in both patience β and endowment persistence ρ.

Together, the two sections show that the model's analytical core is both numerically tractable and dynamically extensible. The smooth equilibrium manifold that is the book's central mathematical object is not merely an abstract topological result: it is the foundation for stable numerical computation, smooth dynamic value functions, and differentiable policy rules that can be taken to data.

Chapter 8

Institutions and Social Choice

One question has been lurking beneath the surface of the analysis: *where are the institutions*? Chapter 5 established that the peaceful equilibrium is Pareto optimal and that equilibria with positive military deployment are generically inefficient (Proposition 5.10), with the inefficiency quantified as an efficiency gap $\Delta^* > 0$ (Proposition 5.11). The Second Welfare Theorem fails in this setting: collectively superior equilibria exist that the uncoordinated market cannot sustain (Proposition 5.16). This failure defines the *institutional enterprise*—the study of what governing arrangements are required to move the system from its conflict-ridden equilibrium toward normatively attractive ones.

A deceptively simple question underlies the enterprise: can the preferences of the I states over institutional arrangements be consistently aggregated into a collective plan? The canonical answer, due to Arrow (1963), is pessimistic: no aggregation procedure can simultaneously satisfy basic conditions of fairness without degenerating into dictatorship. This chapter reframes that obstruction and resolves it. The resolution, due to Chichilnisky (1982), is topological: on spaces with the right geometric structure, a continuous, collectively consistent aggregation always exists. The smooth equilibrium manifold E supplies precisely that structure, and the diffeomorphism $E \cong \mathbb{R}^{IL}$ of Corollary 3.12 is what puts us in a setting where it applies.

The analysis also engages the empirical rational-design program of Koremenos, Lipson and Snidal (2001), who identify five structural dimensions along which international institutions vary and link them to the cooperation problems that motivate design. We show that each KLS dimension has a precise counterpart in the gradient aggregation structure that the theory constructs.

8.1 Remark (The smoothness debate)

Gilpin (1981) distinguishes liberal views of international change—smooth, continuous, long-cycle theories (e.g., Modelski, 1987; Goldstein, 1988)—from Hegelian-Marxist views in which change is fundamentally discontinuous (Moore, 1958). Ruggie (1993, 143–4) laments that IR "is not very good as a discipline at studying the possibility of fundamental discontinuity." The analysis of this chapter does not dissolve this controversy, but it establishes its terms. The smoothness of E is not an empirical claim about world politics but a mathematical consequence of the model's structure. Its institutional significance follows from that: to echo Debreu (1972, 605), our data do not represent our concepts perfectly, and if the map from concepts to predictions is discontinuous, then measurement error of any magnitude yields essentially distinct predictions with no principled basis for choosing among them.

8.1 The Institutional Problem

Koremenos, Lipson and Snidal (2001) treat international institutions as rational, negotiated responses to three cooperation problems: distribution (heterogeneous preferences over which cooperative arrangement to aim for), enforcement (incentives to defect once others have committed), and uncertainty (about others' behavior, future conditions, and the consequences of different arrangements). In their framework, the five structural dimensions along which institutions vary—*membership*, *scope*, *centralization*, *control*, and *flexibility*—are design parameters whose values are explained by the nature of the cooperation problem being addressed.

Each of the three cooperation problems has a direct counterpart in the DCE framework.

Distribution. States have heterogeneous preferences over the equilibrium manifold E: the value function $V_i^*(e)$ assigns different rankings to the same equilibrium for different states. The efficiency frontier is high-dimensional—indexed by the entire distribution of gains from peace—and different states prefer different points on it. Any institutional plan must choose among them, generating distributional conflict even when all states agree that the current equilibrium is inefficient. The efficiency gap Δ^* of equation (5.3) quantifies the aggregate welfare cost of failing to resolve this conflict.

Enforcement. The peace condition (6.1) constrains which equilibria are self-enforcing. Along any credible institutional path, every intermediate equilibrium must satisfy the peace condition for every state: if at any stage state i finds that military exploitation is more profitable than compliance, the arrangement unravels. Proposition 6.8 establishes that the expected value of peaceful exchange is a key determinant of this constraint. A credible institutional path must lie entirely within the *peace set*

$$\mathcal{P} = \left\{ e \in E \;\middle|\; \text{condition (6.1) holds for all } i \right\}, \tag{8.1}$$

the set of equilibria at which no state prefers unilateral deviation. Proposition 6.2 and Definition 6.1 further show that the arrival of an inert entrant—a state with a large non-lootable endowment—expands $\mathcal{P}$ by driving up the prices of contested goods, making exploitation less attractive.

Uncertainty. Chapter 7 embedded the static DCE in a stochastic dynamic environment. Endowments evolve as a Markov process, the equilibrium configuration drifts over time (Proposition 7.3), and the dynamic trade expectation amplifies the static expectation by a factor $1/(1 - \beta\rho)$ (Proposition 7.7). Institutions designed for today's equilibrium must remain viable as the system evolves, motivating institutional flexibility. With many states, coordinating gradient fields that point in diverse directions grows increasingly costly, motivating the design choices KLS label control and membership.

The framework's distinctive contribution is to split the institutional design problem explicitly. *Target selection* asks which equilibrium the institution should aim for, given heterogeneous preferences. *Path selection* asks how the system should navigate from today's equilibrium to the chosen target. Standard political economy focuses almost exclusively on target selection, treating the transition as a black box. The smooth structure of E allows both to be addressed analytically.

8.2 The Space of Institutional Alternatives

An institutional *alternative* is a complete plan: it specifies both a target equilibrium and a trajectory connecting the current state to that target. This double specification is substantively important. Consider two plans that arrive at the same endpoint. One passes through equilibria that are favorable to state i early in the transition; the other achieves the same endpoint by a route that is initially costly for state i. These are genuinely distinct alternatives from i's perspective, even though their endpoints coincide. Under the dynamic extension of Chapter 7, a less patient state (lower discount factor β) places greater weight on the early stages of the path, making this distinction sharp.

8.2 Definition (Space of institutional alternatives)
Fix a current equilibrium $e_0 \in E$. *The* space of institutional alternatives *based at* e_0 *is*

$$\mathcal{A}(e_0) = \{\gamma : [0,1] \to E \mid \gamma(0) = e_0 \text{ and } \gamma \text{ is smooth}\}. \tag{8.2}$$

An element $\gamma \in \mathcal{A}(e_0)$ *is a* credible institutional path*: a smooth trajectory through the equilibrium manifold originating at the current equilibrium and terminating at the target* $\gamma(1) \in E$.

State i has preferences over $\mathcal{A}(e_0)$: it strictly prefers γ to γ' whenever $V_i^*(\gamma(1)) > V_i^*(\gamma'(1))$ (better endpoint), or the endpoints are equivalent but $V_i^*(\gamma(t)) > V_i^*(\gamma'(t))$ for intermediate stages t (better transition dynamics). Preferences encode both where the institution is headed and how it gets there.

The space $\mathcal{A}(e_0)$ is infinite-dimensional. Equipped with the topology of uniform convergence of all derivatives, it is a Fréchet manifold—a smooth, infinite-dimensional manifold without the local compactness of finite-dimensional spaces. The social choice problem asks: can a continuous aggregation of the states' preferences over $\mathcal{A}(e_0)$ be constructed so that the result satisfies basic normative conditions?

8.3 Arrow's Theorem and the IIA Obstruction

Before demonstrating that a consistent aggregation is possible, we must understand why it appears to be impossible. Arrow's theorem provides the sharpest statement of the difficulty.

A *social welfare function* (SWF) is a mapping F from profiles of individual strict preference orderings over $\mathcal{A}(e_0)$ to a collective strict preference ordering over $\mathcal{A}(e_0)$. Three conditions are standard:

Unanimity (P). If every state strictly prefers γ to γ', so does the collective: if $\gamma \succ_i \gamma'$ for all i, then $\gamma \succ_F \gamma'$.

Independence of Irrelevant Alternatives (IIA). The collective ranking of γ versus γ' depends only on how individual states rank γ versus γ'—not on how they rank any other alternatives.

Non-dictatorship (ND). No single state d is such that $\gamma \succ_d \gamma'$ always implies $\gamma \succ_F \gamma'$, regardless of the other states' preferences.

8.3 Theorem (Arrow, 1963)
Let $|\mathcal{A}(e_0)| \geq 3$ *and* $I \geq 2$*. Any social welfare function satisfying* P, IIA, *and* ND *is a dictatorship: there exists a state* d *such that* $F(\succ_1, \ldots, \succ_I) = \succ_d$ *for all preference profiles.*

The theorem applies directly to institutional paths. For any three proposed plans $\gamma_1, \gamma_2, \gamma_3 \in \mathcal{A}(e_0)$, no aggregation of the states' rankings satisfies all three conditions simultaneously. The states may agree on which target equilibrium is desirable while disagreeing over the trajectory—and IIA prevents the aggregation from using information about alternative paths to resolve this disagreement.

Why IIA is the obstruction. The independence condition has a natural topological interpretation in the smooth setting (Chichilnisky, 1982). Preferences over a smooth space can be represented as smooth preference maps—sections of a fiber bundle over the domain. IIA requires the social choice function to depend only on the local pairwise comparison of two alternatives, which forces it to be a *fiber-preserving map* in this bundle. On spaces with non-trivial topology, any fiber-preserving map that respects unanimity must be dictatorial: the topological obstructions in the preference manifold force the aggregation to track a single agent's preferences globally.

The key insight is that IIA requires too much *locality* from the aggregation. It insists that the ranking of γ versus γ' be made without reference to any third alternative—a condition that, in the smooth setting, prevents the aggregation from integrating global information about the preference landscape. When the preference space has the right topology, this restriction is fatal: local comparisons cannot be coherently assembled into a global ranking.

8.4 The Topological Escape: Chichilnisky on the Path Space

Arrow's impossibility is not an impossibility of consistent aggregation as such. It is an impossibility of a specific type of aggregation—one satisfying the ordinal IIA condition. If IIA is replaced by a topological continuity condition, the picture changes entirely.

8.4 Definition (Chichilnisky social choice function)
A mapping $F : \mathrm{Pref}(\mathcal{A}(e_0))^I \to \mathrm{Pref}(\mathcal{A}(e_0))$ *is a* Chichilnisky social choice function *if it is (i)* continuous*: small changes in individual preferences produce small changes in the collective preference; (ii)* anonymous*: the output is unchanged when states' preference profiles are permuted; and (iii)* unanimity-respecting*: if all states share the same preference ordering* $\succ$*, then* $F(\succ, \dots, \succ) = \succ$*.*

Continuity is strictly weaker than IIA. It allows the collective preference to depend on the global structure of individual preferences, not just the pairwise comparison of two alternatives. This is the relaxation that resolves Arrow's impossibility.

8.5 Theorem (Chichilnisky, 1982)
A Chichilnisky social choice function on a smooth domain X *exists if and only if* X *is contractible.*

The theorem reduces the social choice problem to a topological question: is $\mathcal{A}(e_0)$ contractible?

8.6 Proposition
The space of institutional alternatives $\mathcal{A}(e_0)$ *is contractible.*

Proof. Define the homotopy $H : [0,1] \times \mathcal{A}(e_0) \to \mathcal{A}(e_0)$ by

$$H(s, \gamma)(t) = \gamma(st), \qquad s, t \in [0,1]. \tag{8.3}$$

At $s = 1$, $H(1, \gamma) = \gamma$ (the identity on $\mathcal{A}(e_0)$). At $s = 0$, $H(0, \gamma)(t) = \gamma(0) = e_0$ for all t: every path collapses to the constant path at e_0. For each fixed s, the map $t \mapsto \gamma(st)$ is a smooth re-parameterization of γ—it traverses only the initial segment $\gamma([0, s])$ by time $t = 1$. Since γ is smooth and $s \mapsto \gamma(st)$ varies continuously in the topology of $\mathcal{A}(e_0)$, H is a continuous contraction to the constant path. ■

The proof has a direct institutional interpretation. Any institutional path γ can be continuously deformed to the "status quo" path (constant at e_0) simply by slowing it down. As $s \to 0$, the path traverses less and less of its original trajectory, converging to a stationary arrangement. This deformation stays within E at every stage, because a re-parameterization of an equilibrium path is still an equilibrium path.

8.7 Corollary

There exists a continuous, anonymous social choice function on $\mathcal{A}(e_0)$ that respects unanimity.

Proof. Proposition 8.6 establishes that $\mathcal{A}(e_0)$ is contractible. Theorem 8.5 then guarantees the existence of a Chichilnisky social choice function. ∎

The role of $E \cong \mathbb{R}^{IL}$. The contractibility of $\mathcal{A}(e_0)$ holds for any path-connected E: the path-slowing homotopy requires only that re-parameterizations of paths remain in E, which follows from E being path-connected. However, Theorem 8.5 in its full strength requires the domain to carry a *smooth* structure compatible with preferences. The diffeomorphism $E \cong \mathbb{R}^{IL}$ (Corollary 3.12) ensures this: the smooth structure of E is inherited from $\mathbb{R}^{IL}$, making $\mathcal{A}(e_0)$ a smooth Fréchet manifold to which the smooth version of Chichilnisky's theorem applies. Without smoothness of E, only a merely continuous aggregation would be guaranteed; the smooth version is stronger—it ensures that differentiable perturbations to individual preferences produce differentiable changes in the collective preference. This stronger guarantee is what licenses the gradient-field construction of the next section.

The Arrow-to-Chichilnisky transition corresponds to a substantive move in the institutions literature: from the ordinal aggregation of pure game-theoretic bargaining models (which inherit Arrow's impossibility in full force) to the smooth preference aggregation framework, where the differential structure of E becomes a resource for positive institutional design.

8.5 Gradient Fields as Policy Instruments

Corollary 8.7 is an existence result: it guarantees that a collectively consistent aggregation of institutional preferences is possible without specifying one. To make the aggregation constructive—to identify a specific institutional policy that the collective preferences determine—we need a smooth representation of individual preferences over $\mathcal{A}(e_0)$.

The natural candidate comes from the DCE value functions.

8.8 Definition (Value function on the equilibrium manifold)

For each state i, define $\tilde{u}_i : E \to \mathbb{R}$ by $\tilde{u}_i(e) = V_i^*(e)$: the equilibrium payoff of state i at configuration e. By Corollary 3.12 and the smooth dependence of the DCE on parameters (Theorem 4.8), the function $\tilde{u}_i$ is smooth on E.

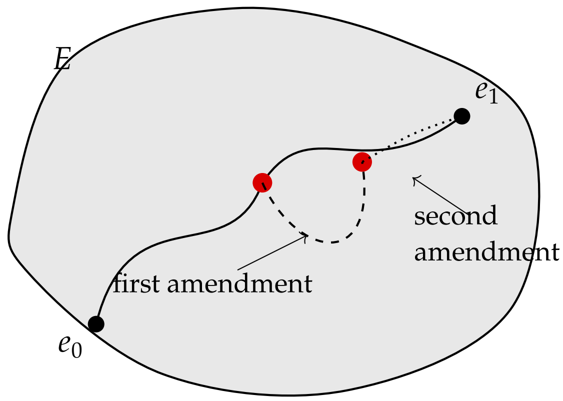

Figure 8.1: *Three elements of the path space $\mathcal{A}(e_0)$ connecting e_0 to a common target e_1: the original proposal (solid line) and two amendments (dashed and dotted). All three paths, along with every other element of $\mathcal{A}(e_0)$, can be continuously deformed into one another and into the constant path at e_0 via the homotopy $H(s, \gamma)(t) = \gamma(st)$. The contractibility of $\mathcal{A}(e_0)$ is the topological condition that licenses Corollary 8.7.*

The function $\tilde{u}_i$ encodes i's ranking of equilibria as a smooth scalar field on E. Its gradient encodes the direction of improvement: at each equilibrium e, the gradient of $\tilde{u}_i$ points toward the nearby equilibria that most benefit state i.

8.9 Definition (Individual gradient field)

The gradient field *of state i is the smooth vector field*

$$\operatorname{grad}(\tilde{u}_i) : E \to TE, \tag{8.4}$$

defined by the condition that for any smooth curve $c : (-\epsilon, \epsilon) \to E$ with $c(0) = e$,

$$\frac{d}{dt}(\tilde{u}_i \circ c)\bigg|_{t=0} = \langle \operatorname{grad}(\tilde{u}_i)(e), \dot{c}(0) \rangle,$$

where $\langle \cdot, \cdot \rangle$ is the inner product on $T_e E$ induced by the diffeomorphism $E \cong \mathbb{R}^{IL}$. The integral curve of $\operatorname{grad}(\tilde{u}_i)$ *from e_0 is state i's* preferred institutional trajectory*: the smooth path through E along which i's welfare increases most rapidly at every stage.*

The gradient field $\operatorname{grad}(\tilde{u}_i)$ translates i's preferences over equilibria into a directed policy instrument: a tangent vector at each point indicating the direction the system should move to benefit state i. The design problem is to aggregate these I vector fields into a single social direction.

8.10 Proposition (Social gradient field)
Let $\{w_i : E \to [0,1]\}_{i=1}^{I}$ be continuous weight functions with $\sum_{i=1}^{I} w_i(e) = 1$ for all $e \in E$. The vector field

$$V^{\text{soc}}(e) = \sum_{i=1}^{I} w_i(e)\ \text{grad}(\tilde{u}_i)(e) \tag{8.5}$$

is a Chichilnisky social choice function on the space of gradient fields of $\{\tilde{u}_i\}$: it is continuous, anonymous when $w_i \equiv 1/I$, and unanimity-respecting. Conversely, any Chichilnisky social choice function on the space of gradient fields takes this form.

Proof (Proof sketch). Continuity follows from the continuity of each w_i and the smoothness of the $\text{grad}(\tilde{u}_i)$. Unanimity is immediate: if $\text{grad}(\tilde{u}_i)(e) = v$ for all i and e, then $V^{\text{soc}}(e) = v$. Anonymity holds when $w_i \equiv 1/I$. For the converse: the space $\Gamma(TE)$ of smooth vector fields on $E \cong \mathbb{R}^{IL}$ is a locally convex Fréchet space, hence contractible. On a contractible locally convex space, any continuous, anonymous, unanimity-respecting map $\Gamma(TE)^I \to \Gamma(TE)$ must be a pointwise convex combination of its arguments—a consequence of the Chichilnisky characterization theorem in the linear setting. The representation (8.5) follows. ■

The integral curves of V^{soc} define the *social institutional trajectory*: the smooth path $\gamma^{\text{soc}} : [0,\infty) \to E$ along which the system evolves under the collectively determined policy. Different weight profiles $(w_i)_i$ represent different configurations of institutional influence. Equal weights ($w_i \equiv 1/I$) represent a democratically symmetric arrangement. A profile concentrated on one state ($w_d \approx 1$) is near-dictatorial. Weight functions that vary across E represent adaptive institutions that rebalance influence as circumstances change.

The credibility constraint. Not every social trajectory is feasible: the path must remain within the peace set $\mathcal{P}$ of equation (8.1).

8.11 Proposition (Credibility constraint)
A social institutional trajectory γ^{soc} is credible if and only if $\gamma^{\text{soc}}(t) \in \mathcal{P}$ for all $t \geq 0$. When $\mathcal{P}$ is non-empty, the credibility constraint restricts the feasible weight profiles $(w_i)_i$: only weight functions whose integral curves remain within $\mathcal{P}$ generate self-enforcing policies.

Proof. By definition, $e \in \mathcal{P}$ if and only if condition (6.1) holds for all i at e. Along the trajectory γ^{soc}, this condition holds at every stage if and only if $\gamma^{\text{soc}}(t) \in \mathcal{P}$ for all t. If the trajectory exits $\mathcal{P}$ at some stage t_0, then some state finds unilateral deviation profitable at $e = \gamma^{\text{soc}}(t_0)$, and the arrangement ceases to be self-enforcing from that point forward. ∎

The interplay between the social gradient and the credibility constraint is the central design tension. The direction $V^{\text{soc}}(e)$ may, at some configurations, point toward regions outside $\mathcal{P}$. The designer must then choose between modifying the weight profile (adjusting the direction of the social gradient to stay within $\mathcal{P}$) or accepting a feasible but collectively less preferred trajectory. This tradeoff—between preference satisfaction and enforcement feasibility—is precisely the distribution-enforcement tension that Koremenos, Lipson and Snidal (2001) identify as central to institutional design. The gradient framework makes it analytically precise.

8.6 KLS Dimensions as Aggregation Parameters

The theoretical framework now admits a direct reinterpretation of the KLS design dimensions. Each of the five dimensions corresponds to a structural parameter of the gradient aggregation problem.

KLS Dimension	Gradient Aggregation Parameter
Membership	The index set over which the gradient fields are aggregated. Restricted membership ($I' \subsetneq I$) drops some gradient fields from the sum in (8.5).
Control	The weight profile $(w_i)_i$. Asymmetric control means $w_d \gg w_j$ for $j \neq d$: one state's gradient dominates. Symmetric control means $w_i \equiv 1/I$.
Centralization	Whether V^{soc} derives from a centralized welfare function $W = \sum_i \alpha_i \tilde{u}_i$ (utilitarian: $V^{\text{soc}} = \text{grad}(W)$) or from a decentralized weighted sum of individual gradients. These coincide for constant w_i but diverge when w_i varies with e.
Scope	Which dimensions of E the gradient is computed over. An institution covering only a subset $\mathcal{L}' \subsetneq \mathcal{L}$ of goods restricts the gradient to the corresponding submanifold of E.
Flexibility	Whether the weight functions $w_i(e)$ depend on the current equilibrium e (adaptive flexibility) or are fixed constants (rigid institution).

With this mapping in place, three KLS conjectures follow from properties of gradient aggregation.

KLS C4: Centralization increases with enforcement problems. Severe enforcement problems mean $\mathcal{P}$ is small: only a narrow region of E satisfies the peace condition for all states. A decentralized weighted sum of gradient fields may produce a social trajectory that exits $\mathcal{P}$: no individual state, following its own gradient, guarantees that the aggregate direction is credibility-feasible. Centralization—constructing V^{soc} from a single welfare function W whose gradient is constrained to remain in $\mathcal{P}$—allows the designer to enforce the credibility constraint directly. Severe enforcement problems therefore call for centralization to coordinate on a feasible social trajectory.

KLS S2: Scope increases with distribution problems. When states' value functions $\tilde{u}_i$ differ substantially, their gradient fields $\text{grad}(\tilde{u}_i)$ point in very different directions on E. Equal-weight aggregation then produces a social gradient that is nearly orthogonal to every individual gradient—a policy that serves no state's interests well. Expanding scope adds issue dimensions to E, enlarging the space over which the gradients are computed. In a higher-dimensional space, gradient fields are more likely to be compatible: a state that cares intensely about one issue dimension can accept an unfavorable direction on another. Scope expansion is therefore a rational response to distribution problems, because it gives the gradient aggregation more room to find a direction that serves all states' interests simultaneously.

KLS M1: Restricted membership increases with enforcement problems. With many states, the social gradient is an average of many individual fields. If even a few states have gradient fields that point toward regions outside $\mathcal{P}$, the aggregate direction may violate the credibility constraint regardless of the other states' weights. Restricting membership to states whose gradients are more closely aligned with the credibility-feasible region of E reduces this problem: fewer conflicting fields must be aggregated, and the credibility constraint is easier to satisfy. Restricted membership is therefore a rational response to severe enforcement problems—it shrinks the aggregation problem to one that is credibility-feasible.

The gradient framework thus connects the formal model to the empirical rational-design program. The design dimensions that KLS identify as varying systematically across international institutions correspond to parameters of the social gradient aggregation that rational designers would adjust in response to the distribution, enforcement, and number problems the institution is created to address.

8.7 Summary

This chapter has constructed the institutional implications of the DCE framework in three steps.

The first step was to define the appropriate space of alternatives. An institutional arrangement is not merely a target equilibrium but a *path in E*—a smooth trajectory from the current configuration to a desired one. The space of such paths, $\mathcal{A}(e_0)$, is the natural domain for the social choice problem, because it encodes both where the institution is headed and how it gets there.

The second step identified the central obstacle and its resolution. Arrow's theorem implies that no ordinal aggregation of the states' rankings can satisfy unanimity, IIA, and non-dictatorship simultaneously; in the smooth setting, IIA forces the aggregation to be a fiber-preserving map, which must be dictatorial on spaces with non-trivial topology. The escape is provided by Chichilnisky's theorem: replace IIA with topological continuity, and a consistent aggregation exists whenever the domain is contractible. The path space $\mathcal{A}(e_0)$ is contractible—by the path-slowing homotopy $H(s, \gamma)(t) = \gamma(st)$—so the Chichilnisky aggregation exists (Corollary 8.7).

The third step made the aggregation constructive. Individual gradient fields $\mathrm{grad}(\tilde{u}_i)$ encode each state's welfare gradient as a smooth vector field on E; their weighted sum V^{soc} defines the social institutional trajectory. The diffeomorphism $E \cong \mathbb{R}^{IL}$ plays three roles in this construction: it guarantees that the $\tilde{u}_i$ are smooth, so gradients are well-defined; it makes $\mathcal{A}(e_0)$ a smooth Fréchet manifold, so the smooth version of Chichilnisky's theorem applies; and it ensures that the peace set $\mathcal{P}$ is a well-defined subset of E, so the credibility constraint is analytically tractable. These are not separate results—they are facets of a single diffeomorphism.

The KLS design dimensions—membership, scope, centralization, control, and flexibility—correspond directly to structural parameters of the gradient aggregation: the index set I, the submanifold of E over which the gradient is computed, the welfare function from which the gradient is derived, the weight profile $(w_i)_i$, and whether those weights are state-dependent. The KLS conjectures follow from properties of gradient aggregation in the DCE, connecting the formal model to the empirical rational-design program.

The obstacles to cooperation remain fully intact. The credibility constraint (Proposition 8.11) expresses their force precisely: a socially preferred trajectory is only implementable if it stays within $\mathcal{P}$, and constructing weight profiles that achieve this under severe enforcement problems may require centralization, scope expansion, or restricted membership—exactly the design responses KLS document empirically. What the topological analysis establishes is that the geometry of E is not itself the obstacle. The space of institutional alternatives is contractible, a collective aggregation exists, and the gradient framework provides a constructive implementation. The obstacles to peace and cooperation are political, not topological.

Chapter 9

Conclusion

This book has argued that a unified treatment of trade and conflict—one that takes seriously the joint determination of market prices and military deployments—produces a richer analytical framework than either tradition provides on its own. The central result is not merely that a dual-competitive equilibrium exists, but that the equilibria form a smooth manifold with exceptional geometric properties. Those properties are the source of the book's substantive contributions.

The diffeomorphism $E \cong \mathbb{R}^{IL}$ (Corollary 3.12) does three things that the separated approaches could not accomplish. It makes the comparative statics enterprise rigorous: the equilibrium mapping is smooth, so derivatives are meaningful, and the signs and magnitudes of equilibrium responses to parameter changes are analytically determinate (Theorem 4.8). It makes welfare measurement possible: the efficiency gap Δ^* (Proposition 5.11) is a smooth, well-defined scalar precisely because E has the geometry of Euclidean space—ordinal comparisons across equilibria become cardinal measurements. And it makes consistent institutional design achievable: the contractibility of the path space follows from E being path-connected, and this is the topological condition that licenses a social choice function over institutional paths (Corollary 8.7) and that makes the gradient-field construction of Proposition 8.10 well-defined.

These are not three separate uses of a technical result. They are three facets of the same underlying claim: that the international system, modeled this way, has enough structure to support precise analysis at the positive, normative, and institutional levels simultaneously.

Without joint determination, these connections are unavailable. The trade expectations result (Proposition 6.8) cannot be stated in a framework where prices and deployments are not simultaneously determined: the question is how a change in one affects the other at equilibrium, and that question only makes sense when both are equilibrium objects. Similarly, the efficiency gap requires comparing the armed equilibrium to the peaceful benchmark in the same space E; the comparison is coherent because the diffeomorphism places both configurations in the same smooth manifold. The systemic conflict result (Proposition 6.2) is a statement about how an exogenous change in the endowment profile shifts equilibrium deployments across all states simultaneously—a genuinely general equilibrium claim that a partial analysis cannot support.

Limitations

The model abstracts from several features of international politics that a complete theory would need to address.

Complete information. States have full knowledge of each other's endowments and utility functions. In practice, uncertainty about capabilities and intentions is central to why wars occur and why institutions are designed the way they are (e.g., Fearon, 1995). Introducing private information would complicate the equilibrium concept considerably—market prices would need to aggregate private signals—but would connect the framework to the extensive literature on information and conflict.

Domestic politics. The state is treated as a unitary actor. The rational-design literature (Koremenos, Lipson and Snidal, 2001) recognizes that domestic distributional conflicts often shape institutional design choices more than international bargaining does. The gradient-field framework of Chapter 8 could in principle be extended to incorporate domestic interest groups as additional agents in the aggregation—expanding the index set I beyond states—but the formal model does not do so.

Strategic dynamics. The dynamic extension of Chapter 7 addresses stochastic variation in endowments but treats the stage game as stationary: the equilibrium concept at each date is the static DCE, with the dynamic layer adding discounted continuation values. Strategic dynamics—arms races, gradual escalation, the acquisition of capabilities over time—are not modeled. These would require a richer state space and a more demanding existence argument than the Blackwell contraction used in Proposition 7.3.

Extensions

Several natural extensions of the framework suggest themselves.

Sanctions and coercion. Economic sanctions represent an institutional instrument that operates within E by distorting the effective endowment profile—reducing the target state's access to markets without requiring the coordinated path-following that institutions in Chapter 8 envision. The comparative statics of Theorem 4.8 provide a natural starting point for analyzing the equilibrium effects of sanctions on prices, deployments, and welfare.

Alliances. An alliance is a commitment that changes the peace condition: by making the marginal cost of deviation higher, an alliance partner effectively expands the peace set $\mathcal{P}$ of equation (8.1). Formalizing this as a perturbation of the peace condition—and studying the comparative statics of alliance formation on equilibrium deployments and market prices—would connect the framework to the formal alliances literature.

Finance. The Walrasian layer of the model covers commodity markets but not capital markets. Sovereign debt, foreign direct investment, and the interaction of financial flows with military competition are natural extensions, particularly given the amplification result of Proposition 7.7: if financial integration increases the effective discount factor β, the dynamic trade expectation grows, and the incentive structure of the peace condition shifts accordingly.

The book's aim has been to give war and peace a common mathematical home. The dual-competitive equilibrium is that home: the setting in which exchange and conflict, market prices and military deployments, welfare analysis and institutional design can all be addressed with the same tools. The topological structure of E is not incidental to this unification—it is what makes the analysis tractable and what connects results across chapters that would otherwise appear to share only a common setting. The obstacles to cooperation identified in Chapter 8 are political, not topological. Whether they are surmountable is an empirical question; that they are analytically tractable is this book's answer.

Appendix A

Mathematical Primer

This primer collects the mathematical background needed to follow the formal arguments of the book. It covers five topics: smooth maps and the chain rule (§A.1); regular values and the preimage theorem (§A.2); smooth manifolds and diffeomorphisms (§A.3); topological degree (§A.4); and contractibility (§A.5). In each section the book's own objects serve as the running illustration: the demand and deployment functions of Chapter 2, the equilibrium manifold E and the diffeomorphism $E \cong \mathbb{R}^{IL}$ from Chapter 3, and the path space $\mathcal{A}(e_0)$ from Chapter 8.

Readers whose background includes differential topology can skip the primer entirely. Readers trained in game theory, general equilibrium theory, or formal international relations theory who find the topology unfamiliar may want to read it before or alongside Chapters 3 and 8. Standard references are Spivak (1965) (smooth maps and the preimage theorem), Milnor (1965) (degree theory), and Balasko (2009) (equilibrium manifold theory, including the degree argument for existence).

A.1 Smooth Maps and the Chain Rule

Smooth functions

A function $F : U \to \mathbb{R}^m$, with $U \subset \mathbb{R}^n$ open, is *smooth* (or C^∞) if all partial derivatives of all orders exist and are continuous on U. Polynomials, exponentials, and compositions of such functions are smooth; functions with kinks or jumps are not.

The derivative as a linear map

The derivative of F at a point $x \in U$ is the linear map

$$DF(x) : \mathbb{R}^n \to \mathbb{R}^m$$

defined by the condition that $F(x+h) = F(x) + DF(x)h + o(\|h\|)$ as $h \to 0$. In coordinates, $DF(x)$ is the $m \times n$ Jacobian matrix $(\partial F^j / \partial x^k)$. The operative point is that $DF(x)$ is *linear*: it approximates F near x to first order. When F is smooth, $DF(x)$ itself varies smoothly as x varies.

The chain rule

If $F : U \to V$ and $G : V \to W$ are smooth, the composition $G \circ F$ is smooth and

$$D(G \circ F)(x) = DG(F(x)) \circ DF(x).$$

In matrix terms: the Jacobian of the composition is the product of the Jacobians, evaluated at the appropriate points.

How this appears in the book

The demand function f_i and the deployment function μ_i of Chapter 2 are smooth maps from the spaces of prices, others' deployments, and endowments into commodity space and the deployment interval. Their smoothness is established in Proposition 2.13: the state's problem has a strictly concave objective and a smooth constraint, so the solution varies smoothly with the parameters.

The chain rule appears prominently in the positive theory. The reduced excess demand is $\widehat{d}(p, \omega) = \overline{d}(p, m^*(p, \omega), \omega)$, where m^* is the Nash equilibrium deployment profile. Differentiating with respect to p via the chain rule gives the decomposition

$$J = J^{\text{ex}} + J^{\text{conf}}, \qquad J^{\text{conf}} = D_m \overline{d}\big|_{m^*} \cdot D_p m^*.$$

The term $D_p m^*$ is itself a chain-rule derivative: it comes from differentiating the fixed-point condition $m^* = \mu(p, m^*, \omega)$ with respect to p, which requires the implicit function theorem described in the next section. The decomposition separates the direct effect of price changes on excess demand (J^{ex}, holding deployments fixed) from the indirect effect that flows through the conflict layer (J^{conf}, as price changes shift the Nash equilibrium).

A.2 Regular Values and the Preimage Theorem

Regular points and regular values

Let $F : U \to \mathbb{R}^m$ be smooth, with $U \subset \mathbb{R}^n$ open. A point $x \in U$ is a *regular point* of F if the Jacobian $DF(x)$ has full row rank m. A value $c \in \mathbb{R}^m$ is a *regular value* of F if every $x \in F^{-1}(c)$ is a regular point; otherwise c is a *critical value*.

Sard's theorem states that the set of critical values has Lebesgue measure zero in $\mathbb{R}^m$ (Spivak, 1965). Regular values are therefore generic: most values in the codomain are regular, and a perturbation of any critical value can be made regular.

The preimage theorem

If c is a regular value of $F : U \subset \mathbb{R}^n \to \mathbb{R}^m$ (with $m \leq n$), then the preimage $F^{-1}(c)$ is a smooth submanifold of $\mathbb{R}^n$ of dimension $n - m$.

This is the workhorse for constructing smooth manifolds. Instead of defining a manifold abstractly via charts, one defines it as the zero set of a smooth map and verifies that zero is a regular value. The dimension count $n - m$ records the degrees of freedom left after the m equations are imposed.

The implicit function theorem

The implicit function theorem is the preimage theorem applied to the parametric setting. Suppose $F(y, z) = 0$ where $F : \mathbb{R}^k \times \mathbb{R}^m \to \mathbb{R}^m$. If the partial Jacobian $D_z F(y_0, z_0)$ is invertible (square and full rank) at a solution (y_0, z_0), then there exists a smooth function $z^* : U \to \mathbb{R}^m$ defined on a neighborhood U of y_0 such that $F(y, z^*(y)) = 0$ and $z^*(y_0) = z_0$. The solution z can be expressed as a smooth function of the parameter y near the solution point.

How this appears in the book

Nash uniqueness (Proposition 3.5). The Nash condition is $r(p, m, \omega) := m - \mu(p, m, \omega) = 0$. The Jacobian $D_m r = I - D_m \mu$ is invertible because the contraction property of Lemma 3.3 ensures all eigenvalues of $D_m \mu$ have modulus less than one, so $I - D_m \mu$ has no zero eigenvalues. The implicit function theorem then gives a smooth function $m^*(p, \omega)$ satisfying $r(p, m^*(p, \omega), \omega) = 0$: the

Nash equilibrium deployment profile exists and varies smoothly with prices and endowments.

The equilibrium manifold E (Proposition 3.7). The equilibrium manifold is $E = \widehat{d}^{-1}(0)$, where $\widehat{d} : S_N \times \Omega_0 \to \mathbb{R}^{L-1}$. The ambient space $S_N \times \Omega_0$ has dimension $L-1+IL$ and the target has dimension $L-1$, so the preimage theorem would give a manifold of dimension IL—if 0 is a regular value. The proof of Proposition 3.7 shows that the partial Jacobian $D_\omega \widehat{d}$ has full rank $L-1$ at every equilibrium, using the freedom to vary each state's endowment independently. Hence 0 is a regular value, $E = \widehat{d}^{-1}(0)$ is a smooth submanifold of dimension IL, and the preimage theorem does the rest.

Generic finiteness (Proposition 3.16). Sard's theorem applied to the natural projection $\pi : E \to \Omega_0$ gives a measure-zero critical set $\Sigma \subset \Omega_0$. Above every regular endowment $\omega \notin \Sigma$, the fiber $\pi^{-1}(\omega)$ is discrete—finitely many equilibrium prices, each isolated.

A.3 Smooth Manifolds and Diffeomorphisms

Manifolds

A *smooth manifold* M of dimension k is a topological space that locally resembles $\mathbb{R}^k$: around every point $p \in M$ there is an open neighborhood homeomorphic to an open subset of $\mathbb{R}^k$, with smooth transition maps where neighborhoods overlap. In this book all manifolds arise as preimages of regular values (§A.2), so the abstract definition is mostly background.

The dimension of a manifold counts the degrees of freedom for moving within it. For the equilibrium manifold, the ambient space $S_N \times \Omega_0$ has dimension $L-1+IL$ and the equilibrium equations impose $L-1$ constraints, leaving IL degrees of freedom.

Diffeomorphisms

A *diffeomorphism* $\Phi : M \to N$ between smooth manifolds is a smooth bijection with a smooth inverse. Diffeomorphic manifolds are indistinguishable from the perspective of smooth analysis: any smooth construction on N (integrals, differential equations, gradients) can be pulled back to M via Φ, and vice versa. In particular, M and N have the same topology, the same dimension, and the same calculus.

The diffeomorphism $E \cong \mathbb{R}^{IL}$ of Proposition 3.10 is constructed explicitly. The map $\Phi : E \to S_N \times \mathbb{R}^I_{++} \times \mathbb{R}^{(L-1)(I-1)}$ sends each equilibrium to the triple $(p, w, \overline{\omega})$: the price vector, the diminished wealth levels, and the endowment components not pinned by the wealth constraints. One verifies directly that Φ is a smooth bijection with a smooth inverse, and that the target space is diffeomorphic to $\mathbb{R}^{IL}$ by coordinate transformations (the simplex S_N is diffeomorphic to $\mathbb{R}^{L-1}$; the positive orthant $\mathbb{R}^I_{++}$ is diffeomorphic to $\mathbb{R}^I$ via the componentwise logarithm).

Tangent spaces and gradient fields

At each point $e \in E$, the *tangent space* $T_e E$ is the vector space of all velocity vectors of smooth curves through e: if $c : (-\epsilon, \epsilon) \to E$ is smooth with $c(0) = e$, then $\dot{c}(0) \in T_e E$. The tangent space captures all infinitesimal directions of motion from e that remain on the manifold. For a smooth submanifold of $\mathbb{R}^n$, $T_e E$ is a linear subspace of $\mathbb{R}^n$ of dimension $\dim E$. The *tangent bundle* $TE = \bigsqcup_{e \in E} T_e E$ assembles all the tangent spaces. A *smooth vector field* is a smooth assignment $V : E \to TE$ with $V(e) \in T_e E$ for every e.

The diffeomorphism $E \cong \mathbb{R}^{IL}$ equips each $T_e E$ with an inner product: for $u, v \in T_e E$, set $\langle u, v \rangle = \langle D\Phi(e)\, u,\ D\Phi(e)\, v \rangle_{\mathbb{R}^{IL}}$, the Euclidean inner product pulled back from $\mathbb{R}^{IL}$ via the derivative of Φ. This inner product makes the gradient well-defined: for a smooth function $f : E \to \mathbb{R}$, the *gradient* $\mathrm{grad}(f)(e) \in T_e E$ is the unique tangent vector satisfying

$$\frac{d}{dt}(f \circ c)\Big|_{t=0} = \big\langle \mathrm{grad}(f)(e),\ \dot{c}(0) \big\rangle$$

for every smooth curve c through e. Intuitively, the gradient points in the direction of steepest ascent of f within E.

How this appears in the book

The value function $\tilde{u}_i : E \to \mathbb{R}$ of Definition 8.8 assigns to each equilibrium e the payoff $V_i^*(e)$ of state i. Its gradient $\mathrm{grad}(\tilde{u}_i) : E \to TE$ is a smooth vector field on E (Definition 8.9). At each equilibrium e, the vector $\mathrm{grad}(\tilde{u}_i)(e)$ points toward the nearby equilibria that most benefit state i. The social gradient field V^{soc} of Proposition 8.10 is a weighted convex combination of these individual gradient fields, and its integral curves define the collectively preferred institutional trajectory through E.

A.4 Topological Degree

Counting preimages

Let $f : M \to N$ be a smooth, proper map between smooth manifolds of equal dimension. (Recall that *proper* means the preimage of every compact set is compact—points cannot escape to infinity.) Choose a regular value $c \in N$. The preimage $f^{-1}(c)$ is a finite set. At each $x \in f^{-1}(c)$, the derivative $Df(x)$ is an invertible linear map from T_xM to T_cN; it either preserves or reverses orientation. Assign +1 or −1 accordingly. The *degree* $\deg(f)$ is the signed sum of these signs. It is independent of the regular value chosen.

The *modulo*-2 *degree* $\deg_2(f)$ counts only the parity of $|f^{-1}(c)|$. It requires no orientation and applies to all smooth manifolds.

Key properties

1. **Homotopy invariance.** If f and g are connected by a proper smooth homotopy $H : M \times [0, 1] \to N$, then $\deg(f) = \deg(g)$ and $\deg_2(f) = \deg_2(g)$.

2. **Surjectivity.** If $\deg_2(f) = 1$, then f is surjective: every point in N has at least one preimage.

3. **Normalization.** A diffeomorphism has $\deg = \pm 1$ (and $\deg_2 = 1$).

Homotopy invariance is the operative property. To compute the degree of a hard map, deform it continuously to a simpler map whose degree is known; the degree cannot change during a proper deformation.

How this appears in the book

The existence of a DCE (Theorem 3.14) is proved by showing that the natural projection $\pi : E \to \Omega_0$ has $\deg_2(\pi) = 1$, which implies surjectivity: every endowment profile $\omega \in \Omega_0$ has at least one equilibrium price above it.

The degree is computed by continuous deformation. Define the scaled battle function $s_i^t = t \cdot s_i$, scaling conflict from absent ($t = 0$) to actual ($t = 1$). The corresponding natural projection $\pi^t : E^t \to \Omega_0$ varies smoothly in t, and the desirability condition ensures each π^t is proper. The family $\{\pi^t\}$ is therefore a proper homotopy from π^0 to $\pi^1 = \pi$. At $t = 0$ there is no conflict and E^0 is the manifold of a standard smooth exchange economy; its natural projection has $\deg_2 = 1$ by Balasko (2009, Proposition 4.6.3). Homotopy invariance then gives $\deg_2(\pi) = 1$, and existence follows.

The argument has a conceptual interpretation: the equilibria of the dual-competitive economy are continuous deformations of the equilibria of the peaceful economy. Conflict distorts prices and reallocates resources but does not create or destroy equilibria in the parity sense—the modulo-2 count is preserved because the deformation is proper and the initial count is odd.

A.5 Contractibility and Social Choice

Homotopy of maps

Let X and Y be topological spaces, and let $f, g : X \to Y$ be continuous maps. A *homotopy* from f to g is a continuous map $H : X \times [0, 1] \to Y$ with $H(x, 0) = f(x)$ and $H(x, 1) = g(x)$ for all $x \in X$. We say f and g are *homotopic* (written $f \simeq g$).

Contractible spaces

A topological space X is *contractible* if the identity map $\mathrm{id}_X : X \to X$ is homotopic to a constant map: there exists a point $x_0 \in X$ and a continuous $H : X \times [0, 1] \to X$ with $H(x, 0) = x$ and $H(x, 1) = x_0$ for all x. Informally, the entire space can be continuously squeezed to a single point without leaving X.

Contractible spaces have rich topological structure:

- Every convex subset of $\mathbb{R}^n$ is contractible: the homotopy $H(x, t) = (1 - t)x + tx_0$ squeezes everything toward any fixed point x_0.
- Any space diffeomorphic to $\mathbb{R}^n$ is contractible (since $\mathbb{R}^n$ is convex).

- Contractibility implies path-connectedness: any point x can be connected to x_0 by the path $t \mapsto H(x, t)$.
- Contractibility implies simple connectedness: any loop in X can be continuously contracted to a point.

Chichilnisky's theorem

Arrow (1963) established that no social welfare function can simultaneously satisfy Unanimity (P), Independence of Irrelevant Alternatives (IIA), and Nondictatorship (ND) when there are three or more alternatives and two or more agents.

Chichilnisky (1982) reformulates social choice in the smooth setting and identifies the topological source of the obstruction. IIA requires the collective ranking of any two alternatives to depend only on individual rankings of those two alternatives—a pairwise separation condition. In the smooth setting, this forces the aggregation to be a *fiber-preserving map* between preference fiber bundles over the domain. On spaces with non-trivial topology (non-vanishing first cohomology), fiber-preserving maps that respect unanimity are necessarily dictatorial: the topology of the preference manifold forces the aggregation to track a single agent's preferences globally.

Chichilnisky replaces IIA with a topological continuity condition: small changes in individual preferences produce small changes in the collective preference. Her key theorem is:

> A continuous, anonymous social choice function on a smooth domain X that respects unanimity exists if and only if X is contractible.

Contractibility eliminates the topological obstructions that make IIA fatal: on a contractible domain, the preference fiber bundle is trivial and continuous aggregations satisfying unanimity need not be dictatorial.

How this appears in the book

The equilibrium manifold E (Corollary 3.12). Since $E \cong \mathbb{R}^{IL}$ and $\mathbb{R}^{IL}$ is convex (hence contractible), E is contractible. The explicit homotopy squeezes E to a point by pulling each $e \in E$ along the straight line in $\mathbb{R}^{IL}$ toward the origin via Φ, then pushing back via Φ^{-1}.

The path space $\mathcal{A}(e_0)$ (Proposition 8.6). The *path-slowing homotopy* is

$$H(s,\gamma)(t) = \gamma(st), \qquad s,t \in [0,1].$$

At $s = 1$, $H(1,\gamma) = \gamma$ (the identity on $\mathcal{A}(e_0)$). At $s = 0$, $H(0,\gamma)(t) = \gamma(0) = e_0$ for all t: every path collapses to the constant path at e_0. For intermediate s, the map $t \mapsto \gamma(st)$ is a re-parameterization of γ that traverses only the initial segment $\gamma([0,s])$ by time $t = 1$. Since re-parameterizations of smooth paths from e_0 remain in $\mathcal{A}(e_0)$, the homotopy stays within the domain. Hence $\mathcal{A}(e_0)$ is contractible.

Chichilnisky's theorem then guarantees a continuous, anonymous social choice function over institutional alternatives (Corollary 8.7). The contrast with Arrow is precise: Arrow's IIA imposes pairwise separation—a local condition on the aggregation that, in the smooth setting, requires fiber-preservation and leads to dictatorship on spaces with non-trivial topology. Chichilnisky's continuity condition imposes no pairwise separation. On a contractible domain, the weaker condition is compatible with fairness. The contractibility of $\mathcal{A}(e_0)$ rests on the path-slowing homotopy, which in turn requires only that E be path-connected. The full strength of Theorem 8.5—the smooth version guaranteeing that differentiable perturbations of individual preferences produce differentiable changes in the collective preference—additionally requires E to carry a smooth structure, which the diffeomorphism $E \cong \mathbb{R}^{IL}$ supplies.

Bibliography

Arrow, Kenneth J. 1963. *Social Choice and Individual Values*. 2nd ed. New Haven, CT: Yale University Press.

Balasko, Yves. 2009. *The Equilibrium Manifold: Postmodern Developments in the Theory of General Economic Equilibrium*. Cambridge, MA: The MIT Press.

Balasko, Yves. 2011. *General Equilibrium Theory of Value*. Princeton, NJ: Princeton University Press.

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

Carroll, Robert J. 2026. "The Shape of Possible States." Working paper, University of Illinois at Urbana-Champaign.
URL: *https://robertjcarroll.com/shape.pdf*

Chichilnisky, Graciela. 1982. "Social Aggregation Rules and Continuity." *Quarterly Journal of Economics* 97(2):337–352.

Collier, Paul and Anke Hoeffler. 2004. "Greed and Grievance in Civil War." *Oxford Economic Papers* 56(4):563–595.

Copeland, Dale C. 2015. *Economic Interdependence and War*. Princeton, NJ: Princeton University Press.

Debreu, Gerard. 1972. "Smooth Preferences." *Econometrica* 40(4):603–615.

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

Findlay, Ronald and Kevin H. O'Rourke. 2007. *Power and Plenty: Trade, War, and the World Economy in the Second Millennium*. Princeton, NJ: Princeton University Press.

Friedman, Milton. 1949. "The Marshallian Demand Curve." *Journal of Political Economy* 57(6):463–495.

Garfinkel, Michelle R., Stergios Skaperdas and Constantinos Syropoulos. 2015. "Trade and Insecure Resources." *Journal of International Economics* 95(1):98–114.

Gilpin, Robert. 1981. *War and Change in World Politics*. Cambridge, UK: Cambridge University Press.

Goldstein, Joshua S. 1988. *Long Cycles: Prosperity and War in the Modern Age*. New Haven, CT: Yale University Press.

Hirshleifer, Jack. 2001. *The Dark Side of the Force: Economic Foundations of Conflict Theory*. Cambridge, UK: Cambridge University Press.

Humphreys, Macartan. 2005. "Natural Resources, Conflict, and Conflict Resolution: Uncovering the Mechanisms." *Journal of Conflict Resolution* 49(4):508–537.

Keohane, Robert O. 1984. *After Hegemony: Cooperation and Discord in the World Political Economy*. Princeton, NJ: Princeton University Press.

Koremenos, Barbara, Charles Lipson and Duncan Snidal. 2001. "The Rational Design of International Institutions." *International Organization* 55(4):761–799.

Le Billon, Philippe. 2013. *Wars of Plunder: Conflicts, Profits, and the Politics of Resources*. Oxford, UK: Oxford University Press.

Mas-Colell, Andreu. 1985. *The Theory of General Economic Equilibrium: A Differentiable Approach*. Cambridge, UK: Cambridge University Press.

Mas-Colell, Andreu, Michael D. Whinston and Jerry R. Green. 1995. *Microeconomic Theory*. New York: Oxford University Press.

Milnor, John W. 1965. *Topology from the Differentiable Viewpoint*. Charlottesville: University Press of Virginia.

Modelski, George. 1987. *Long Cycles in World Politics*. Seattle, WA: University of Washington Press.

Moore, Jr., Barrington. 1958. *Political Power and Social Theory: Seven Studies*. Cambridge, MA: Harvard University Press.

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

Ruggie, John Gerard. 1993. "Territoriality and Beyond: Problematizing Modernity in International Relations." *International Organization* 47(1):139–174.

Spivak, Michael. 1965. *Calculus on Manifolds*. New York: W.A. Benjamin.

Tian, Nan, Aude Fleurant, Alexandra Kuimova, Pieter D. Wezeman and Siemon T. Wezeman. 2018. Trends in World Military Expenditure, 2017. Sipri fact sheet Stockholm International Peace Research Institute.

Waltz, Kenneth N. 1979. *Theory of International Politics*. Boston, MA: McGraw Hill.