A Theory of Force

Robert J. Carroll

University of Illinois at Urbana-Champaign

Draft — March 24, 2026

Contents

Chapter 0

What is Force?

The true hero, the true subject, the center of the Iliad is force. Force employed by man, force that enslaves man, force before which man's flesh shrinks away. In this work, at all times, the human spirit is shown as modified by its relations with force, as swept away, blinded, by the very force it imagined it could handle, as deformed by the weight of the force it submits to. For those dreamers who considered that force, thanks to progress, would soon be a thing of the past, the Iliad could appear as an historical document; for others, whose powers of recognition are more acute and who perceive force, today as yesterday, at the very center of human history, the Iliad is the purest and the loveliest of mirrors.

Simone Weil, "The Iliad, or the Poem of Force"

This book takes Weil's challenge seriously. Force stands at the very center of human history, and yet we have remarkably little clarity about what sort of thing it is. States possess "armed forces"[1] and engage in "uses of force;" they "project" force and "balance" it; they "apply" force and "threaten" it. More dynamically, they "mobilize," "stockpile," and "build up" their forces, not to mention "combining" or "allying" them. Whatever the word points to, it is salt in the waters of statecraft—perhaps of statehood itself. Scholars and policymakers alike speak in its language, and formal models assign it symbols and variables, but rarely is it specified as an object with properties of its own. Clausewitz (1832) defined war as "an act of force to compel our enemy to do our will;" this book asks the prior question of what force itself *is*.

[1] Weber (1978) defined the state as the entity claiming a monopoly on the legitimate use of physical force within a territory. Tilly (1985) reversed the arrow: "war made the state, and the state made war." This book takes the state as given and asks what the force it monopolizes *is*.

Simple questions often beget complicated answers, and "What is force?" is no exception. This is not just because of the subtlety of the concept but also because of the form of the question. Questions of the form "What is x?" are *ontological* questions: they ask what kind of thing x is—what structures it requires, what distinctions it supports, what it can and cannot do. We are not searching for an essence. We are looking for the structures without which reasoning about force becomes incoherent. Definitions offer clarity and useful intuitions, but they cannot trace out the contours of the class of things to which a concept refers. They treat only the symptoms of our ignorance.

Our answer will be constructive. Rather than defining force, we build it: we identify the features that any plausible force must exhibit, ask what mathematical structures those features require, and assemble them from the ground up. This chapter clears the ground for that construction. We begin by showing that formal models already carry ontological commitments about force—commitments they rarely make explicit and that cause real trouble. We then fix the scope of what we mean by force, introduce a companion for the journey, set criteria for what a successful construction must achieve, and outline the dual program of structure and forgetting that animates the chapters to come.

0.1 Ontological Commitment in Models of War

Every model announces an ontology, whether it admits as much or not. To write down a variable is already to say that something exists, that it can be measured, related, or combined in particular ways, and that those operations matter for the outcome of interest. To specify a payoff function is not just to set strategic logic; it is to stipulate what the world is made of and how its pieces can be moved.

Robert Powell (2017) reminds us that the distinguishing virtue of rationalist international relations lies in its "tight deductive links"—the fact that carefully specified assumptions yield tractable models whose conclusions follow cleanly from their premises. As he puts it, "a fundamentally important prerequisite to understanding the implications of any set of assumptions is a tight deductive link between assumptions and conclusions." This virtue is real: without deductive discipline, there is no clarity about what follows from what.

But alas, deductive tightness is not sufficient. A chain of reasoning can be impeccable internally while its anchoring concepts drift. What good is a tight deductive link if the symbol it carries—"force," say—means one thing in one application and another in the next? Logic can guarantee that implications follow from assumptions, but it cannot guarantee that the assumptions are written in a language that fits the world. In the syntactic view of scientific theories, the model is an interpretation of a theory; in the semantic view, it is a structure that represents a signature. In either case, the mathematics is only the image. The adequacy of that image depends on the mapping between symbol and world. When that mapping is loose—when many different and even incommensurable things can be pushed through the same variable—the semantic link slackens, and the entire deductive chain is left hanging in midair.

This is the terrain on which ontology matters. The problem is not that the rules of inference fail; it is that the language itself may be unfit for purpose. A constructive ontology secures the semantic link, rendering variables not only manipulable but meaningful. Only once we have categories built to capture the relevant features of force can the vaunted tight deductive links of rationalism do the work they promise.

This looseness applies, and perhaps applies *especially,* to formal models of force. Consider one of the simplest and most widely used tools in the rationalist repertoire: the contest model of militarization. Two players, $i \in \{1, 2\}$, simultaneously choose a militarization level $m_i \in \mathbb{R}_{\geqslant 0}$. Their payoffs are given by von Neumann–Morgenstern expected utility functions:

$$U_1(m_1, m_2) = \frac{\lambda m_1^\alpha}{\lambda m_1^\alpha + m_2^\alpha} \times (V - k(m_1 + m_2)), \tag{1}$$

$$U_2(m_1, m_2) = \frac{m_2^\alpha}{\lambda m_1^\alpha + m_2^\alpha} \times (V - k(m_1 + m_2)), \tag{2}$$

where $\lambda \in \mathbb{R}_{>0}$ captures the relative effectiveness of the forces, $\alpha \in (0, 1]$ captures the decisiveness of superior force, $V \in \mathbb{R}_{>0}$ captures the value of the prize, and $k \in (0, 1]$ captures the inverse-recuperability of militarization costs. The game has a unique Nash equilibrium in pure strategies.[2]

Here each player chooses a militarization level $m_i \in \mathbb{R}_{\geqslant 0}$, which then determines both the probability of prevailing and the costs of fighting. The mathematics is clean and the equilibrium unique.

Formally, m_i enters payoffs in two ways. First, it determines the probability of winning the contest, via a variant of the Tullock contest success function, $\frac{\lambda m_i^\alpha}{\lambda m_i^\alpha + m_j^\alpha}$. "Fighting harder" helps a player by raising the chance of securing

[2]The contest success function is a variant of Tullock's (1980). Equilibrium existence and uniqueness follow from standard arguments; see Skaperdas (1996) for a general treatment.

the prize—a territory, a concession, an advantage. Second, it determines the cost of fighting, $k(m_1 + m_2)$, where k captures the inverse-recuperability of militarization costs. "Fighting harder" hurts a player by making the conflict more expensive. In this sense, m_i *acts like force acts*: committing more troops or materiel both improves prospects of success and increases the costs of battle. But resemblance is not identity. The model treats m_i as if it were force, without ever specifying what force must be. Soldiers, war goods, spending, capacity, power—all of these, at different times, have been pushed through the same symbol. The deductive link is tight—the model works—but the semantic link drifts, and with it the comparability of results across applications.

Ronald Giere's philosophy of scientific models clarifies why this drift is so damaging. In his account, scientific representation is not a bare relation between language and world, but a pragmatic activity: *scientists use models to represent aspects of the world for specific purposes* (Giere, 2004). Models succeed not because they are intrinsically similar to reality—anything is similar to anything else in countless respects—but because scientists designate *which respects of similarity* are relevant, and to what *degree of fit*. It is only by specifying these respects and degrees that a symbol in a model can be said to represent something in the world at all. Seen through this lens, the problem with contest models is clear: the symbol m_i is used without agreement on what respect of force it is meant to capture. Sometimes it is soldiers, sometimes spending, sometimes capacity, sometimes power. Without explicit designation, the model's internal deductions may be tight, but its representational connection to the world is slack. Giere thus helps us see that semantic discipline is not an optional extra: it is what allows a model's symbols to hook onto reality in the first place.

This looseness presents itself in three ways.

Semantic consistency. Does m_i mean the same thing across applications of the contest model? The deductive form is the same, but the literature is rife with variation in what exactly is being armed. For some, m_i is "war effort," a deliberately generic resource that could be measured in francs, pounds, or battalions (Beviá and Corchón, 2010; Hirshleifer, 1991). For others, it is a "war good" or "military good" (Garfinkel, 1990; Powell, 1993), an output of economic spending that does not accumulate across periods. Elsewhere it is called "capacity" or simply "arms" (Meirowitz and Sartori, 2008; Skaperdas, 1992; Coe and Vaynman, 2020), a category elastic enough to include new technologies, strategies, or organizational innovations. Sometimes the variable is written as "arms levels" or even "force levels" (Fearon, 2018), while in other places the same functional role is played by what authors explicitly call "military power" (Hodler and Yektaş, 2012).

Each of these is intelligible on its own. But taken together, they pose a problem. The contest model has a tight deductive link from assumptions to conclusions, just as Powell celebrates. Yet the very object of deduction—m_i—slides from one meaning to another. The links are tight in form but slack in content. If different literatures push different objects through the same symbol, then the comparability of results is undermined, and the rationalist project fails to achieve the very clarity it seeks. Without semantic discipline, tight deductive links are not tight enough.

Structural fitness. The second dimension of looseness lies not in what m_i is taken to mean, but in the mathematical space in which it is placed. By convention, contest models set $m_i \in \mathbb{R}_{\geqslant 0}$. This choice is convenient: it permits smooth variation, aggregation, and calculus. But convenience is not ontology. Do we really believe that force admits all the operations that the real numbers do? That it can be continuously scaled—as is required for the stated interpretation of λ to make sense? That it can be commensurably aggregated across domains, and meaningfully added and subtracted?

Consider a mixture of 10,000 infantry, one nuclear submarine, and a cyber operations unit. What is the scalar sum of this force? If the answer depends on context, doctrine, or technological interaction, then additivity fails. In other settings, force is capped, discrete, or governed by thresholds, making continuous scaling questionable. The choice of $\mathbb{R}_{\geqslant 0}$ smuggles in strong assumptions about divisibility, commensurability, and substitutability. These may be justified in particular cases, but they cannot be left implicit.

Modelers often select the mathematical space that best serves tractability. Jackson and Morelli (2009), for example, toggle between discrete and continuous representations of militarization depending on the demands of the argument. This flexibility is not itself a flaw; sometimes discrete models clarify intuition before more elaborate continuous accounts. But it underscores the point: the choice of space is rarely driven by a theory of force. It is driven by the needs of the model.

Saunders Mac Lane's philosophy of mathematics sharpens the point. Mathematical structures, he argues, arise from basic human activities of reasoning: collecting, comparing, computing, arranging, measuring (Mac Lane, 1986, pp.

35–36)—indeed, Mac Lane goes so far as to call these "human cultural activities." Different activities can yield different formalizations, and no single space is privileged outside its context. A set, a vector, or a metric space is not a mirror of nature but a tool for making particular distinctions precise. Seen in this light, placing m_i in $\mathbb{R}_{\geqslant 0}$ is not a neutral move but a substantive one: it assumes that the reasoning needs of modeling force are best served by continuity, additivity, and scalar comparability. That may be right in some contexts, but it cannot be taken for granted. The ontology of force must tell us which reasoning needs we actually have, and only then which mathematical spaces can serve them.

Ontology reverses the order. Our choice of space must follow from a theory of force, not the other way around. Ideally, the construction we develop will port back to applied models like the contest game above, but only after the primitives of force have been secured. Only then can we construct something not only useful in practice but also reliable under scrutiny—semantically consistent across applications and structurally fit for the models we ask it to support.

Domain identification. A third problem is that the very same contest form is used to study domains far removed from war. Lobbyists competing for influence, firms competing for monopoly rents, litigants vying in court, and laboratories racing to patent a discovery have all been modeled with variants of the same contest success function. The mathematics is indistinguishable. What makes the contest game above a model of *force*, rather than a model of lobbying, rent-seeking, litigation, or innovation?

The answer cannot lie in the form itself, which is identical across fields. Nor can it lie in the deductive tightness of the model, which applies equally well to lobbying contests as to arms races. The difference lies only in the interpretation of the inputs and outputs. To call m_i "force" is to claim that the resources being committed are specifically resources of organized violence, and that the outcomes they determine are contests of coercion between states. Without such a designation, the contest model is a free-floating mathematical skeleton, equally at home in microeconomics or in international relations.

It should be said that this generality is a strength of the rent-seeking framework. Instrumentally, the fact that a single form can represent such disparate contests is precisely what has made it valuable: one tool relates lobbying to war, litigation to innovation, rent-seeking to arming. But the very breadth that makes the framework powerful also makes it fragile. Each new application requires us to ask how best to map the model to and from the domain at hand. Without explicit guidance, we are left with the danger that the same formalism is doing very different work under the same name.

This raises a problem of domain identification: what grounds our claim that a particular instantiation of the contest form is about force at all? A constructive ontology must do more than secure semantic consistency and structural fitness; it must also provide a principled basis for domain identification, so that we can say when a contest is a contest of force and not something else.

These concerns are not only philosophical; they also matter for empirical practice. Without knowing what kind of object force is, we do not even know which data to collect or how to interpret the results we obtain. If force is well represented by an n-dimensional vector of real numbers—each denoting a different variant of a "war good"—then data on the numbers of tanks, ships, and aircraft are relevant. If instead force is a scalar, then some aggregation procedure is required, and any such procedure will be theory-dependent in ways we must make explicit.

For all our present confusion, three commitments are clear. First, we should want the variable we construct to be *useful*—able, when meaningfully mapped to $\mathbb{R}_{\geqslant 0}$ or some other space, to behave as force ought to behave in models of militarization, bargaining, and war. Second, we should want it to be *reliable*—its mathematical space chosen because it follows from a theory of force, not simply from the demands of tractability. Third, we should want it to be *domain-secure*—anchored in a principled account of when a contest is in fact a contest of force, and not lobbying, litigation, or rent-seeking by another name. Only with usefulness, reliability, and principled domain identification secured can the deductive links of rationalism do the work they promise. The next section outlines what a proper answer would look like.

0.2 The Program

In his classic essay "On What There Is," Willard Van Orman Quine (1948) argues that a theory is committed to those entities to which its bound variables must refer for its claims to be true. What has come to be called *Quine's criterion of ontological commitment* is often summarized by the slogan: "to be is to be the value of a bound variable."[3] To answer "What is force?" is therefore to *construct*

[3] See relevant articles in the *Stanford Encyclopedia of Philosophy* (*e.g.*, Bricker, 2016; Hylton and Kemp, 2023). George Boolos (1984) extends this slightly to "to be is to be a value of a variable

a variable that takes all possible forces as values: values past, present, and future; values actual, possible, and impossible; values real, ideal, and imaginary. The variable must be structurally sound—it must capture the structures that any account of force requires. And it must be logically sound—it must come equipped with a grammar that allows meaningful reasoning about force.

The variable ranges over a set. Call it X for now—the set of all forces. We do not yet know what its elements look like, how many there are, or what structure the set carries. We only know that we want to compare its members.

The natural predicate is "is at least as forceful as": given two forces F and G, we should be able to say whether F is at least as forceful as G, or that the two cannot be distinguished. Write this $F \succsim G$. If $\succsim$ is complete (any two forces can be compared) and transitive (if $F \succsim G$ and $G \succsim H$, then $F \succsim H$), we have a complete preorder on X. The natural ambition is then to represent this ordering by a real-valued function: a map $m\colon X \to \mathbb{R}$ such that

$$F \succsim G \iff m(F) \geq m(G).$$

This is what anyone who assigns a number to a force is doing, whether they know it or not. When Trevor Dupuy (1979) compressed twenty-eight weapon types and seventy-three battlefield modifiers into a single Combat Power Potential, he was constructing a force function $m\colon X \to \mathbb{R}$. When a contest model assigns a militarization level $m_i \in \mathbb{R}_{\geqslant 0}$, it is doing the same thing: projecting a structured force onto a scalar. We call m a *force function*—not a utility function, because the predicate that generates it is not "is at least as good as" but "is at least as forceful as."[4]

But can m always be found? Not in general. Completeness and transitivity of $\succsim$ are necessary but not sufficient. On uncountable sets, complete preorders can

(or to be some values of some variables)." For Carnap's (1950) complementary view—that the choice of a linguistic framework is a practical decision, not a factual one—see the same *SEP* articles.

[4]The distinction matters. A utility function represents a preference ordering over outcomes: the agent prefers more of u to less. A force function represents a *forcefulness* ordering over force objects: one force is at least as forceful as another in the sense that it is at least as capable of doing the things that forces do—coercing, defending, projecting, deterring. Doctrine is to the force scale as preference is to the utility scale: it is the substantive judgment that determines which forces rank above which. The two scales need not agree—a doctrine may select a smaller, cheaper force over a larger one, just as a preference may select a less forceful option because it costs less to maintain. Doctrine is the subject of Chapter 4; forcefulness is the subject of this chapter and the next three.

fail to admit any real-valued representation at all—lexicographic preferences on $\mathbb{R}^2$ are the canonical example.[5] Whether m exists depends on the structure of X itself—how large it is, what topology it carries, whether the ordering $\succsim$ is well-behaved with respect to that topology. We do not yet know any of this.

The reals, then, are not the natural home of force. They are the natural home of *force as measured*—force projected through a lens of forcefulness, or cost, or probability of success. The structured force F lives in X; the scalar $m(F)$ lives on a line. Whether that projection is possible, and whether it is any good, depends entirely on what X turns out to be.

This gives us the program.

1. *Force is recursive* (Chapter 1). We build the set X from the ground up: elements bond into molecules, molecules gather into configurations, configurations are organized under command hierarchies. At each level, we discover how large the set is and what structure it carries.

2. *Force is dynamic* (Chapter 2). Forces change—molecules are built and broken, organizations are restructured and reformed—and the grammar of those changes is double-pushout graph rewriting. The rewrite rules compose into a *category of force*, **Force**: a symmetric monoidal category whose objects are structured forces and whose morphisms are valid transformation sequences. This is a resource theory in the sense of Coecke, Fritz and Spekkens (2016): the convertibility preorder tells us which forces are reachable from which, resource monotones measure different dimensions of capability, and the Force-Maker's strategic position is encoded by the slice of the category at her current force.

3. *Force is costly* (Chapter 3). The Force-Maker assigns costs to transformations; those costs compose via a *quantale*—a structure that accommodates asymmetry (building up costs differently from tearing down) and non-additivity (the cost of A then B need not be the sum of the parts). The cost of the cheapest path between two forces is a *Lawvere metric* on the force space, and from that metric we construct a second-countable topology on X—the topological condition that Debreu's representation theorem requires.

[5] Define $(x_1, y_1) \succsim (x_2, y_2)$ iff $x_1 > x_2$, or $x_1 = x_2$ and $y_1 \geq y_2$. This is complete and transitive, but no function $m\colon \mathbb{R}^2 \to \mathbb{R}$ can represent it: between any two distinct equivalence classes there would need to be a distinct rational number, but the equivalence classes are uncountable.

4. *Force is subjective* (Chapter 4). The Force-Maker's ranking of force configurations is her *doctrine*—a preorder on the force space—and whether that ranking can be compressed into a continuous scalar m is a doctrinal question, not a mathematical one. Debreu's (1954) theorem gives the answer: a complete, continuous preorder on a second-countable space admits a continuous real-valued representation, unique up to strictly increasing transformation—which is the index number problem in military assessment, formalized. When completeness fails, no scalar suffices, but a countable family of continuous functions represents the partial doctrine (Richter, 1966; Peleg, 1970). We close by returning to the contest model with which we opened: the Debreu theorem makes it possible to say, precisely, what doctrine the variable m_i encodes, what its level sets look like, and what it is committed to ignoring. The projection from structured force to scalar is *earned*, not assumed; and the semantic link is no longer slack.

The rest of this chapter develops the tools we need before the construction can begin: a working stipulation that fixes what counts as force, a companion for the road, and the methodological commitments that guide how we build and when we forget.

0.3 What We Mean by Force

Definitions alone will not do the work we need; still, we must fix scope before we build. We therefore begin with a minimal stipulation, not an essence claim.

Working stipulation. By "force" we mean the organized means of *physical* coercion. This includes personnel (soldiers, police, gendarmes), platforms (ships, aircraft, armored vehicles), munitions, and enabling systems (logistics, ISR, C^2, cyber units) configured for coercive use and the infrastructures that sustain them (bases, depots, ports, airfields, supply chains). It covers both internal and external applications: police within borders, military across them. It is a decidedly material notion. When Weber speaks of a monopoly on the legitimate use of *physical* force, he does not mean social, economic, or political "force." We will not use the word in those extended senses here.[6]

[6] As the *Oxford English Dictionary* records, the relevant senses of "force" denote military strength, organized bodies of armed men (or police), and the physical power of coercion. These

Why not just say "power"? Because "power" is broader and contested. Already in Thucydides' *History* (432 B.C.E.), the canonical line about the origins of the Peloponnesian War is remembered as "the growth of Athenian power"—yet the Greek itself never mentions power.[7] From the beginning, then, "power" has stretched across material, organizational, and reputational dimensions. Power, as Susan Strange (1988) has argued, operates through at least four structures: security, production, finance, and knowledge. Kenneth Waltz's (1979) catalogue—population, territory, resources, economic capability, military strength, political stability, competence—makes clear that military force is but one component of a larger composite. John Mearsheimer (2001) pushes further toward materialism, effectively equating effective power with military forces and their comparison. Our focus here is deliberately narrower: we model the *security* structure's material core. Nothing in what follows denies that other structures shape or condition force; indeed later chapters will link the construction to production, finance, and knowledge inputs. But the object we construct is *force*, not power writ large.[8]

Exclusions and edge cases. We exclude purely social, economic, or diplomatic leverage (sanctions, audience costs, status) except insofar as they alter the organization, availability, or effectiveness of physical force. We also exclude metaphoric uses ("the force of argument"), and we treat ambiguous instruments (*e.g.*, dual-use cyber capabilities) under force only when configured for coercive physical effects or for the command, control, and employment of coercive means.

---

ordinary-language anchors provide orientation, but they are not theoretical foundations.

[7] Thucydides I.23.6: τὴν μὲν γὰρ ἀληχεστάτην πρόφασιν, ἀφανεστάτην δὲ λόγῳ, τοὺς Ἀχηναίους ἡγο υμαι μεγάλους γιγνμένους καὶ φόβον παρέχοντας το ις Λακεδαιμονίοις ἀναγκάσαι ἐς τὸ πολεμε ιν. Both forms mean, literally, "the Athenians becoming great" or "the growth of the Athenians." Notably, Thucydides does not use the standard Greek terms for "power"—δύναμις, ἰσχύς, or ῥώμη. The word "power" entered the tradition through Thomas Hobbes's 1629 English translation ("the growth of the Athenian power") and has persisted in nearly all subsequent versions. The canonical phrase "the growth of Athenian *power*" is thus at least partly a translator's interpolation.

[8] Dahl's (1957) canonical relational definition of power—A gets B to do what B would not otherwise do—is about the *effect* of power. Strange's structural definition is about its *sources*. This book is about the *material substrate* of one such source. For a recent taxonomy distinguishing compulsory, institutional, structural, and productive power, see Barnett and Duvall (2005); for an earlier radical critique, see Lukes (2005).

Police, prisons, and the choice of examples. Nothing in the construction that follows is specific to military force. A police department is recursively structured: it has lieutenants, sergeants, captains, precincts, and specialized units organized under a command hierarchy. It is dynamic: its organization changes over time. It is costly: maintaining it entails tradeoffs. And it is subjective: what counts as adequate policing depends on who is asking. The same is true of a prison guard force, a gendarmerie, a coast guard. Every principle developed in this book—recursive structure, dynamics, cost, subjectivity—applies with full generality to the organized means of physical coercion in any domain.

We spend most of our time on military force for practical rather than principled reasons: its history is deep and well-documented, its connections to the international relations literature are immediate, and its global scope lets us draw on examples from Shaka's *amabutho* to NATO carrier strike groups. But the reader should understand that the construction is not a theory of armies. It is a theory of force.

Why this stipulation matters for the program. Section 0.1 identified three adequacy conditions: semantic consistency, structural fitness, and domain identification. The stipulation above: *(i)* imposes semantic discipline (what m_i can legitimately denote); *(ii)* anticipates which mathematical structures may be appropriate (additivity, thresholds, partial orders) given the *material, organized* character of force; and *(iii)* anchors domain identification (when a contest is a contest of *force* rather than lobbying or rent-seeking). This is not the last word on meaning; it is the first constraint on the construction that follows.

0.4 Building and Forgetting

The construction must meet several criteria simultaneously: it should err wide (admitting more than it must, so that edge cases and future technologies are not prematurely excluded); it should be structurally rich (representing composition and organization, not just scalars); it should be portable (traveling across literatures without semantic drift); it should admit disciplined abstraction (projections and equivalence classes that preserve what matters and forget what does not); and it should connect to observable indicators (troops, tanks, budgets, bases, cyber units) so that theories can be tested.[9]

To meet these criteria, we adopt a dual stance of *building and forgetting*.

The constructive program begins with graphs.[10] Graphs provide a concrete, data-ready language for representing forces as structured objects: nodes for elements, edges for relationships, labels for properties. Starting here is deliberate: graphs anchor the ontology in a representation that could, in principle, be used to organize actual data on personnel, platforms, munitions, enablers, and infrastructures.

Yet graphs are a means rather than an end. Once the structure is in place, we will often "forget" the graphical details and reason at a higher level of abstraction. For many purposes, what matters is not the full diagram but the relationships it supports: equivalences, aggregations, thresholds, or comparative measures. The result is a dual-level ontology: rich enough to represent composition and organization in detail, yet flexible enough to project into simpler spaces when models require.

This is not merely practical. It signals something deeper, what we might call an *ontology of use*: rather than asking which things exist in the abstract, we ask which distinctions must exist for a theory or model to do its job. This is the stance adopted in many areas of formal modeling and in much of category theory: structure is not a mirror of metaphysics, but a scaffold for reasoning.[11]

---

[9] Two further criteria deserve mention. *Representational adequacy*: the construction must admit projections into spaces that standard models can use—Section 0.2 identified this as depending on the structure of the set, and Chapter 1 will settle the question. *Principled correspondence*: following Giere, the construction must designate which respects of similarity between model and world are relevant and to what degree.

[10] Much of Chapter 1 will be devoted to building up the graph-based theory of force objects. But for present purposes, it suffices to remind the reader that a graph is a collection of nodes (or vertices) connected by edges (or links). It therefore includes information both about things (the nodes) and about relationships between things (the edges). For example, a graph might include nodes for a soldier, his shield, and his sword, with edges indicating that the soldier *carries* the shield and the sword. Importantly, the edges are first-class objects in their own right, so that the same three nodes without the edges represent a fundamentally different force object. The formal construction is deferred to Chapter 1.

[11] Category theory is a branch of mathematics that studies the abstract properties of mathematical structures and the relationships between them. We will be using it at several points in what follows, though it is worth noting that the term "category" has a different lineage in philosophy. As Lowe observes, philosophy has its own tradition of category theory, one that "lies at the heart of ontology—but, properly understood, concerns categories conceived as categories of being, not, in Kantian style, as categories of thought. (There is, of course, also a branch of mathematics called 'category theory', but since ontology has the first claim on the term, I use it here without

We build more than we need so we can decide, carefully, what to leave behind. We learn well so that we can forget well. Forgetting is far from free—it requires judgment about which distinctions are preserved under projection and which mappings commute with the logics of our models—but by making the structure explicit, we put ourselves in a position to forget with discipline rather than with carelessness.

0.5 Devotion

We need a companion for the road. In his *Notes on the Theory of Choice*, David M. Kreps (1988) orients decision-theoretic reasoning around a character, the famous "trade-off talking rational economic person," or TOTREP. We do something similar. Our *Force-Maker* is the one who must act in a world defined by force: she classifies, she combines, she organizes, she decides. Unlike TOTREP, however, she is not a type. She is not representative of a class of agents. She is a dweller in the world that the model tries to describe. She theorizes, but does not do so formally. She decides, but often under duress or under constraints not of her own making.

She is our muse. Whatever we discover in her company is something she already knows; we are not making her smarter but simply trying to keep up.[12] Her presence reminds us that theory is not for our amusement alone. The problems we analyze are hers, and she will have to live with whatever clarity or confusion our models produce.

The Force-Maker will grow more prominent as subjectivity enters—when costs must be calibrated, when doctrine must be assessed, when the theory must confront the question of what she *wants*. For now, she watches us build. At each level of the construction, we will pause to ask what she can see and what she can distinguish—her *scope* and her *discrimination*. These are not formal parameters but a recurring question: how does the world look from where she stands?

---

apology to the mathematicians concerned)" (Lowe, 2006, p. 5). In this manuscript the line between the two traditions will blur; the construction will make clear that this is not accidental.

[12] This posture of humility echoes John Muth's (1961) rational expectations: the assumption is not that agents are omniscient but that they can do no worse than the model that describes them. As Deirdre McCloskey (1985, p. 53) puts it, "Muth's notion was that the professors, even if correct in their model of man, could do no better in predicting than could the hog farmer or steelmaker." The notion is one of intellectual modesty. Her decision-theoretic ancestry runs through Savage's (1954) axiomatization of subjective expected utility: she has preferences, she acts as if she holds beliefs, and the two cohere.

Why devote ourselves to a muse at all? Plato tells us that the gods, pitying our labor, gave us festivals and the Muses together so that we might return upright: in his account, leisure and the muse are a joint gift. Leisure invites divine company.[13] In devoting ourselves to the Force-Maker, we are not fleeing force and violence. We are turning toward them, as rigorously and humanely as we can—and we are inviting that divine company.

One of our goals, then, is to record when the company has arrived. Clausewitz devotes a famous chapter to *genius* in war—not genius as a vague compliment, but as the specific capacity to see what others cannot: the *coup d'oeil* that grasps a situation whole, the resolution that holds firm in the fog.[14] If the construction we build is any good, genius should be *visible in the structure.* We should be able to point to a place in the formalism and say: here is where the leap was made.

And indeed we can. When Shaka kaSenzangakhona replaced the throwing spear with the stabbing spear and retrained the *impi* around it, he created a new *molecule*—a new graph on the same elements, optimized for close-quarters combat. That is molecular genius: seeing a better graph and building it. When the Romans invented the legion and Napoleon the corps system, they created new ways of *nesting* units inside one another—organizational genius, operating at the level of command structure and recursive composition. When Andrew Marshall at the Office of Net Assessment asked how to compare force structures across adversaries—not by counting tanks but by understanding how systems of systems interact—he was asking the cost question: what is the *price* of military capability, and who is paying too much?[15] And when Dupuy tried to compress twenty-eight weapon types and seventy-three battlefield modifiers into a single number, he was performing an act of *doctrine*—the imposition of a forcefulness ordering on a structured object.

Each of these is a different kind of genius, and each appears at a different level of the construction:

- *Recursive genius* (Chapter 1): new molecules, new configurations, new

[13] Josef Pieper, *Leisure, the Basis of Culture* (1948). The Greek *scholē* gives us "school": learning is leisure. The Latin *licere* ("to be free") gives us "leisure" as license—not escape from work, but freedom *to* work in the most human way.

[14] *On War*, Book I, Chapter 3: "On Military Genius." Clausewitz writes that genius is "a harmonious combination of elements, in which one or the other ability may predominate, but none may be in conflict with the rest" (1832).

[15] Marshall's framework is most accessible through Krepinevich (1997) and Krepinevich and Watts (2015).

ways of composing force.

- *Dynamic genius* (Chapter 2): the double-pushout method gives us a microfoundation for how force changes, and genius appears in three forms—*match-finding* (seeing that a known transformation applies to the current situation), *scheme-invention* (designing a transformation no one has written down), and *preparedness* (judging which transformations the force should be able to execute without the leader present).
- *Cost genius* (Chapter 3): finding cheaper paths to the same capability, or recognizing that an adversary's path is ruinously expensive.
- *Subjective genius* (Chapter 4): the doctrines that compress structure into action—seeing which compression preserves what matters and which discards it.

If the formal apparatus does its job, genius is not a mystery. It is a specific structural innovation at a specific level of analysis, legible in the mathematics. The Force-Maker already knows this; we are trying to keep up.

Chapter 1

Force is Recursive

To define force—it is that x that turns anybody who is subjected to it into a thing. Exercised to the limit, it turns man into a thing in the most literal sense: it makes a corpse out of him.

Simone Weil, "The Iliad, or the Poem of Force"

When the Roman legions met the Macedonian phalanx at Cynoscephalae in 197 BC, the encounter tested two different theories of what military force is. The phalanx was the dominant formation of the Hellenistic world: a dense mass of pike-bearing infantry, trained to advance in strict alignment and nearly invulnerable from the front. The Roman legion was organized differently—not as a single bloc but as a hierarchy of *manipuli*, each small enough to maneuver semi-independently and commanded by officers with genuine tactical discretion. When the battle moved onto broken ground and the phalanx's flanks became exposed, individual Roman units wheeled to exploit the gap without waiting for orders from the center. The phalanx, designed to operate as a unified mass, could not respond in kind; it collapsed piecemeal. The armies were comparable in size. What was not comparable was their structure.

The pattern at Cynoscephalae—similar raw material, categorically different military capability, explained by differences in organization—recurs persistently across military history, and it poses a direct challenge to the most natural way of representing force. The natural representation is scalar: a number, a score, an aggregate index. Such scalars capture the output of national-accounts thinking about military power, and they are not without value; but they systematically obscure what matters. The capability of the Roman legion is not a property of any individual soldier; it is a property of the hierarchy that organizes soldiers into manipuli, manipuli into cohorts, and cohorts into legions—together with the allocation of those configurations to specific tasks and the command relations that permit semi-independent action at every level. None of that structure is captured by a scalar. A theory of force adequate to explaining military outcomes must represent it directly.

The construction proceeds through four recursive levels, each built from the last: *elements*, the atomic building blocks; *molecules*, structured combinations of elements; *configurations*, gatherings of molecules; and *structured forces*, organizations that command, support, and allocate configurations. At each level, the same pattern recurs: we build the objects, show that there are countably many of them, and pause to ask what the Force-Maker can see. By the chapter's end, we will have constructed a variable—in the Quinean sense of Chapter 0—comprehensive enough to accommodate any force, past, present, future, or imagined.

1.1 Elements

Figure 1.1: *The atomic level: elements are the indivisible building blocks of force.*

The Cynoscephalae argument shows that structure matters; the first task is to say what force is structured *out of*. A state's force is not a blob but a collection of distinct, nameable pieces: tanks, ships, planes, soldiers, each a little piece of force in its own right. Force is atomic in the sense that it is made up of many distinct pieces, which—for lack of a better term—we call *elements.* For all its science envy, the metaphor purchases a great deal of intuition:

1. *Elementariness.* Elements come in sorts, each with its own properties. Just as hydrogen and helium are different sorts of physical atoms, so too are tanks and planes different sorts of military atoms.

2. *Familiality.* These elements may be organized into families. Helium and neon are both noble gases; the Dreadnought and Iowa battleship classes are both battleships. Being constructed by different states in different eras, they might plausibly be different force elements; nevertheless, their battleshippiness is a commonality.

3. *Configurability.* Elements can be arranged in different ways. A tank might arrive alone, or with air support, or as part of an armored division. The properties of the element are the same, but the properties of the arrangement are different.

The theory begins with a set of force elements.

1.1 Primitive (Elements of Force)

There is a nonempty, countable index set L enumerating the elements of force.

The primitive L is a set of force elements, each of which is a sort of thing that makes up the state's force. Its elements include tanks, ships, planes, soldiers, and so on; we write these in fixed-width font and capital letters, as in `TANK`, `SHIP`, `PLANE`, and `SOLDIER`. Whatever different sorts of things convey force, they are enumerated in L.

What does it mean to say that force arrives in packets and that these packets come in a countable number of sorts? Consider the case where $|L| = 1$, meaning that there is only one sort of force element. As we know only that this element bears force, we can say nothing else about it and may only name it `FORCE`. This is a perfectly reasonable approach, and it is in essence what is asserted by formal models where militarization or mobilization is unidimensional. Unpacking L serves to examine all of the assumptions implicit in the unidimensional approach. This is not to say that the unidimensional approach is wrong; it is simply a choice, and one that should be made consciously. Lanchester's (1916) square law—the foundational attrition model—is the purest $|L| = 1$ case: force is a number, and numbers fight numbers. Lanchester began with one sort and stayed there. Dupuy (1979) went the other direction: his Theoretical Lethality Index distinguishes twenty-eight weapon types—from hand-to-hand combat through the javelin and the ordinary bow, up to the V-2 ballistic missile and the megaton nuclear airburst—each assigned a numerical lethality score. This is $|L| = 28$: a real attempt at multidimensionality. But Dupuy then multiplied these scores by seventy-three modifying factors for terrain, weather, morale, leadership, logistics, and the rest, collapsing everything back into a single Combat Power Potential. Lanchester had $|L| = 1$ by design; Dupuy had $|L| = 28$ by observation and then compressed it to $|L| = 1$ by choice.[1] It may be the case that one would

[1] The seventy-three modifying factors are not elements of force but properties of the *battlefield*—closer to what Debreu (1954) would call states of the world: a tank in the Sinai is a different good from a tank in the Ardennes. We return to this distinction in Chapter 4, where the act of compressing structure into a scalar is formalized as a doctrine.

want to know about `LAND FORCE`, `SEA FORCE`, and `AIR FORCE`, or about `M1 ABRAMS`, `USS ARIZONA`, and `F-22 RAPTOR`s. Conversely, if $|L| = \aleph_0$, then there are infinitely many sorts of force elements, which leaves open the possibility that a new tool of war might yet be invented. In our interpretation, L is however granular it needs to be to encode the thinking of the Force-Maker meaningfully. And countability is not a simplifying assumption—it is a consequence of what force *is*. Force is the thing wielded by states, and states are bureaucracies. Tilly tells us this much. The Force-Maker does not describe the caliber of a rifle with an arbitrary real number; she says 5.56mm NATO. She classifies, she lists, she counts—because the institution she serves can only operate on things it can name.[2] Countability is baked into force by the nature of the entity that wields it.

Though L is a primitive, we retain wide latitude in how it is specified, used, and interpreted. The point is less to identify actual elements than to acknowledge that a state's force is made up of many different sorts of things. Here the natural familiality of force elements is a key feature.

1.2 Definition (Classification of Force Elements)

A classification scheme *is a partition of the set of force elements into equivalence classes called* families. *Given two classification schemes $\mathcal{L}_1$ and $\mathcal{L}_2$, we say that $\mathcal{L}_1$ is* finer *than $\mathcal{L}_2$ if every family in $\mathcal{L}_1$ is a subset of a family in $\mathcal{L}_2$.*

The complete enumeration of elements is a classification scheme, but so too is the division of elements into land, sea, and air forces. Different classification schemes answer different questions, and their refinement ordering allows us to compare them: one scheme is "more detailed" than another if and only if it is finer. The Composite Index of National Capability (Singer, Bremer and Stuckey, 1972), the dominant measure in empirical international relations for decades, is a specific classification scheme: it partitions national resources into six families (military expenditure, military personnel, iron and steel production, energy consumption, total population, urban population). It is a coarse partition; the construction asks what happens when you refine it.

[2] The formal point: every element of L that the Force-Maker can name is a finite string over a finite institutional vocabulary—a designation like `5.56mm NATO` or `M1A2 ABRAMS`. Any finite string over a finite alphabet is a natural number under binary coding, so the set of all such strings is countable. Uncountability would require the Force-Maker to distinguish a continuum of force elements, each requiring an irreducibly infinite specification. But she is a bureaucrat, not a deity: her lists are finite, her alphabet is finite, and therefore her index set is countable. The assumption $|L| \leq \aleph_0$ is not a modeling convenience; it is a theorem about the expressive capacity of the institutions that wield force.

The Force-Maker's scope. Which elements does the Force-Maker know about? Which does she treat as interchangeable? A Force-Maker with a narrow scope—one who knows of `TANK` and `PLANE` but not of `DRONE`—inhabits a smaller world than one who knows of all three. And a Force-Maker who cannot distinguish a `DESTROYER` from a `FRIGATE` uses a coarser classification than one who can. These two dimensions—what she can *see* and what she can *distinguish*—will recur at every level of the construction.

1.3 Definition (Epistemic Stance)

An epistemic stance *on a label set L is a pair $(S, \approx)$ where $S \subseteq L$ is the* scope *(which element types the Force-Maker recognizes) and $\approx$ is an equivalence relation on S (which types she treats as interchangeable). The Force-Maker with stance $(S, \approx)$ inhabits the restricted universe built from $S/\approx$—the quotient of her scope by her discrimination.*

The set of epistemic stances on L forms a lattice: broader scope and finer discrimination yield a higher stance, narrower scope and coarser discrimination yield a lower one.[3] The point is that the variable we are constructing may be seen through different eyes, and that this matters. Two Force-Makers with different stances on the same label set L will build different structured forces from the same raw material—not because the material differs, but because their windows onto it differ.

1.2 Molecules

Elements do not arrive in isolation. A hoplite without a shield is not a hoplite but a ῥιψάσπις (*ripsaspis*), a shield-caster—a term of derision.[4] A tank crew without a tank is an infantry squad with delusions of grandeur. The meaningful unit is not the element but the *molecule*: a structured combination of elements, where the structure records which elements are combined and how they relate to one another. John Keegan's (1976) insistence on reconstructing what soldiers

---

[3] The lattice structure is developed in full in Carroll (2026), where epistemic stances are shown to form a complete lattice under the natural ordering. Here, the definition suffices; the lattice-theoretic overhead is not needed until Chapter 4, where doctrine determines which stance the Force-Maker adopts.

[4] See the edited volume by Donald Kagan and Gregory F. Viggiano, *Men of Bronze* (2013), especially the chapter by Viggiano and Hans van Wees (2013).

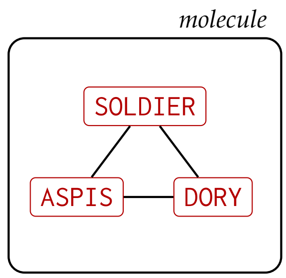

Figure 1.2: *The molecular level: elements (red, from the previous level) are connected by edges (black) into a molecule. The edges—the internal structure—are the new thing.*

actually carry, fire, and shelter behind—the "face" of battle—is the molecular perspective *avant la lettre*.[5]

1.4 Definition (Force Molecule)

A force molecule *is a connected labeled graph*

$$M = \left(n, E, \ell \colon \underline{n} \to L\right),$$

where $n \in \mathbb{N}$ *is the number of nodes,* $\underline{n} = \{1, \ldots, n\}$ *is the set of nodes,* $E \subseteq \binom{n}{2}$ *is a set of undirected edges, and* ℓ *is a labeling function assigning to each node an element of force. The graph* (n, E) *is required to be connected.*

A force molecule, then, is a little network of force elements. The nodes represent individual items—a soldier, a helmet, a spear—and the edges represent relationships between them: who carries what, who operates what, who supports whom.[6] The labeling function ℓ tells us what sort of element each node is.

[5] Clausewitz makes a similar point in Book V, Chapter 4 of *On War*: infantry, cavalry, and artillery are distinct arms with qualitatively different combat properties, and "the greatest strength is produced by a combination" (1832, p. 280).

[6] These relationships are plainly of different types—carrying, operating, supporting—and one could formalize this by introducing a set of *bond types* B and a labeling function $\beta \colon E \to B$ alongside the node labeling $\ell \colon \underline{n} \to L$. The construction goes through unchanged: countability is preserved (finitely many bond types on finitely many edges), configurations still decompose uniquely (Lemma 1.12), and the free commutative monoid structure is unaffected. We omit bond types here because none of the results that follow depend on them, but the reader who wants a richer graph loses nothing by adding them. The same tolerance extends upward: command edges (Section 1.4) could carry types distinguishing operational from administrative control, and allocation triples could be typed by mission. At each level the formalism is ready for the refinement; we simply do not require it yet.

Connectivity ensures that a molecule is a single, self-contained unit rather than a disjoint collection of parts.

1.5 Remark (Isomorphism of Molecules)

Two force molecules $M_1 = (n, E_1, \ell_1)$ *and* $M_2 = (m, E_2, \ell_2)$ *are* isomorphic *if there is a bijection* $\varphi\colon \underline{n} \to \underline{m}$ *such that* $(i, j) \in E_1$ *if and only if* $(\varphi(i), \varphi(j)) \in E_2$ *and* $\ell_1(i) = \ell_2(\varphi(i))$ *for all* $i \in \underline{n}$. *We write this* $M_1 \cong M_2$. *Isomorphic molecules share all structural properties; we treat them as the same.*

To give an example, consider the hoplite molecule in Figure 1.3.

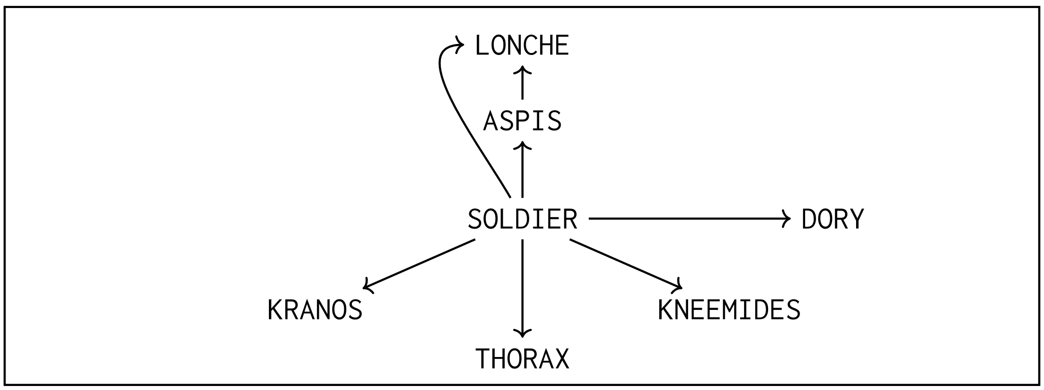

Figure 1.3: *A hoplite molecule, dory/lonche variant. The hoplite (*`SOLDIER`*) is connected to each piece of equipment; the* `ASPIS` *(shield) is additionally connected to the* `LONCHE` *(spear point), as the spear was often braced against the shield.*

The hoplite was the standard infantry soldier of ancient Greece. Its most characteristic element was a heavy wooden shield, the ἀσπίς (*aspis*), often adorned with an identifying device. The panoply also included a helmet (κράνος, *kranos*), a breastplate (θώραξ, *thorax*), greaves (κνημῖδες, *kneemides*), and a spear (δόρυ, *dory*) whose point (λόγχη, *lonche*) was often braced against the shield. Absent a panoply, a hoplite was simply not a hoplite; the molecule captures this: the `SOLDIER` node is connected to each piece of equipment, and the graph is connected as required.

The graph-theoretic approach earns its keep here. A molecule is more than a bag of elements. The edges encode relationships that matter: removing the `ASPIS` from the hoplite molecule does not merely subtract a shield—it severs the connections that make the panoply a functioning whole, including the ASPIS–LONCHE connection. This is generativity in miniature: the whole is more than the sum of its parts, and the graph is where the "more" lives.

Not all molecular innovation occurs on the battlefield. Sometimes the genius is in the design. Shaka kaSenzangakhona, founder of the Zulu Kingdom, transformed his force's combat performance by redesigning its most essential molecule, the IMPI.[7] He replaced the traditional ASSEGAI, a long throwing spear, with the IKLWA, a short stabbing spear; he enlarged the standard shield; and he ordered all warriors to fight barefoot. These changes were not cosmetic—they were a new molecule, a new graph on the same (or similar) elements, optimized for close-quarters combat. In the language we are developing, Shaka's genius operated at the molecular level: he saw a better graph and built it.

1.6 Primitive (The Set of Force Molecules)

The set of all force molecules on L, up to isomorphism, is denoted $\mathbb{M}_L$.

1.7 Lemma (Countability of $\mathbb{M}_L$)

$\mathbb{M}_L$ *is countable.*

Proof. For each $n \in \mathbb{N}$, the set of connected graphs on $\underline{n}$ is finite (there are at most $2^{\binom{n}{2}}$ graphs on n nodes, of which finitely many are connected). For each such graph, the set of labelings $\ell \colon \underline{n} \to L$ is countable (L is countable, and there are $|L|^n$ labelings, a countable number since L is countable and n is finite). Taking isomorphism classes does not increase the count. Thus $\mathbb{M}_L = \bigcup_{n=1}^{\infty} \mathbb{M}_L^{(n)}$ is a countable union of countable sets, and therefore countable. ■

The Force-Maker's scope, again. The epistemic stance of Definition 1.3 lifts from atoms to molecules. Which molecules can the Force-Maker conceive? The hoplite molecule is obvious to anyone who has seen a hoplite, but the IMPI molecule was obvious only to Shaka. Parker's (1988) thesis that European military dominance

[7] For the historical detail, see Morris (1965) and Laband (1995). Rosen (1991, pp. 19–22) discusses Shaka's reforms as a paradigm case of peacetime military innovation driven by a single innovating officer rather than by external threat—an observation that, in our language, reduces to the question of who has the authority to rewrite the molecule.

rested on molecular innovations—the *trace italienne*, the broadside ship, the musket-equipped infantry line—is a claim about this kind of molecular vision; McNeill (1982) traces the same co-evolution of military technology and social organization across a longer arc. A Force-Maker with narrow molecular scope inhabits a sparser $\mathbb{M}_L$—not because the molecules do not exist, but because she cannot see them. A Force-Maker who cannot distinguish two molecules operates with a coarser equivalence than isomorphism—her discrimination merges what a finer stance would separate.

1.3 Configurations

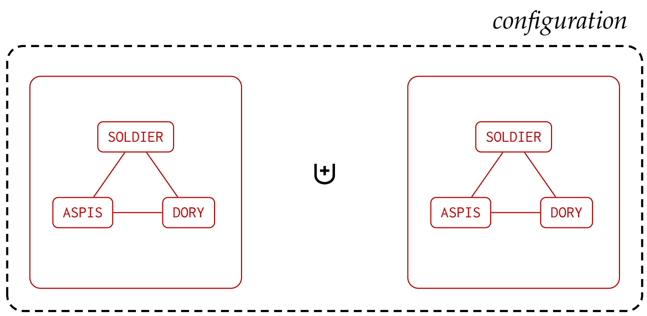

Figure 1.4: *The configurational level: molecules (red, from the previous level) are gathered via ⊎ (black). The gathering—co-presence without coordination—is the new thing.*

Having established that force arrives in molecules, we must now account for the fact that the Force-Maker does not deploy a single molecule at a time. She gathers molecules together: two hoplites, a trireme full of marines, an armored division's worth of tanks and support vehicles. The first task is simply to *mise en place*—to gather the necessary ingredients.

1.8 Primitive (Force Configuration)

There is a binary operation $\uplus$ *representing the* configuration *of two force molecules. It has the following properties:*

1. Unitality: *letting* $\mathbb{0}_{\mathbb{M}_L} := (0, \varnothing, \varnothing)$ *be the* void force, *we have*

$$M \uplus \mathbb{0}_{\mathbb{M}_L} = M = \mathbb{0}_{\mathbb{M}_L} \uplus M \quad \textit{for all } M \in \mathbb{M}_L;$$

2. Associativity: *for all* $M_1, M_2, M_3 \in \mathbb{M}_L$, *we have*[8]

$$(M_1 \uplus M_2) \uplus M_3 = M_1 \uplus (M_2 \uplus M_3)\,;\ \textit{and}$$

3. Commutativity: *for all* $M_1, M_2 \in \mathbb{M}_L$, *we have*

$$M_1 \uplus M_2 \cong M_2 \uplus M_1.$$

Configuration is gathering without transformation. The operation $\uplus$ places molecules side by side: they are co-present, not yet coordinated or combined in any deeper sense. The void force $\mathbb{0}_{\mathbb{M}_L}$ is the identity element—the Force-Maker's empty table before she begins. Associativity says the order of gathering does not matter. Commutativity (up to isomorphism) says that the arrangement within the gathering does not yet matter either—that is for organization to address.[9]

Crucially, configuration does *not* involve conversion: taking one configuration and turning it into another is a transformation, and transformations are the subject of Chapter 2. Here, the Force-Maker simply gathers her ingredients. *Mise en place.*

1.9 Primitive (The Free Commutative Monoid of Configurations)

The set of all force configurations on L, *generated by taking all finite* $\uplus$*-products of molecules in* $\mathbb{M}_L$, *is denoted* $\mathbb{M}_L^\star$. *Under* $\uplus$, *the pair* $(\mathbb{M}_L^\star, \uplus)$ *is the free commutative monoid on* $\mathbb{M}_L$.

Every configuration is a finite multiset of molecules: a precise inventory of how many of each molecule the Force-Maker has gathered. The word "free" means that no non-trivial relations hold among the generators—two configurations are the same if and only if they contain the same molecules in the same multiplicities.

1.10 Definition (Graph Union)

The graph-theoretic realization of $\uplus$ *is the* graph union*: given two molecules* $M_1 = (n_1, E_1, \ell_1)$ *and* $M_2 = (n_2, E_2, \ell_2)$, *their configuration is*

$$M_1 \uplus M_2 = \big(n_1 + n_2,\ E_1 \cup E_2',\ \ell_1 \cup \ell_2'\big),$$

---

[8] Associativity allows us to write $M_1 \uplus M_2 \uplus M_3$ without ambiguity, and indeed to write $\biguplus_{i=1}^{n} M_i \coloneqq M_1 \uplus \cdots \uplus M_n$.

[9] Commutativity holds up to isomorphism, not strict equality, because the graph-theoretic realization of $\uplus$ introduces new node labels that depend on ordering. Substantively, two hoplites gathered left-then-right are isomorphic to two hoplites gathered right-then-left; the molecules are the same and the gathering is the same.

where $E_2' = \{(i+n_1, j+n_1) : (i,j) \in E_2\}$ *and* $\ell_2'(i+n_1) = \ell_2(i)$ *for all* $i \in \underline{n_2}$. *In other words,* $M_1 \uplus M_2$ *is the disjoint union of the two graphs, with node indices shifted to avoid collision.*

Note that $M_1 \uplus M_2$ is *not* required to be connected. This is by design: a configuration is typically a disconnected graph, each connected component being a molecule. The molecules do not interact within a configuration; they merely coexist. Most quantitative security studies work with configurations alone. Epstein's (1990) attrition-exchange models treat forces as aggregations of units—configurations in our sense—and show how attrition rates depend on the composition of the multiset. The offense-defense balance literature (Glaser and Kaufmann, 1998) treats force as aggregate ratios—configurations without organizational structure—and the persistent difficulty of defining the balance precisely illustrates the insufficiency of working with the monoid alone.

1.11 Lemma (Countability of $\mathbb{M}_L^\star$)
$\mathbb{M}_L^\star$ *is countably infinite.*

Proof. A configuration is a finite multiset over $\mathbb{M}_L$. Since $\mathbb{M}_L$ is countable (Lemma 1.7), the set of all finite multisets over $\mathbb{M}_L$ is countable. It is infinite because, for any molecule M, the configurations $M, M \uplus M, M \uplus M \uplus M, \ldots$ are all distinct. ■

Figure 1.5 shows the configuration of two hoplite molecules.
The two hoplites are co-present but not coordinated. There are no edges between the two molecules—they simply occupy the same configuration. Given a configuration, we can always decompose it into its constituent molecules by identifying connected components.

1.12 Lemma (Deconfiguration)
For every configuration $C \in \mathbb{M}_L^\star$*, there exists a unique multiset of molecules* $\{M_1, \ldots, M_k\}$ *such that* $C = M_1 \uplus \cdots \uplus M_k$ *and each* M_i *is connected.*

Proof. Write $C = (\overline{n}, \overline{E}, \overline{\ell})$. We construct the decomposition by finding the maximal connected subgraphs of C. Initialize VISITED $\leftarrow \varnothing$ and COMPONENTS $\leftarrow \varnothing$. While VISITED $\neq \{1, \ldots, \overline{n}\}$, select a vertex $v \notin$ VISITED and perform a

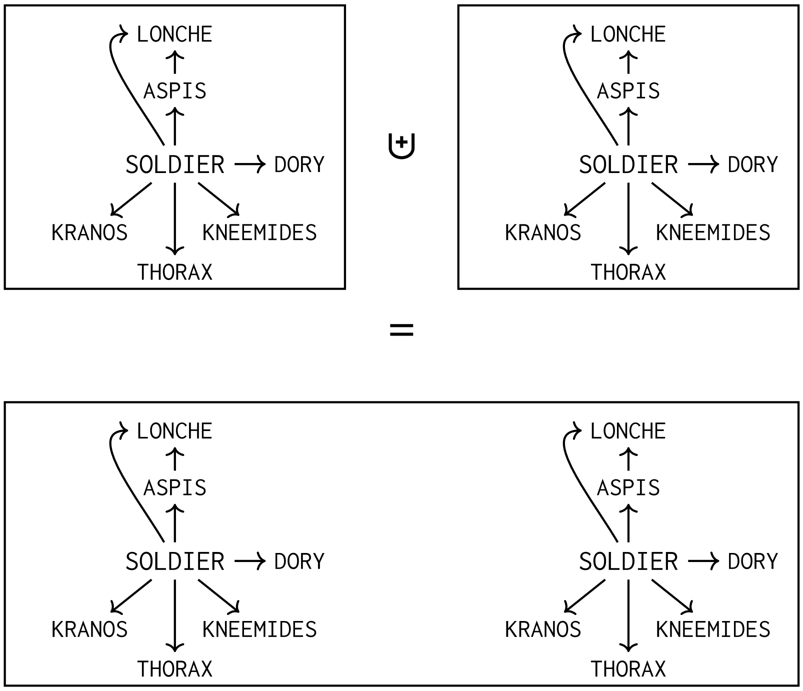

Figure 1.5: *The configuration* $\uplus$ *illustrated.* Top: *two hoplite molecules in separate boxes.* Bottom: *their graph union—the same two molecules now coexist in a single configuration. No edges connect the two components; they are co-present but do not interact.*

depth-first traversal: push v onto a stack; while the stack is nonempty, pop a vertex u, mark it visited, add it to the current component, and push all unvisited neighbors of u onto the stack. When the stack empties, record the current component and repeat.

The vertex set being finite, this terminates. Each recorded component $\widehat{V}_i$ is connected (any two of its vertices are linked by the traversal path) and maximal (any neighbor of a vertex in $\widehat{V}_i$ was pushed onto the stack and therefore included). The components partition $\{1, \ldots, \overline{n}\}$. Restricting $\overline{E}$ and $\overline{\ell}$ to each $\widehat{V}_i$ and renumbering vertices yields a molecule $\widehat{M}_i = (\widehat{n}_i, \widehat{E}_i, \widehat{\ell}_i)$. By construction, $C = \widehat{M}_1 \uplus \cdots \uplus \widehat{M}_{\widehat{k}}$, and each $\widehat{M}_i$ is connected.

Uniqueness (up to isomorphism and reordering) follows from the uniqueness of the decomposition of a graph into maximal connected subgraphs. ∎

The *multiplicity vector* $\mu_C \colon \mathbb{M}_L \to \mathbb{N}_0$ records the outcome: $\mu_C(M)$ is the number of connected components of C isomorphic to M. Two configurations are the same if and only if their multiplicity vectors agree.

Deconfiguration matters because it is what an intelligence analyst does. Given a satellite image of an adversary's deployment, the analyst's first task is to identify the component molecules—a tank platoon here, an infantry squad there, an artillery battery in the rear—and count them. This is the multiplicity vector. It is also what an arms control treaty verifies: the INF Treaty did not count "force"; it counted specific types of missiles and launchers—molecules—and imposed ceilings on their multiplicities. Configuration builds up; deconfiguration takes apart. Both operations are lossless: the configuration can always be recovered from its multiplicity vector, and vice versa.

The Force-Maker's scope, once more. The epistemic stance (Definition 1.3) lifts again to configurations. Which configurations can the Force-Maker envision? She may know that she has two hoplites but not realize that there exists a configuration involving a `TRIREME` and a `MARINE` molecule. Scope limits what she can see; discrimination limits what she can tell apart. The construction itself is the same regardless; what changes is the Force-Maker's window onto it.

1.4 Structured Forces

Configuration alone is not enough. An armored division is not merely a bag of tank molecules and infantry molecules; it is an *organized* bag—a hierarchy of command, with subordination, coordination, and support relationships among units. Some of the most famous Force-Makers in history have innovated not at the molecular level but at the organizational level: Julius Cæsar's legions, Napoleon's corps system, Shaka's *amabutho* regiments, George C. Marshall's reorganization of the Army for global war. If the Force-Maker pushes a button, these innovations reflect improvements in the mechanism that takes button press to desired outcome.

Organization is not a separate assertion about force; it is part of the recursive structure. At the atomic level, we defined what force is made *of*. At the molecular level, we defined how elements combine. At the configurational level, we defined how molecules gather. Now, at the organizational level, we define how gatherings are *structured*—and the same recursive pattern will hold.

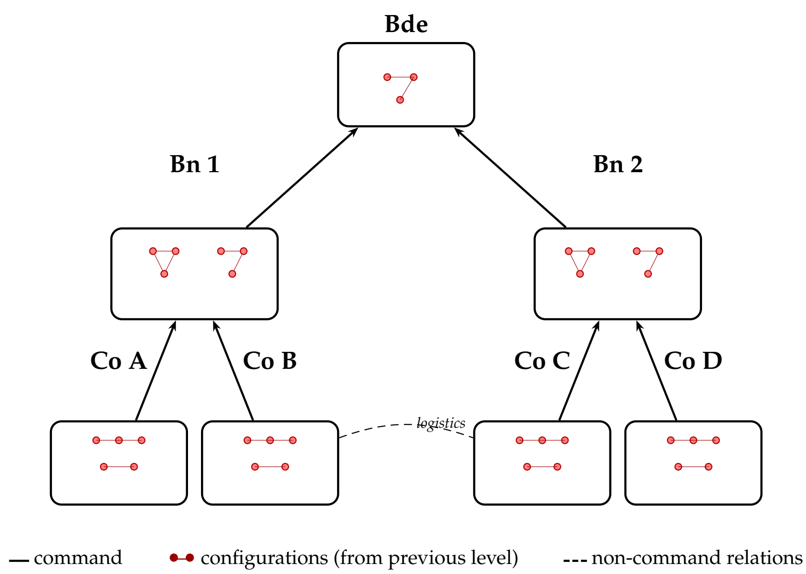

Figure 1.6: *The organizational level: each unit (black box) contains a configuration of molecules (red, from the previous level). Command (black arrows) and non-command relations (dashed) organize the units into a structured force. The organizational structure is the new thing; the configurations are what was already built.*

1.4.1 Command Structures

We begin with command, the primary organizational relationship that converts a gathering of molecules into a structured force.

1.13 Primitive (Command Structure)

A command structure *is a finite partial order* $(V, \leq_{cc})$, *where each* $v \in V$ *is a* force unit *and* $\leq_{cc}$ *is a* subordination relation *on* V: $v_1 \leq_{cc} v_2$ *means that unit* v_1 *is subordinate to unit* v_2. *The set of all command structures is*

$$\mathbf{ComStruct} \cong \bigcup_{n \in \mathbb{N}} \mathbf{Part}(\underline{n}),$$

where $\mathbf{Part}(\underline{n})$ *is the set of all partial orders on* $\underline{n} = \{1, \ldots, n\}$.

Being a partial order, the subordination relation $\leq_{cc}$ is reflexive, transitive, and antisymmetric. Reflexivity is trivial: every unit is (in some minimal sense) subordinate to itself, a reflection of whatever autonomy being called a “unit” confers. Transitivity ensures a coherent chain of command: if a company is subordinate to a battalion and the battalion to a brigade, then the company is subordinate to the brigade. Antisymmetry ensures that no two distinct units are mutually subordinate—no circular chains of command. These seem reasonable for organized force, where much of the point of being organized is to preclude violations of exactly these properties.[10] Van Creveld (1985) traces the historical evolution of this structure from Napoleon's personal direction—where the entire partial order terminated in a single brain—through Moltke's general staff, which distributed intelligence across the command structure, to the pathologies of over-centralization in Vietnam. His central finding is that the *shape* of the partial order matters as much as the equipment it commands: armies that pushed decision thresholds downward and allowed subordinates to exercise initiative—what the German tradition calls *Auftragstaktik*—have consistently outperformed those that centralized control, regardless of technological era.[11]

Note that a command structure need not be connected. A partial order on V may have multiple connected components: units that share no chain of command whatsoever. This is not a pathology—it is normal. Coalition forces are the clearest case: an American brigade and a British brigade operating in the same theater may have no subordination relation between them, only lateral coordination through a joint headquarters that is itself a separate component. Even within a single state's forces, distinct service branches may deploy to the same theater under separate chains of command—the U.S. Army and Marine Corps in the Pacific in 1944, for example, reported to different unified commanders. The partial order captures this: units in different components are organizationally independent. Connectedness, when it obtains, is an empirical feature of a particular force, not a structural requirement.

Figure 1.7 shows a Roman cohort as a command structure.
The fractal nature is visible: the cohort subdivides into maniples, which subdivide into centuries. Caesar commanded at least twelve legions during the

[10] The antisymmetry requirement encodes substantive motivations more than technical ones. Nearly all results hold for preordered sets (merely reflexive and transitive relations) without difficulty. But modern chains of command do not tolerate circularity, and so we keep antisymmetry.

[11] Van Creveld identifies two fundamental responses to the uncertainty inherent in command: increase the system's information-processing capacity, or restructure the organization to operate with less information. He argues that the second approach—self-contained units, decentralized authority, mission-type orders—has been superior across twenty-five centuries of Western military history (1985, pp. 268–274).

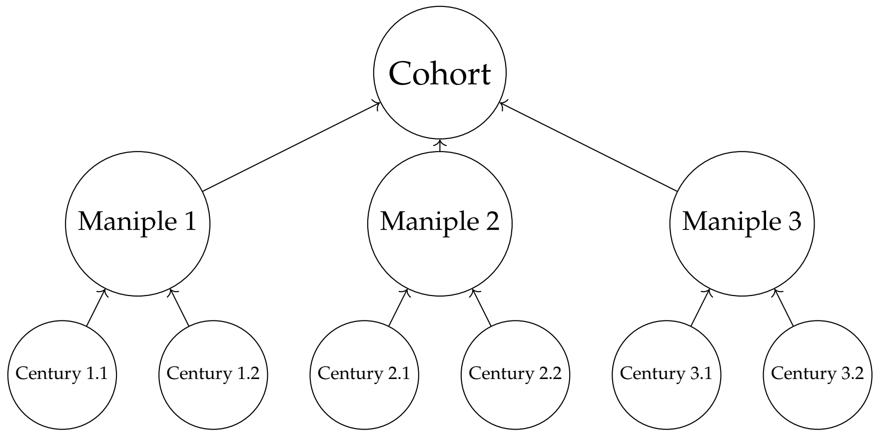

Figure 1.7: *Command structure of a Roman cohort: one cohort commanding three maniples, each commanding two centuries. Arrows point from subordinate to superior. The fractal pattern—units subdividing into smaller units—is the organizational counterpart of the recursive structure seen at lower levels.*

Gallic Wars, each containing ten such cohorts.[12] The organizational innovation—the *way* units are nested—is itself a kind of genius, distinct from molecular innovation but no less consequential.

1.4.2 Non-Command Relations

Command is not the only organizational relationship that matters. A medical company might support multiple combat units without being subordinate to any of them. An intelligence unit might serve the entire division. A logistics network might connect units across different branches of the command hierarchy. These non-command relationships are not subordination, but they are structure.

1.14 Definition (Non-Command Relations)

Given a command structure $(V, \preccurlyeq_{\mathrm{cc}})$*, a set of* non-command relations *is a finite collection* $R = \{H_1, \ldots, H_k\}$ *of hypergraphs on* V*. Each* $H_j = (V, \mathcal{E}_j)$ *is a hypergraph with* $\mathcal{E}_j \subseteq 2^V \setminus \{\varnothing\}$ *a collection of nonempty subsets of* V *called* hyperedges.

---

[12] See Judson (1888) for a detailed account of Caesar's army. Machiavelli (1521) admired the Roman legion so deeply that he modeled his Florentine reforms on it. Napoleon's corps system, too, echoed the hierarchical self-similarity of the legion (see, *e.g.*, Chandler, 1966).

Each hypergraph in R captures a different type of non-command relationship: one for logistics, one for intelligence, one for medical support, and so on. Hyperedges connect *sets* of units rather than pairs, because support relationships are often many-to-many: a single medical company might support three combat battalions, and that relationship is best captured as a single hyperedge {medical company, battalion 1, battalion 2, battalion 3} rather than three separate binary edges.

1.4.3 Force Allocation and the Complete Construction

A command structure with non-command relations tells us *how* force is organized, but not *what* is in it. For that, we need to assign configurations to units.

1.15 Definition (Force Allocation)

Given a command structure $(V, \preccurlyeq_{\text{cc}})$*, a* force allocation *is a function*

$$\text{assign}\colon V \to \mathbb{M}_L^{\star}$$

that equips each force unit $v \in V$ *with a force configuration* $\text{assign}(v) \in \mathbb{M}_L^{\star}$.

A century might be assigned 80 hoplite molecules and 20 support-staff molecules. A headquarters unit might be assigned a configuration of officers, communications equipment, and intelligence-gathering apparatus. The allocation function is where the recursive levels meet: configurations (level 3) are assigned to units in a command structure (level 4).

1.16 Definition (Structured Force)

A structured force *is a triple*

$$F = \big((V, \preccurlyeq_{\text{cc}}),\ R,\ \text{assign}\big),$$

where $(V, \preccurlyeq_{\text{cc}})$ *is a command structure,* R *is a set of non-command relations on* V*, and* $\text{assign}\colon V \to \mathbb{M}_L^{\star}$ *is a force allocation. The set of all structured forces is denoted* $\mathbb{F}^{\star}(\mathbb{M}_L^{\star})$.

This is the full ontological object. A structured force is a command hierarchy, a collection of non-command relationships, and an assignment of configurations to every unit in the hierarchy. It is the constructive answer to the Quinean question posed in Chapter 0.

1.17 Proposition (Countability of Structured Forces)
$\mathbb{F}^\star(\mathbb{M}_L^\star)$ *is countable.*

Proof. For each $n \in \mathbb{N}$, the number of partial orders on $\underline{n}$ is finite. The number of finite collections of hypergraphs on $\underline{n}$ is finite (each hypergraph is a subset of $2^{\underline{n}}$, and there are finitely many such subsets; a finite collection drawn from a finite set is itself one of finitely many possibilities). The number of allocation functions assign: $\underline{n} \to \mathbb{M}_L^\star$ is countable ($\mathbb{M}_L^\star$ is countable by Lemma 1.11, and there are $|\mathbb{M}_L^\star|^n$ such functions, a countable number). Thus the set of structured forces with $|V| = n$ is countable, and $\mathbb{F}^\star(\mathbb{M}_L^\star) = \bigcup_{n=1}^{\infty} \mathbb{F}^\star(\mathbb{M}_L^\star)^{(n)}$ is a countable union of countable sets. ∎

The recursive pattern. Let us pause to observe what has happened. At each level, the same pattern has played out: we introduced new objects built from the previous level, showed that they are countable, and asked what the Force-Maker can see. Elements are to molecules as molecules are to configurations as configurations are to structured forces. At each transition, new structure emerges—edges, then gathering, then command and support—but the underlying logic is the same. This is the fractal nature of force: the same mathematical machinery, applied recursively, at each scale.[13]

1.5 From Soldier to Armed Forces

The construction is not an abstraction awaiting illustration. It *is* the architecture of every organized fighting force that has ever existed. To see this, we will trace a single path through the recursive structure: from one soldier in the United States Army, up through every echelon, to the entirety of the armed forces of the most powerful state in the international system. At every level, the same pattern recurs. At every level, the same mathematics applies.[14]

[13] Every recursive definition must eventually bottom out in a base case that avoids self-reference; in our construction, this is Primitive 1, the elements L. Without a bottom, the recursion is not a construction but a circle. Hofstadter (1979, p. 127) calls this general pattern "nesting, and variations on nesting."

[14] Data in this section are drawn from the Congressional Budget Office (2021), which provides the most detailed publicly available accounting of the U.S. military's force structure.

Element. Start with a soldier. One human being: an element of L. She is not yet a force. She is a label, a point in the index set, a name in a roster. She carries nothing, commands no one, belongs to no unit. She is the irreducible atom of organized violence: necessary but not sufficient.

Beside her, an M1A2 Abrams main battle tank. Another element of L—different in every physical respect, but formally identical in status. A tank without a crew is not a force any more than a soldier without a weapon. Both are elements, waiting to be composed.

Molecule.

Put a crew in the tank: a tank commander, a gunner, a loader, a driver. Now we have a *molecule* in the sense of Definition 1.4: a connected labeled graph with five nodes (four soldiers, one tank) and edges encoding the crew relationship—who operates what, who commands whom inside the vehicle. This is the atomic unit of armored warfare: not the tank alone, not the soldiers alone, but the connected graph of their relationship.

A rifleman with his weapon is a simpler molecule—two nodes, one edge. A mortar team is three soldiers and a tube. The molecular level is where elements become *capable*: the edges encode the knowledge, training, and physical configuration that turn a collection of elements into something that can act.

Configuration.

A heavy platoon gathers four tank molecules. No command structure yet, no subordination—just four tank crews, four copies of the same molecular graph placed side by side in the free commutative monoid $\mathbb{M}_L^\star$. This is a configuration: the $\uplus$ operation of Primitive 1.8, applied three times. Sixteen soldiers, four tanks, and the platoon has not yet been *organized*—it has only been *gathered*.

Company.

Now add command. A company is two to five platoons under a captain, with a headquarters element: a command structure $(V, \leqslant_{cc})$ in the sense of Primitive 1.13, where each node v is a platoon or headquarters section, and $\leqslant_{cc}$ encodes subordination to the company commander. A heavy company has about 60 to 200 personnel and 14 armored vehicles. The allocation function assign gives each platoon its configuration. The company is the first level at which we have a *structured force*: the full triple $\big((V, \leqslant_{cc}), R, \text{assign}\big)$.

At this level, the formal apparatus begins to bite. A rifle company and an armor company are *different structured forces*—different molecules in the configurations, different allocation functions—even if they have the same number of nodes in the command structure. The variable distinguishes them.

Battalion.

A battalion is a structured force whose nodes are companies. It is commanded by a lieutenant colonel and typically includes three to five combat companies and one support company: 400 to 1,000 personnel. The command structure is a partial order on the companies. The allocation function assigns each company its own structured force—and here the recursion becomes explicit. Each node in the battalion's command structure is itself a structured force at the company level. The fractal has begun.

This is where *combined arms* enters. The Combined-Arms Battalion in an Armored BCT contains both rifle companies and armor companies: different types of combat molecules deliberately mixed within a single unit's configuration, on the theory that different arms are more effective in combination than in isolation. In the language of the construction, combined arms is a claim about molecular diversity within the configuration assigned to a node of the command structure.

Brigade.

Figures 1.8 and 1.9 show the next level: the Armored Brigade Combat Team, the basic tactical unit of the modern United States Army. Table 1.1 records the same information numerically.

			Trucks			
Company	Personnel	AVs	L	M	H	Misc.
Brigade HQ						
HQ Company	137	4	29	6	0	0
Battalion Total	137	4	29	6	0	0
Field Artillery						
HQ Battery	233	19	33	10	0	0
Field Artillery (× 3)	91	14	7	1	6	0
Battalion Total	506	61	54	13	18	0
Cavalry						
HQ Troop	116	17	14	6	0	0
Cavalry Troop (× 3)	94	17	1	1	0	0
Armor Company	63	15	2	1	0	0
Battalion Total	461	83	19	10	0	0

Infantry						
HQ Company	177	25	19	5	1	0
Rifle Company (× 2)	137	15	2	1	0	0
Armor Company	64	15	2	1	0	0
Battalion Total	515	70	25	8	1	0
Armor I						
HQ Company	176	25	19	5	1	0
Rifle Company	137	15	2	1	0	0
Armor Company (× 2)	64	15	2	1	0	0
Battalion Total	515	70	25	8	1	0
Armor II						
HQ Company	176	25	19	5	1	0
Rifle Company	137	15	2	1	0	0
Armor Company (× 2)	64	15	2	1	0	0
Battalion Total	515	70	25	8	1	0
Brigade Engineer						
HQ Company	85	5	14	7	0	0
Signal Company	35	0	16	4	0	0
Military Intelligence	118	0	20	6	0	10
Combat Engineer I	113	16	5	1	4	8
Combat Engineer II	98	12	4	1	4	17
Battalion Total	449	33	59	19	8	35
Brigade Support						
HQ Company	85	0	12	9	0	0
Medical Company	82	8	14	10	0	0
Field Maintenance	118	4	10	20	12	0
Distribution Company	140	0	5	1	64	8
Field Artillery FSC	154	5	14	14	21	1
Cavalry FSC	111	6	7	9	7	1
Infantry FSC	147	6	12	12	21	1
Armor I FSC	147	7	12	12	21	1
Armor II FSC	147	7	12	12	21	1
Brigade Engineer FSC	141	3	14	14	13	1
Battalion Total	1,272	46	112	113	180	14
Brigade Combat Team Total	4,222	437	348	185	209	49

Table 1.1: Personnel and equipment of the United States Army Armored Brigade Combat Team, by company. Each row is a force allocation $\mathrm{assign}(v)$ for a unit $v \in V$; the indentation encodes the command structure $(V, \preccurlyeq_{\mathrm{cc}})$; the columns reflect a six-family classification scheme on L. "FSC" denotes a forward support company. "AV" denotes armored vehicles.

The brigade is a structured force with seven immediate subordinates—Field Artillery, Cavalry, Infantry, two Armor battalions, Brigade Engineer, Brigade Support—each itself a structured force. Every component of the triple is visible in Table 1.1.

Classification schemes. The columns—Personnel, Armored Vehicles, Light/Medium/Heavy Trucks, Miscellaneous—are a classification scheme on L in the sense of Definition 1.2: a coarse partition of the elements of force into six families. The Congressional Budget Office does not distinguish an M1 Abrams from a Bradley Fighting Vehicle, lumping both under "armored vehicles." A finer classification scheme would break armored vehicles into tanks, infantry fighting vehicles, and armored personnel carriers; a coarser one might simply count "vehicles." Classification schemes are not idle formalism: they are how real institutions organize their accounting of force.

Configurations as allocation. Each row is a force allocation in the sense of Definition 1.15: the function assign evaluated at a unit $v \in V$. A Cavalry Troop is assigned 94 personnel, 17 armored vehicles, and 1 light truck—that is $\mathrm{assign}(\texttt{CAVALRY TROOP})$, a member of $\mathbb{M}_L^\star$ expressed through the six-family classification scheme. Battalion totals are what one obtains by summing the configurations of subordinate units.

Command structure. The indentation encodes the command structure $(V, \preccurlyeq_{\mathrm{cc}})$. The Brigade HQ sits at the top of the partial order. Below it, seven battalions are its immediate subordinates; each battalion commands its constituent companies. The *span of control* at the brigade level is seven—the same order of magnitude as in Caesar's legions, but with a crucial modern addition: every battalion has a formal headquarters, staffed by its own company of personnel and equipment. The development of formal headquarters at multiple echelons reflects the professionalization of the officer corps traced by Huntington (1957) and Janowitz (1960). Feaver (2003) models civil-military relations as a principal-agent problem:

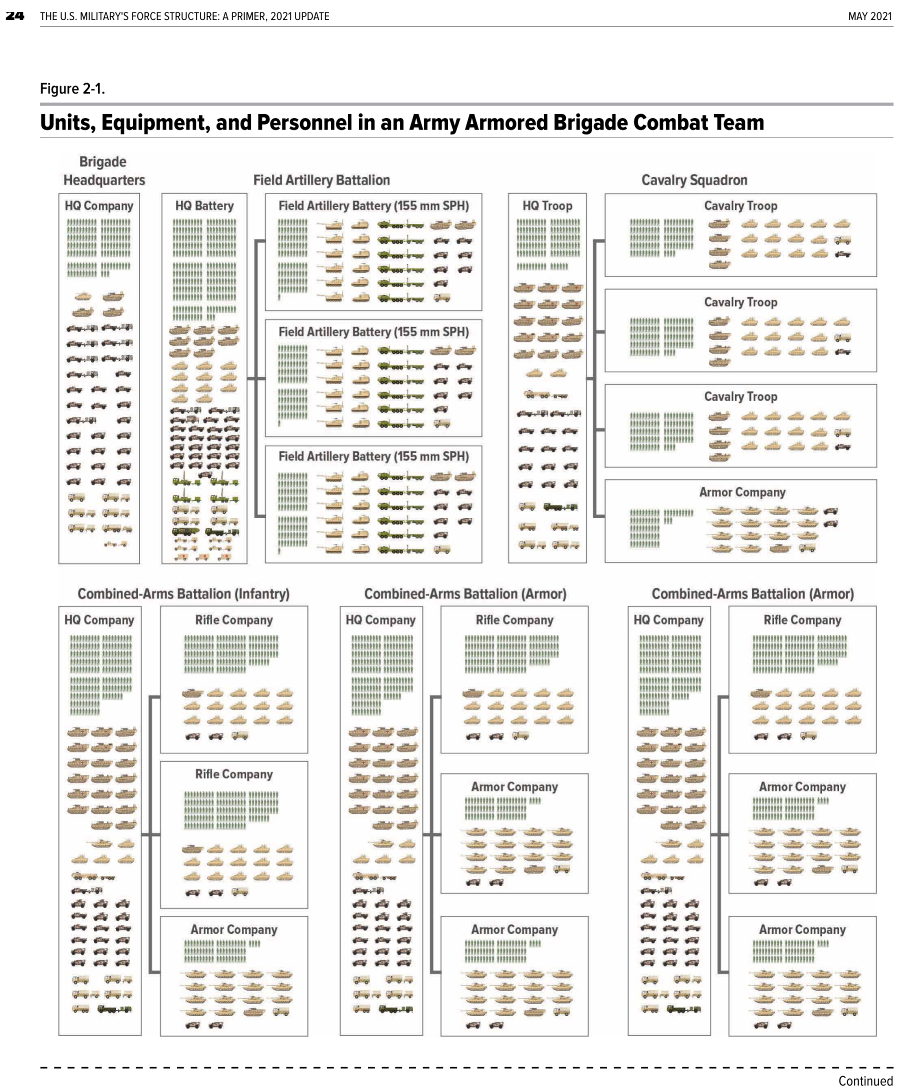

Figure 1.8: *Units, equipment, and personnel in an Army Armored Brigade Combat Team—combat and reconnaissance units. Source: Congressional Budget Office (2021), Figure 2-1.*

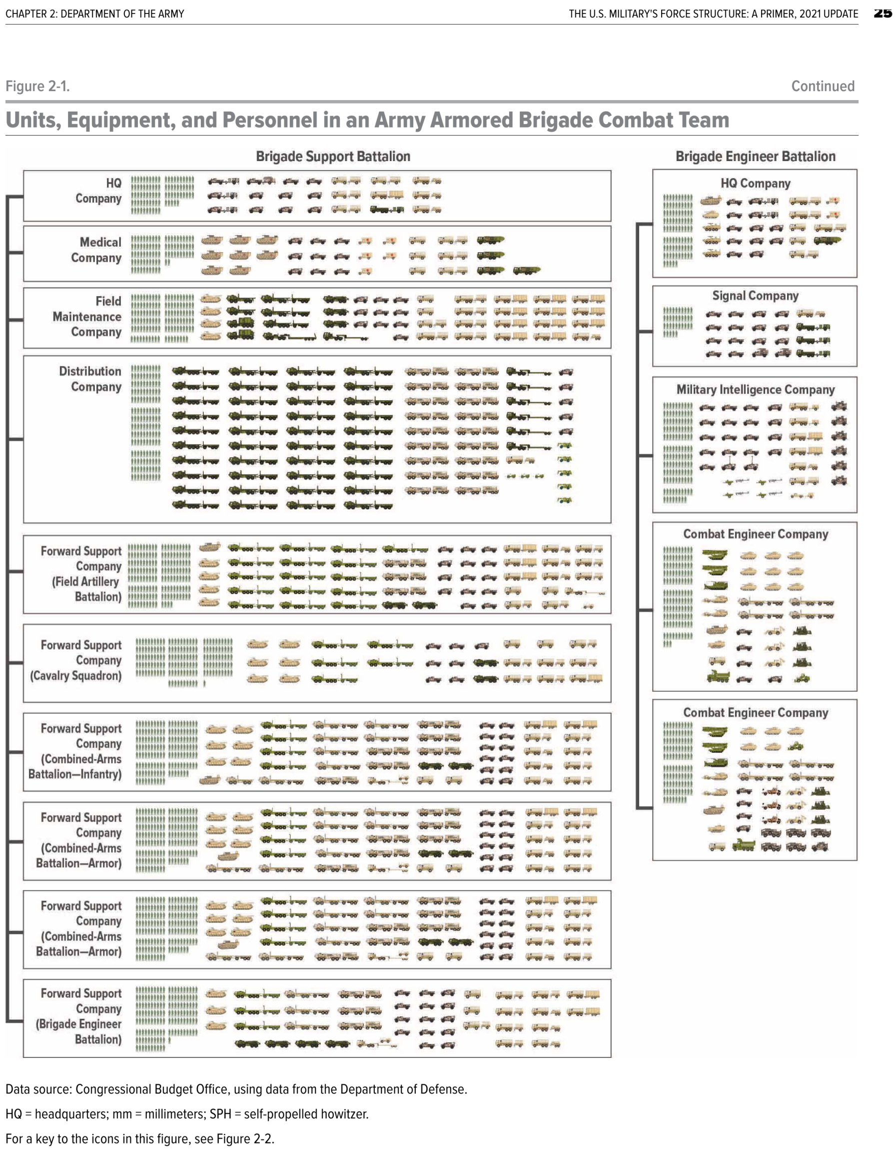

Figure 1.9: *Units, equipment, and personnel in an Army Armored Brigade Combat Team (continued)—support and engineer units. Note the forward support companies, each linked to a specific combat battalion: these are the non-command relations of Definition 1.14. Source: Congressional Budget Office (2021), Figure 2-1 (continued).*

the civilian principal designs monitoring mechanisms—constraints on $(V, \preccurlyeq_{cc})$ imposed from above the military hierarchy.[15]

Non-command relations. The Brigade Support Battalion contains six forward support companies, each named for the combat battalion it supports: the Cavalry FSC supports the Cavalry Battalion, the Infantry FSC supports the Infantry Battalion, and so on. These are non-command relations in the sense of Definition 1.14: the Cavalry FSC is *subordinate* to the Brigade Support Battalion (a command relation) but *supports* the Cavalry Battalion (a non-command relation). These support relationships are hyperedges connecting support units to combat units across branches of the command hierarchy. The Medical Company provides general support to the entire brigade—a single hyperedge connecting one support unit to all combat units at once. As van Creveld (2004) has argued, the revolution in military logistics is an underappreciated dimension of the revolution in military affairs; the BCT's formalized support structure is its institutional expression.[16]

Tooth-to-tail. The Brigade Support Battalion, with 1,272 personnel, accounts for roughly 30% of the BCT's total strength of 4,222. This is the *tooth-to-tail ratio*: the proportion of a force devoted to combat versus support. In the language of the construction, it is a property of the allocation function—the ratio of combat molecules to support molecules summed across the command structure.

Division, Corps, Army.

The recursive structure does not stop at the brigade. A *division* is commanded by a major general and typically contains two to five BCTs, plus an aviation brigade, an artillery brigade, an engineer brigade, and a logistics brigade: 12,000 to 16,000 personnel. It is a structured force whose nodes are brigades. The non-command relations multiply: division-level intelligence assets serve all subordinate BCTs; division artillery provides fires across the entire front.

A *corps* is commanded by a lieutenant general and contains two to five divisions along with numerous support brigades and commands: 40,000 to 100,000 personnel. It is a structured force whose nodes are divisions.

[15] Brooks (2008) extends the civil-military analysis to strategic assessment itself: the quality of information flowing through the command structure depends on the institutional relationship between civilian and military leaders. Avant (1994) shows that domestic political institutions shape which kinds of military change—and therefore which transitions between structured forces—are politically feasible.

[16] See also Erbel and Kinsey (2018) and Bury (2021) for recent assessments of the logistics revolution.

An *army* is the highest command level in a given theater of operations: 100,000 to 300,000 personnel under a general, supporting one or more corps. It is a structured force whose nodes are corps.

In 2021 the United States Army maintained 32 active-component and 28 National Guard BCTs—twelve armored, seven Stryker, and thirteen infantry in the active component alone. The Army's total military personnel numbered approximately 1,016,000, of whom 363,000 served in combat units, 469,000 in support units, and 183,000 in overhead.

Beyond the Army.

The construction is not limited to ground forces. Every branch of the United States military is a structured force, and the variable must accommodate all of them.

The Department of the Navy fields 11 aircraft carriers, each crewed by roughly 3,360 sailors—and each the center of a *carrier strike group* that includes a carrier air wing (1,750 personnel, organized into squadrons of fighters, electronic-attack aircraft, helicopters, and airborne early-warning planes), several escort destroyers (the Navy maintained 72 Arleigh Burke–class destroyers, each with a crew of 350), and supply ships. A carrier strike group is a structured force: the carrier is one node, the air wing another, each destroyer another, with $\preccurlyeq_{cc}$ encoding the strike group commander's authority and non-command relations encoding the coordination between air defense, anti-submarine warfare, and logistics networks. The molecules are radically different from the Army's—an F/A-18E Super Hornet bears no resemblance to an M1 Abrams—but the *structure* of the triple is the same.

The Navy also maintained 53 attack submarines (crews of 200, organized into submarine squadrons) and 33 amphibious ships supporting 24 active-component Marine Corps infantry battalions. The Marine Corps organizes its forces into *Marine air-ground task forces*: integrated combinations of ground combat, air combat, and logistics units tailored to specific operations. A Marine expeditionary unit is a structured force at roughly the brigade level; a Marine expeditionary force is a structured force at the corps level.

The Department of the Air Force organized its combat power into squadrons: 41 fighter squadrons (nominally 12 aircraft each), 3 bomber squadrons, 16 cargo squadrons, 28 tanker squadrons, and 23 unmanned aerial system squadrons. A fighter squadron is a structured force whose molecules—pilots, aircraft, weapons, maintenance crews—differ in every physical particular from the Army's, but whose formal architecture is identical: a command structure, non-command maintenance and logistics relationships, and an allocation of configurations to each subordinate unit. The new Space Force, organized under the Department of the Air Force, added yet another domain in which the same construction applies. Nuclear weapons are the most consequential elements of L in the modern era. Schelling's (1966) insight—that their value lies in the threat of use, not use itself—means that for nuclear forces, the molecular and organizational structure matters less than the *credibility* of the command structure: the partial order must be believed to terminate in a decision that will actually be executed.

The Department of Defense spent more than $400 billion per year on operation and support across all five services (Congressional Budget Office, 2021). The variable $\mathbb{F}^\star(\mathbb{M}_L^\star)$ must accommodate all of this. From the single soldier—one element of L—to the entirety of the United States Armed Forces, every level is an instance of the same construction. A tank crew is a molecule. A platoon is a configuration. A company is a structured force. A battalion is a structured force whose nodes are structured forces. A brigade is a structured force whose nodes are structured forces whose nodes are structured forces. A carrier strike group is a structured force whose nodes include an air wing that is itself a structured force whose nodes are squadrons. And so it goes, all the way up: the same triple $\big((V, \preccurlyeq_{\mathrm{cc}}), R, \mathrm{assign}\big)$, applied recursively, at every scale. This is what force being recursive means: not a claim about any single level, but a structural fact about the relationship between levels that persists from the smallest molecule to the largest joint force.

Does the recursion terminate? The pattern invites a question: what comes *next*? Elements bond into molecules; molecules gather into configurations; configurations are assigned to nodes in a command structure. Each step introduces a new kind of relation—bonding, gathering, commanding. Is there a fifth?

There is, and it is empirically familiar: *alliance.* When sovereign states coordinate their structured forces—NATO, the Coalition in the Gulf War, the Allied command in World War II—the result is a structured force whose nodes are themselves entire national structured forces. Formally, this is $\mathbb{F}^\star(\mathbb{F}^\star(\mathbb{M}_L^\star))$: the construction applied to its own output. But the fifth level differs from the fourth in a way the fourth did not differ from the third. Within a state's military, the command relation $\preccurlyeq_{\mathrm{cc}}$ is *authoritative*: the sovereign designs it, and subordinates obey it. Between sovereign states, command becomes *negotiated, conditional, and revocable.* Weitsman (2014) documents how alliance cohesion and command integration vary dramatically across coalition types; Snyder's (1997) alliance security dilemma—the tradeoff between entrapment and abandonment—is a claim about the inherent incompleteness of the coalition's partial order. National caveats—restrictions on how a contributing nation's forces may be employed—mean that the coalition's partial order is systematically incomplete: some units cannot be ordered to perform certain operations, not because of any military consideration but because a sovereign has withheld consent. The verb sequence is: *exist, bond, gather, command, ally.* And "ally" marks a qualitative break, because the command relation now carries political constraints that are not the Force-Maker's to remove. The recursion does not stop; but the nature of the relation changes, and the analysis of that change belongs to a different inquiry.

1.6 The Variable Defined

The construction is complete. A *force* is a finite partial order representing an organizational chart, along with a finite collection of hypergraphs on the same set of nodes representing non-command relationships, where each node in the partial order has been assigned some collection of connected graphs representing the molecules of force, where the nodes of these molecules are assigned one of a countable set of elements of force.

Formally: a force is a member of $\mathbb{F}^\star(\mathbb{M}_L^\star)$. The variable ranges over all such objects. It is comprehensive enough to accommodate the Roman legion and the U.S. Marine Corps, the Zulu *amabutho* and the Mongol *tumen*, the Athenian trireme fleet and a modern carrier strike group. It accommodates forces that have never existed, forces that could exist, and forces that should never exist. It is the value space of the bound variable in "there exists some x such that x is a force."

The system of systems.

What has this construction accomplished? Andrew Marshall, who directed the Pentagon's Office of Net Assessment for over four decades, insisted that military capability could not be understood by examining any single weapon system, unit, or doctrine in isolation. What mattered was the *system of systems*: the emergent properties arising from the interaction of weapons, organizations, doctrine, training, and support infrastructure.[17]

[17] Marshall's framework, developed through decades of classified net assessments, is most accessible through the work of his protégés. See Krepinevich (1997) for the system-of-systems concept in the context of the revolution in military affairs, and Krepinevich and Watts (2015) for

The construction formalizes exactly this insight. A structured force is not a list of weapons; it is not a headcount; it is not an organizational chart. It is all three simultaneously: the command structure $(V, \leqslant_{cc})$, the non-command relations R, and the allocation function assign together constitute a system in which every component interacts with every other. The same 4,222 soldiers and 437 armored vehicles, organized under a different partial order or with different non-command relations or with different allocation of molecules to units, would be a *different force*—and, as Biddle (2004) has argued at length, potentially a vastly more or less effective one. The variable captures this: two members of $\mathbb{F}^{\star}(\mathbb{M}_L^{\star})$ that agree on total headcount and equipment but differ in organization are formally distinct.

Why it matters.

Talmadge (2015) has shown just how consequential this distinction is: authoritarian leaders who fear military coups deliberately degrade their own command structures—restricting lateral communication, fragmenting the chain of command, inserting political officers—producing structured forces that are organizationally crippled despite possessing the same equipment as their more effective rivals. In our language, coup-proofing is a constraint on the partial order $(V, \leqslant_{cc})$ and the non-command relations R, imposed for political rather than military reasons. Reiter and Stam (2002) argue the positive complement: democracies fight more effectively because of superior initiative at the tactical level—their structured forces feature more delegation in the partial order.[18] The variable distinguishes the resulting forces; the battlefield confirms the distinction. Pollack (2019) extends this point across the entire modern history of Arab warfare: Egyptian, Jordanian, and Syrian armies entered the 1967 war with material superiority over Israel—more troops, more tanks, comparable or superior equipment—and were crushed in six days. The same patterns recurred in 1973, in the Iran–Iraq War, in Libya's wars in Chad. Pollack systematically evaluates four candidate explanations (Soviet doctrine, politicization, underdevelopment, culture) and finds that the answer lies not in any single input but in how organizational practices—exactly our $(V, \leqslant_{cc})$ and R—transform the same materiel into radically different fighting forces.

This is what the military effectiveness literature has been circling for decades.

a broader treatment of Marshall's approach to net assessment.

[18] On the micro-foundations of cohesion, see Shils and Janowitz (1948), who showed that Wehrmacht soldiers fought for primary-group bonds, not ideology—non-command relations R that the construction captures structurally but not motivationally.

The landmark collection is Millett and Murray's (1988) three-volume *Military Effectiveness*, whose framework distinguishes political, strategic, operational, and tactical effectiveness—levels that map onto different scales of the recursive construction. Posen (1984) showed that doctrine—the principles governing how force is employed—is shaped by organizational politics as much as by strategic necessity. Doctrine, in our language, constrains which structured forces the Force-Maker will consider: it narrows her scope. Allison and Zelikow's (1999) Model II—the organizational process model—makes a related point: military organizations are bureaucracies with standard operating procedures that constrain which members of $\mathbb{F}^\star(\mathbb{M}_L^\star)$ are reachable from any given starting point. Biddle (2004) demonstrated that force employment—the interaction of cover, concealment, dispersion, suppression, and combined arms at the tactical level—matters far more than raw materiel advantage. Force employment, in our language, is a claim about the joint structure of the allocation function and the command hierarchy: not how *much* force you have, but how it is *composed*. The relevant variable is not $\sum_v |\text{assign}(v)|$—total inventory—but the molecular diversity within each $\text{assign}(v)$ and the command structure that enables different molecule types to coordinate. A force with rifle and armor molecules assigned to the same company-level node under a command structure that authorizes combined-arms maneuver is a different structured force—and, Biddle demonstrates, a dramatically more effective one—than the same molecules disaggregated into single-arm units under a rigid hierarchy, even when the two forces are identical in total headcount and equipment. Rosen (1991) asked a different but complementary question: not which forces are effective, but how military organizations *innovate*—how they move from one structured force to another. His definition of major innovation—"a change in one of the primary combat arms of a service in the way it fights, or alternatively, the creation of a new combat arm"—is, in our language, a transition from one member of $\mathbb{F}^\star(\mathbb{M}_L^\star)$ to another: new molecules, new configurations, new command structures. The British Army's adoption of the tank in 1916–1918, the U.S. Navy's development of carrier aviation in the interwar period, the creation of the submarine force—each is a rewriting of the structured force triple. In Chapter 2 we will formalize exactly this kind of transition.

The question of why some armies fight effectively and others do not—the question that animates Posen, Biddle, Rosen, and the entire military effectiveness tradition—is, in the language of the construction, a question about which regions of $\mathbb{F}^\star(\mathbb{M}_L^\star)$ produce victory and which do not. That question is beyond the scope of this book. But the variable must be built before the question can be asked precisely.

The countability chain is complete. L is countable by fiat. $\mathbb{M}_L$ is countable because it is a countable union of finite sets of finite connected labeled graphs. $\mathbb{M}_L^\star$ is countable because it is the set of finite multisets over a countable set. $\mathbb{F}^\star(\mathbb{M}_L^\star)$ is countable because it is a countable union of countable products of finite partial orders, finite hypergraph collections, and countable allocation functions. This matters.

Section 0.2 asked whether a force function $m\colon X \to \mathbb{R}$ representing "at least as forceful as" can always be found, and warned that on uncountable sets the answer is no. Countability changes the answer.

1.18 Corollary

Any complete preorder defined on $\mathbb{F}^\star(\mathbb{M}_L^\star)$ *admits a monotone function to* $\mathbb{Q}$*.*

Proof. Quotient by indifference to obtain a countable total order; embed it in $(\mathbb{Q}, \leq)$ by induction, using density to place each new element. ∎

Dupuy's Combat Power Potential is such a function. So is the contest model's m_i. The corollary says maps like these must exist—and they land in $\mathbb{Q}$, not $\mathbb{R}$: the rationals suffice.

But existence is cheap. A force function that assigns wildly different values to forces separated by a single reassignment—one lieutenant moved from one platoon to another—is valid in the sense of the corollary, but vacuous. What we actually need is *continuity*: small changes in force should produce small changes in m. And the bare countable set carries no natural topology in which continuity means anything. The question is never whether a force function exists but whether it respects the structure it compresses.

But force is anything but still. The variable is defined, but a photograph of an army is not a theory of war. Forces change: molecules are built and broken, configurations are gathered and dispersed, organizations are restructured and reformed. The hoplite's panoply was eventually obsoleted by the Macedonian *sarissa*; the Roman legion gave way to the medieval host; the massed infantry of the First World War was shattered by the tank. These transformations are not merely historically interesting. They are the raw material from which the cost metric of Chapter 3 will be built: two forces are close when the cheapest way to transform one into the other is cheap. The grammar of change is also the geometry of the space.

In the next chapter, we turn to the question of how one force becomes another, approached through the compositional language of double-pushout graph rewriting. Each rewrite rule specifies what to delete, what to preserve, and what to create; rules compose into a *category of force*, **Force**, whose morphisms are valid transformation sequences and whose monoidal product is the disjoint union $\oplus$ we built here. The void force $\mathbb{0}_{\mathbb{M}_L}$—the empty table at which this chapter began—will reappear as the monoidal unit and, under unrestricted rules, as a universal waypoint through which any force can reach any other. The category is a resource theory: the convertibility preorder tells us which forces are reachable from which, and resource monotones—real-valued functions that respect convertibility—will show that no single number can capture everything about force that matters. That multidimensionality is not a deficiency of measurement but a structural fact, and the formal apparatus of the next chapter will make it precise.

Chapter 2

Force is Dynamic

The art of war is simply the art of producing such transformations, and its equipment, its processes, even the casualties it inflicts on the enemy, are only means directed toward this end.

Simone Weil, "The Iliad, or the Poem of Force"

On 5 December 1757, Frederick II of Prussia engaged an Austrian army at Leuthen in Silesia under a numerical disparity severe enough that his own generals counseled against the attack: roughly 33,000 Prussians against a defending force nearly twice that size. Frederick attacked anyway, executing the oblique order of battle—a formation in which the attacker concentrates against one flank of the defending line *en échelon*, refusing the opposite wing, and drives against the chosen point with locally superior force before the enemy can redeploy. The maneuver was complex; it required battalions to advance at different rates along converging axes, to shift artillery forward at the correct moment, and to exploit local successes without waiting for orders from the center. Executed against a prepared defender of superior numbers, it required that the battalions be able to do all of this more or less automatically. Frederick had spent the autumn of 1755 at Potsdam ensuring that they could: drilling the oblique order until the movements were habitual. At Leuthen, the Austrians' southern flank was shattered before their northern wing had received orders to move. The battle was over in a few hours.

What the autumn of 1755 had changed was not the elements of the Prussian force—the same soldiers, muskets, and regiments were there before and after—but the *structure* of the force. A transformation rule of considerable complexity had been installed, at the cost of weeks of drill, into the battalions themselves;

the force carried the rule, and could fire it autonomously. Frederick had altered not what his force *had* but what it could *do*—and, crucially, what it could do without him watching every step. That change was as real and as consequential as any addition of troops or materiel; it simply belongs to a different category of fact about force.

The stylized facts this episode illustrates are the organizing commitments of the present chapter. Force transforms: it does not merely exist in some configuration and remain there. Transformations are events with a before and an after, and the relationship between those two states is not arbitrary; it is governed by what rules the force carries, what operations are available, and what constraints must be satisfied for a transformation to be valid. Some transformations are reversible, others are not; some are cheap, others costly; some require specialized structure to execute at all. A theory of force that represents only what force *is* and not how force *changes* is a theory that cannot explain Leuthen, or Cynoscephalae, or Goldwater–Nichols—the cases that most reward explanation.

Chapter 1 built the variable $\mathbb{F}^\star(\mathbb{M}_L^\star)$—the set of all structured forces—from elements to molecules to configurations to organizations, and thereby gave the Force-Maker a precise account of what force *is*: a structured force $F = ((V, \preccurlyeq_{\mathrm{cc}}), R, \mathrm{assign})$ comprising a command hierarchy, a collection of non-command relations, and an allocation of molecular configurations to organizational units. An ontology of objects is, however, only half of what she needs. The Force-Maker does not merely *possess* a force; she *changes* it. She trains soldiers, reorganizes divisions, replaces obsolete equipment, invents new formations. Weil understood what this implies: the art of war is the art of producing *transformations*, and everything else—equipment, processes, even casualties—is a means to that end.

This chapter builds the mathematical structure that will carry the rest of the book: the *category of force*, **Force**. We state its shape immediately, because the destination is not the question here.

A *category* consists of a collection of objects together with a collection of morphisms—arrows, or processes—between them, subject to two requirements: morphisms compose (a process $f : A \to B$ and a process $g : B \to C$ together determine a composite process $g \circ f : A \to C$), and every object carries an identity morphism (the process of doing nothing, which composes trivially on either side). The axioms are deliberately minimal; they impose exactly what is needed to reason about transformation in general, without presupposing anything about what the objects or morphisms happen to be. A category is a mathematical formalization of the idea that transformation is compositional: the meaningful question is not what a single step can do but what a chain of steps can reach—and composition is how chains are built.

Force is a *symmetric monoidal category*.[1] Its objects are the members of $\mathbb{F}^\star(\mathbb{M}_L^\star)$. Its morphisms are *force transformations*: the processes by which one structured force becomes another. Composition chains transformations; the identity morphism is the option of doing nothing. The monoidal product is the disjoint union $\oplus$ from Chapter 1, with the void force $\nvdash_{\mathbb{M}_L}$ as the monoidal unit. The symmetry $F_1 \oplus F_2 \cong F_2 \oplus F_1$ says that independent forces can be placed alongside each other in either order. All of this follows directly from Chapter 1 and requires no further argument.

The question that occupies the rest of the chapter is harder: what, precisely, are the morphisms of **Force**? Forces are complex objects—partial orders, hypergraphs, and labeled molecular graphs all interlocked—and there is more than one way to think about how they change. This chapter enacts the program Section 0.4 introduced: it builds more than it needs, and then forgets with discipline. The next four sections examine specific accounts of what the morphisms of **Force** might be—elementary graph edits, double-pushout rewrites, catalytic transformations, and organizational restructuring—each illuminating features of force dynamics that the others leave in shadow, none of them the final word. We begin with the most conservative possible answer and work toward the cases where the theory is genuinely open; then we step back, and what remains after the forgetting is the convertibility preorder, the theory of resource monotones, and three distinct faculties of military genius that the framework makes formally precise.

2.1 The Minimal Account: Elementary Edits

The most conservative possible account of force transformation identifies a morphism with a finite sequence of *elementary edits*: primitive operations that add or remove nodes and bonds one at a time.

[1] The framework we adopt draws on Coecke, Fritz and Spekkens (2016), who formalize a resource theory as a symmetric monoidal category: objects are resources, morphisms are transformations between them, and a monoidal product captures parallel composition. Gonda and Spekkens (2023) develop the theory of monotones—real-valued functions that respect the convertibility ordering induced by the category's morphisms—which we will connect to the force function m from Corollary 1.18.

Military historians and IR scholars have catalogued the types of force change in various ways, but the lists converge on a small number of qualitatively distinct operations. Rosen (1991) distinguishes quantity from quality from combination; Biddle (2004) distinguishes inventory from employment; McNeill (1982) traces the interaction of configurational change (new weapons and platforms) with organizational change (new command relationships and doctrines) across a millennium of military history. Behind these different framings lies the same underlying partition: a force can be changed by adding elements, removing elements, forging relationships between elements, or severing relationships between elements. The four elementary edits are the formal crystallization of this vocabulary.

Each operation has a direct historical correlate. `+Node` is the commissioning of a new platform, the recruitment of soldiers, the activation of a reserve unit; `-Node` is decommissioning, disbandment, or battlefield attrition. `+Bond` is training in the widest sense—the forging of the coordination relationship between a soldier and her weapons, between units that must maneuver together, or between allied forces that must operate jointly. `-Bond` is the severing of trained relationships: through combat losses, forced officer rotation, or the deliberate disruption that authoritarian coup-proofing strategies impose on the bonds of collective military effectiveness. The four primitives are not a mathematical imposition on history; they are the analytical vocabulary that emerges when historians describe force change with enough precision to be useful.

A structured force is built from labeled graphs at every level: molecules are connected labeled graphs, configurations are finite multisets of them, and organizational structures are partial orders and hypergraphs over the same node sets. The most natural way to transform a graph is to edit it: add a node, remove a node, add an edge, remove an edge. Give us four primitive operations—`+Node`, `-Node`, `+Bond`, `-Bond`—and any transformation can be expressed as a finite sequence of them.[2]

Figure 2.1 illustrates the idea. A soldier acquires equipment (five `+Node` operations), then training forges bonds between soldier and equipment (five `+Bond` operations). Each step is local and easy to understand.

The elementary-edit account is the most conservative possible foundation for

[2] Any two configurations in $\mathbb{M}_L^\star$ can be connected by such a sequence: tear the source down to the void force one operation at a time, then build up the target. The proof is straightforward: remove all bonds to get a collection of dangling nodes, remove all nodes, then add the target's nodes and bonds in any order that respects connectivity.

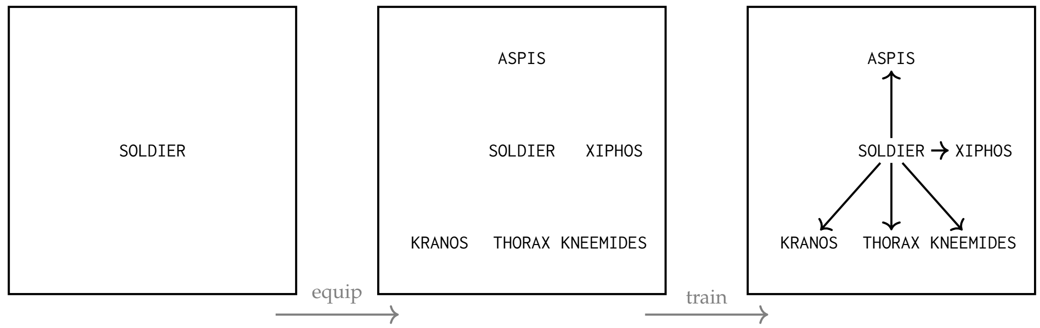

Figure 2.1: *Making a hoplite in two stages: equipping adds atoms (left to middle), training adds bonds (middle to right). The soldier is preserved throughout; everything else is added.*

Force. It commits to almost nothing: it says only that whatever transformations are, they can be *cashed out* in terms of adding and removing nodes and bonds. This is the bedrock. Whatever more sophisticated account we develop in the sections that follow should be consistent with it—a way of organizing these primitive moves, not replacing them.

The elementary-edit account remains the semantic bedrock: whatever more sophisticated account we develop in the sections that follow must be consistent with it, and Remark 2.20 establishes that every morphism in **Force** factors through elementary edits. But the picture operates at a grain size below the level at which the Force-Maker actually plans.

She does not issue orders in primitive operations. When Frederick orders the oblique attack at Leuthen, he does not dictate each footstep; he specifies the maneuver as a unit. When a staff officer writes the order to convert a rifle company to a mechanized one, she does not enumerate each `+Node` and `+Bond`; she specifies a pattern to match and a replacement to install. The Force-Maker operates at the level of complete transformations: she sees a substructure in her current force, she knows what should occupy that position instead, and she issues the move as a single act.

This is the more efficient view—and it has a natural mathematical expression. A complete operational move is a *rewrite rule*: a span $L \xleftarrow{\ell} K \xrightarrow{r} R$ specifying what to remove ($L \setminus \ell(K)$), what to preserve (K), and what to install ($R \setminus r(K)$) in one step. The elementary edits of this section are degenerate rewrite rules, each touching a single node or bond; the general case touches whole substructures at once. The next section develops this account for the case of *configurational* *rewrites*—changes to the molecular content of a structured force.

2.2 Equipping: Double-Pushout Rewriting

The Force-Maker who thinks in complete operational moves is already thinking in rewrite rules. *Double-pushout (DPO) graph rewriting* is the standard algebraic formalization of this mode of operation for labeled hypergraphs.[3] A single DPO step locates a match of the left-hand side L in the current graph G, cuts out $L \setminus \ell(K)$ along the interface K, and glues in $R \setminus r(K)$—all as one coherent act. The surrounding context remains untouched; only the designated substructure is replaced.

As a structural consequence, validity no longer needs to be policed from outside. Two categorical pushouts—one to cut, one to glue—ensure that when an algebraic condition on the match is satisfied, the result is automatically a valid graph. Validity becomes a property of the rule and its match, checkable before the transformation fires, rather than a constraint on sequences enforced after the fact.

We develop this formalism for the class of morphisms we call *configurational rewrites*: transformations that change the molecular content of a structured force. The story is clean here, and the cleanness is itself informative—it tells us what makes configurational change tractable in ways that organizational restructuring (Section 2.4) will turn out not to share.

2.2.1 Rewrite Rules as Spans

We work in a category $\mathcal{G}$ of finite directed labeled hypergraphs with node labels drawn from L and bond types drawn from the appropriate relation set.[4] The objects of $\mathcal{G}$ include (the underlying graphs of) all molecules in $\mathbb{M}_L$ and all configurations in $\mathbb{M}_L^\star$.

[3]The DPO approach originates with Ehrig, Pfender and Schneider (1973) and was developed extensively in Ehrig et al. (2006). For a modern survey, see Corradini et al. (1997).

[4]Formally: objects of $\mathcal{G}$ are pairs (V, E, λ) where V is a finite set of nodes, E is a finite set of directed hyperedges (S, T, r) with $S, T \subseteq V$ nonempty and r a bond type, and $\lambda : V \to L$ is a labeling function. Morphisms are label-preserving maps that send nodes to nodes and edges to edges, preserving source sets, target sets, and bond types. This is the category in which the molecular and configurational structures of Chapter 1 live.

2.1 Definition (Rewrite Rule)

A rewrite rule *(or* production*) is a span in $\mathcal{G}$:*

$$p = \left(L \xleftarrow{\ell} K \xrightarrow{r} R\right),$$

where L, K, and R are objects of $\mathcal{G}$ and ℓ, r are morphisms (injective graph homomorphisms). We call L the left-hand side *(the pattern to be matched), K the* interface *(what is preserved), and R the* right-hand side *(the replacement).*

The interface K is the structural common ground between what is removed and what takes its place. It encodes the nodes and bonds that persist through the transformation, anchoring the replacement in the surrounding context.

2.2 Example (Equipping a Soldier)

The equipping transformation from Figure 2.1 (left to middle panel) is a rewrite rule with:

- $L = \{\texttt{SOLDIER}\}$*: a single node, no edges.*
- $K = \{\texttt{SOLDIER}\}$*: the soldier is preserved (same as L).*
- $R = \{\texttt{SOLDIER}, \texttt{KRANOS}, \texttt{THORAX}, \texttt{KNEEMIDES}, \texttt{ASPIS}, \texttt{XIPHOS}\}$*: the soldier plus five equipment nodes, no bonds yet.*

The interface maps $\ell : K \hookrightarrow L$ and $r : K \hookrightarrow R$ both send the soldier to itself. The rule says: wherever you find a bare soldier, you may replace him with a soldier surrounded by equipment.

2.3 Example (Training a Hoplite)

The training transformation (middle to right panel of Figure 2.1) is a different rule:

- $L = \{\texttt{SOLDIER}, \texttt{KRANOS}, \texttt{THORAX}, \texttt{KNEEMIDES}, \texttt{ASPIS}, \texttt{XIPHOS}\}$*: six nodes, no bonds.*
- $K = L$*: all six nodes are preserved.*
- R*: the same six nodes, now connected by five bonds from soldier to equipment.*

The interface maps ℓ and r are both the identity on nodes. The rule says: wherever you find an equipped-but-untrained soldier, you may forge the bonds of training.

Notice that the four primitive edit operations of Section 2.1 are all special cases of rewrite rules. $\texttt{+Node}_\ell$ is the rule with $L = K = \varnothing$ and R a single node labeled ℓ. $\texttt{-Node}_\ell$ reverses this span. $\texttt{+Bond}_{(S,T,r)}$ has $L = K = R$ on the same nodes, with R carrying one additional edge. $\texttt{-Bond}_{(S,T,r)}$ reverses. DPO subsumes the primitives; its power lies in rules that are not primitive—rules that simultaneously delete substructures, preserve interfaces, and insert replacements.

2.2.2 Applying a Rule: The Double Pushout

A rewrite rule is an abstract template. Applying it to a specific force requires finding the pattern inside the force and then performing the replacement.

2.4 Definition (Match)
Let $p = (L \overset{\ell}{\leftarrow} K \overset{r}{\rightarrow} R)$ *be a rewrite rule and let* G *be an object of* $\mathcal{G}$ *(a host graph representing a force configuration). A* match *of* p *in* G *is a morphism* $m : L \to G$*—an embedding of the left-hand side into the host graph.*

A match identifies *where* in the current force the rule's pattern occurs. The Force-Maker who looks at her order of battle and sees that a particular formation matches the left-hand side of a known rewrite rule has found a match.

Given a rule and a match, the DPO construction performs the replacement in two steps, each a categorical pushout. The construction is summarized by the following diagram, in which the top row is the rule, the bottom row is the rewrite, and the vertical arrows connect them:

$$\begin{array}{ccccc} L & \xleftarrow{\ \ell\ } & K & \xrightarrow{\ r\ } & R \\ \downarrow{\scriptstyle m} & & \downarrow{\scriptstyle d} & & \downarrow{\scriptstyle m'} \\ G & \xleftarrow{\ g\ } & D & \xrightarrow{\ g'\ } & G' \end{array}$$

2.5 Definition (DPO Rewrite Step)
Given a rewrite rule $p = (L \overset{\ell}{\leftarrow} K \overset{r}{\rightarrow} R)$ *and a match* $m : L \to G$*, a* DPO rewrite step $G \Rightarrow_p G'$ *consists of:*

1. ***Deletion.*** *Construct the* pushout complement *D of $\ell : K \to L$ and $m : L \to G$: the unique (up to isomorphism) object D and morphisms $d : K \to D$, $g : D \to G$ making the left square a pushout. Intuitively, D is G with the matched portion of L removed, retaining the interface K.*

2. ***Addition.*** *Construct the* pushout *G' of $r : K \to R$ and $d : K \to D$: the unique (up to isomorphism) object G' and morphisms $m' : R \to G'$, $g' : D \to G'$ making the right square a pushout. Intuitively, G' is D with the replacement R glued in along the interface K.*

The result G' is the rewritten graph.

The construction has a beautiful feature: it may fail, and when it does, it fails precisely because the proposed application is structurally incoherent. This is the algebraic content of the *gluing condition*.

2.6 Proposition (Gluing Condition)

The pushout complement in Step 1 exists if and only if the gluing condition *is satisfied:*

(a) ***Dangling condition.*** *No edge in G that is not in the image of m is incident to a node in $m(L) \setminus m(\ell(K))$—that is, removing the matched nodes would not leave any edges dangling.*

(b) ***Identification condition.*** *The match m restricted to $L \setminus \ell(K)$ is injective—distinct nodes and edges to be deleted are matched to distinct nodes and edges in G.*

Proof. This is the standard result for DPO rewriting in adhesive categories; see Ehrig et al. (2006, Theorem 3.11). The category $\mathcal{G}$ of finite directed labeled hypergraphs is adhesive (pushouts along monomorphisms are van Kampen squares), so the result applies. ■

When the gluing condition fails, the rewrite step does not produce an error or an invalid intermediate—it simply *does not apply*. The pushout complement fails to exist, and the transformation is undefined. This is not a bug but the core feature. The validity problem of Section 2.1 is dissolved, not solved: it never arises. A rule that would break the surrounding structure by leaving dangling edges produces no output at all.

The gluing condition is also the formal expression of an observation that runs through the IR literature on force employment, usually stated informally. Biddle (2004) argues that military effectiveness depends not on material inventory alone but on whether the components of a force are operationally integrated: removing or replacing one component of a well-integrated force often degrades or destroys the value of the others. A force trained to fight as a combined-arms team cannot simply have its artillery stripped out without losing the infantry doctrine that assumed artillery support—the infantry bonds are dangling edges. The gluing condition makes this claim algebraic: you cannot delete a substructure while it remains load-bearing for the surrounding graph. What Biddle describes as "force employment" is, in part, the condition that the gluing condition is satisfied at every step of the transformation—that no component is removed while it is still structurally integral to the force's operational pattern.

To see the double pushout in action, consider the training rule from Example 2.3 applied to a host graph that contains not just the soldier-with-equipment but also a second soldier standing nearby. The second soldier is *context*: she is not part of the match and must be preserved untouched by the rewrite.
The second soldier, shown in gray, illustrates the fundamental promise of DPO: context is preserved. She has no part in the match m, so she appears unchanged in D and in G'. The gluing condition ensures this: because no edge in G outside the match is incident to a node that would be deleted, the surrounding structure is untouched. If the second soldier had been bonded to the first soldier's equipment—say, by a `shares` edge—and the rule tried to delete that equipment, the dangling condition would fail and the rule would not apply. The surrounding structure protects itself.

Parallel independence. The Force-Maker rarely transforms her force one step at a time. She trains one battalion while reorganizing another; she upgrades artillery in the south while repairing bridges in the north. When two transformations touch different parts of the force, we should expect them to be independent—and the DPO formalism makes this precise.

Two DPO rewrite steps $G \Rightarrow_{p_1,m_1} H_1$ and $G \Rightarrow_{p_2,m_2} H_2$ are *parallel independent* if their matches do not interfere: the image of m_1 in G, minus the interface nodes that are preserved, is disjoint from the image of m_2, and vice versa.[5]

[5]Formally: $m_1(L_1) \cap (m_2(L_2) \setminus m_2(\ell_2(K_2))) = \varnothing$ and $m_2(L_2) \cap (m_1(L_1) \setminus m_1(\ell_1(K_1))) = \varnothing$. The condition says that neither step tries to delete or modify anything the other step needs.

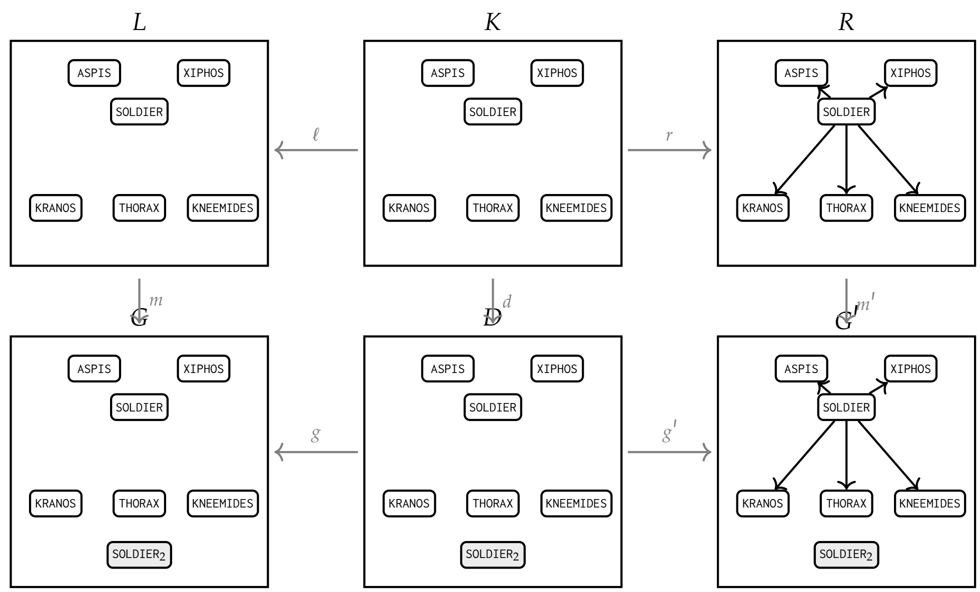

Figure 2.2: *The training rule applied to a host graph G containing one soldier with equipment and a second soldier (gray, context).* Top row*: the rule* $L \xleftarrow{\ell} K \xrightarrow{r} R$*. L and K contain the same six nodes (no bonds); R adds training bonds.* Bottom row*: the rewrite. G is the host graph. D is the pushout complement—G with the matched part of L removed, but since* $L = K$ *(nothing is deleted),* $D \cong G$*.* G' *is the pushout—D with R's bonds glued in. The second soldier (gray) passes through untouched: she is context, not part of the match.*

2.7 Proposition (Parallel Independence)

If two DPO rewrite steps $G \Rightarrow_{p_1,m_1} H_1$ *and* $G \Rightarrow_{p_2,m_2} H_2$ *are parallel independent, then there exist rewrite steps* $H_1 \Rightarrow_{p_2,m_2'} G'$ *and* $H_2 \Rightarrow_{p_1,m_1'} G'$ *yielding the same result* G'*. That is, the two steps can be applied in either order with the same outcome.*

Proof. This is the Church–Rosser theorem for DPO rewriting in adhesive categories; see Ehrig et al. (2006, Theorem 5.12). Since $\mathcal{G}$ is adhesive, the result applies directly. ■

The theorem says exactly what the Force-Maker needs to hear: if two transformations operate on disjoint parts of the force, the order in which she executes them does not matter. Training the 1st Battalion's riflemen while upgrading the 3rd Battalion's vehicles produces the same result regardless of which happens first.

When two configurational rewrites do *not* commute—when the same equipment must be part of both matches—the DPO formalism locates the conflict precisely in the overlap.

Configurational closure. A DPO rewrite applied to a valid configuration produces another valid configuration.

2.8 Proposition (Configurational Closure)
Let $C \in \mathbb{M}_L^\star$ *be a configuration and let* $p = (L \xleftarrow{\ell} K \xrightarrow{r} R)$ *be a rewrite rule in* $\mathcal{G}$*. If there exists a match* $m : L \to C$ *satisfying the gluing condition, then the result* $C' = C \Rightarrow_p C'$ *is again a member of* $\mathbb{M}_L^\star$*.*

Proof. The pushout in $\mathcal{G}$ preserves finiteness (pushouts of finite objects are finite), label-correctness (pushouts preserve the labeling functor), and the multiset-of-connected-components structure (the gluing condition ensures no dangling edges, so connected components are either modified coherently or left intact). Since $\mathbb{M}_L^\star$ is exactly the set of finite multisets of connected finite labeled graphs over L, the result C' is again a member. ∎

2.9 Example (Rifle Upgrade)
An infantry platoon replaces its bolt-action rifles with semi-automatic ones.

- *L: a soldier node bonded to a* `BOLTRIFLE` *node by a* `wields` *edge.*
- *K: the soldier node alone—the rifle is not preserved.*
- *R: the soldier node bonded to a* `SEMIRIFLE` *node by a* `wields` *edge.*

The bolt-action rifle is in $L \setminus \ell(K)$*: it is consumed by the rule. The semi-automatic rifle is in* $R \setminus r(K)$*: it is created. The soldier passes through unchanged, but the weapon she carries is different.*

The scope of configurational rewriting extends well beyond individual equipment changes. When the Royal Navy launched HMS *Dreadnought* in 1906, it applied a single configurational rewrite to the capital ship: L = a mixed battery of large and small guns with distributed fire control, K = hull, armor, engines, and crew (preserved), R = a uniform main battery of twelve-inch guns with centralized fire-control systems. The rule was not complicated in its statement, but its consequences rewrote the convertibility preorder of every navy in the world. Every pre-Dreadnought battleship fell in the preorder: no force built to the old design could reach the new frontier without applying the same rule or an improved successor. Germany, Japan, the United States, France, Russia, and Austria-Hungary all applied it—the Anglo-German naval race of 1906–1914 is, in our language, a forced expansion of the global rule library, a single configurational rule whose adoption by one power compelled adoption by every other that sought to maintain competitive convertibility. The Dreadnought episode is the formal content of what the "military revolution" literature means when it says a single innovation "obsolesces" existing forces: not merely that existing forces become less effective in some vague sense, but that they fall in the convertibility preorder relative to forces that have adopted the new rule—and that adoption, under any realistic rule library, is compulsory.

The configurational side of **Force** is the more tractable one. The category $\mathcal{G}$ is adhesive, the DPO theory is well-developed, and the closure property is clean. Before turning to the harder problem, there is a specific class of configurational transformation that deserves its own treatment.

2.3 Training and Catalysis

When a unit is trained in a particular maneuver, something happens that goes beyond the forging of bonds between a soldier and her equipment. The soldiers have internalized a pattern—a rewrite rule—that they can execute on the battlefield without external direction. A well-drilled company does not need the Force-Maker standing behind it to execute a flanking maneuver; the rule has been installed through repetition, and the unit carries it as part of its own structure. Training, in this sense, is not merely configurational rewriting: it is *rule installation*, the transfer of transformational authority from the Force-Maker to the units themselves. The distinction matters enormously for how we think about military capability—and it requires a new formal concept to capture.

Consider again the training transformation of Example 2.3: the rule takes six nodes (soldier plus equipment) and adds five bonds; as stated, it is self-executing—soldier plus equipment go in, a trained hoplite comes out. The Force-Maker knows this is not how training works. A drill instructor is required—one of *her* agents, a node in the force graph placed there precisely because she enables this transformation. The instructor must be present for the rule to fire, she shapes the bonds that are forged, and—crucially—she is not consumed by the process. She walks away at the end, unchanged, ready to train the next recruit. In the language of DPO, the drill instructor is a *catalyst*: a substructure that must appear in the match for the rule to fire but that passes through the interface into the output unaltered.

2.10 Definition (Catalytic Rule)
A rewrite rule $p = (L \xleftarrow{\ell} K \xrightarrow{r} R)$ *is* catalytic *with respect to a subgraph* $C \hookrightarrow K$ *if:*

1. C *embeds into both* L *(via* ℓ*) and* R *(via* r*) as an isomorphic copy: the catalyst appears on both sides of the rule, unchanged.*

2. *The match* $m : L \to G$ *requires* C *to be present in the host graph: removing the image of* C *from* G *would cause the match to fail.*

We call C *a* catalyst *for* p*.*

The definition captures a pattern that is everywhere in military transformation: the structural requirement that a certain kind of agent be present not because she is transformed by the process but because, without her, the process cannot occur at all.

2.11 Example (The Drill Instructor)
Rewrite the training rule from Example 2.3 with a catalyst.

- *L: soldier, equipment (five nodes), and a drill instructor node* `DI`*, bonded to the soldier by a* `trains` *edge.*

- *K: all seven nodes—soldier, equipment, drill instructor—preserved. The* `trains` *bond is also in K.*

- *R: all seven nodes, with the five training bonds from soldier to equipment added, and the* `trains` *bond retained.*

The drill instructor is catalytic: she must be present (the rule won't fire without a `trains` *bond in the host graph), she is fully preserved (she passes through K unchanged), and she can catalyze the same rule again on a different soldier.*

Catalysts and force multiplication. A single drill instructor can catalyze the training of an entire cohort by repeated application of the same rule to different matches in the host graph. The instructor is not consumed, so the rule can fire as many times as there are matches. This is the formal content of the phrase "force multiplier": a catalyst is an element of force whose value lies not in what it *is* but in how many transformations it *enables*. Dupuy's (1979) combat power factors—leadership, training quality, morale—are, in this language, claims about the catalytic capacity of certain nodes in the force graph.

King (2013) has documented this process in detail for modern infantry. The transformation from a collection of individual soldiers into a tactically effective small unit is not primarily a matter of equipping (configurational rewriting) but of collective training: the repeated, catalyzed forging of bonds—mutual trust, shared tactical reflexes, coordinated movement under fire—that turn a group into a team. The bonds are invisible in any inventory of equipment, but they are the difference between a force that fights and a force that disintegrates on contact. King's point is that these bonds are not natural or automatic; they require a catalyst (an experienced NCO, a training program, a doctrinal framework) and sustained repetition. In our language: the rule must fire many times, at many matches, before the host graph carries the density of bonds that collective effectiveness requires.

Training establishments as catalyst stocks. A training base is a standing collection of catalytic nodes—drill instructors, ranges, simulators, doctrine manuals—that exist not to fight but to enable the rewrite rules that turn raw recruits into soldiers and soldiers into units. The base is a permanent subgraph of the host graph whose purpose is to be matched, repeatedly, by catalytic rules. If the base is destroyed—if the catalytic nodes are removed from the host graph—the training rules can no longer fire. The force still has its current configurations, but it can no longer produce new ones. This is why states invest so heavily in institutions that never see combat: the peacetime military is, in large part, a standing stock of catalysts.

Talmadge (2015) has shown, devastatingly, what happens when the catalytic stock is deliberately degraded. Authoritarian leaders who fear military coups systematically attack exactly the catalytic nodes: they rotate officers to prevent the accumulation of unit-level bonds, restrict realistic collective training to prevent the emergence of autonomous capability, insert political commissars whose presence *inhibits* rather than catalyzes the forging of professional relationships.

The equipment is unchanged—the tanks are the same tanks, the rifles the same rifles—but the rewrite rules that turn equipment into capability can no longer fire. In our formalism, coup-proofing is the deliberate removal of catalytic subgraphs from the host graph, or worse, the insertion of *inhibitory* nodes that cause the gluing condition to fail where it would otherwise succeed.[6]

The same catalytic logic applies at higher institutional levels. Staff colleges, war-gaming centers, and doctrinal development commands are catalysts for organizational transformation: they do not fight, but without them, the rewrite rules that restructure command hierarchies and develop new operational concepts cannot be applied. The German General Staff system of the nineteenth century was, in this language, a standing catalyst for scheme-invention—a permanent subgraph of the Prussian military whose purpose was to author new rewrite rules, not merely to apply existing ones (Citino, 2005). Murray (2011) has argued that the capacity for wartime adaptation—the ability to write and install new rewrite rules under the pressure of ongoing combat—is the single most important determinant of military success. The armies that adapted in the First World War (the German stormtroop tactics, the British combined-arms methods of 1918) did so because they retained functioning catalytic institutions—general staffs, training commands, doctrinal centers—that could author new rules and install them in combat units faster than the enemy.

Catalytic necessity. The gluing condition makes catalytic necessity a *theorem*, not an assumption.

2.12 Proposition (Catalytic Necessity)
There exist configurations $A, B \in \mathbb{M}_L^\star$ and a rewrite rule p such that no match of p in A satisfies the gluing condition (the transformation $A \Rightarrow_p B$ is impossible), but there exists a catalyst C and a match of p in $A \oplus C$ that satisfies the gluing condition and yields $B \oplus C$.

Proof. By construction. Let A consist of a single node v labeled `RECRUIT`. Let B consist of the node v relabeled `SOLDIER` with five equipment nodes and

[6] An inhibitory node is, in a sense, an anti-catalyst: its presence in the match causes the pushout complement not to exist. A political officer who must approve every lateral communication between units is such a node: the rewrite rule for coordinated maneuver requires certain bonds between units, but the officer's presence introduces additional edges that violate the dangling condition when the rule tries to fire.

five training bonds (a trained hoplite). Let p be the catalytic training rule of Example 2.11: L requires a recruit, equipment, *and* a drill instructor bonded to the recruit; K preserves all of them; R adds training bonds. In A alone, the left-hand side L cannot be matched: A contains no drill instructor node and no `trains` bond, so no morphism $m : L \to A$ exists. Now let C consist of a drill instructor node with the appropriate label. In $A \oplus C$, the match exists: L embeds into the combined graph, the gluing condition is satisfied (no dangling edges, no identification conflicts), and the DPO step produces $B \oplus C$—the trained hoplite plus the drill instructor, unchanged. ■

The proposition says that the force space has transformations that are *structurally impossible* without a catalyst—not costly, not difficult, but literally undefined. No amount of cleverness in sequencing primitive edits can get from A to B through a valid DPO path if the catalytic node is absent from the host graph. The drill instructor is not a convenience; she is a structural prerequisite.

This gives formal teeth to an observation that military historians have made informally: you cannot train an army without trainers. Talmadge's dictators understand this perfectly—which is why they attack the catalytic stock rather than the equipment. Destroy the drill instructors, and the rewrite rules stop firing, regardless of how many recruits and how much equipment you have. The catalytic necessity proposition makes the mechanism precise: it is not that training becomes *harder* without the catalyst but that the relevant DPO step *ceases to exist*.

Installed vs. catalyzed rules. Catalysis explains what enables a transformation in the moment: the catalyst must be present. Rule installation explains what training produces: a force that no longer *needs* the catalyst.

A catalytic rule requires the catalyst to be present in the match every time it fires. An *installed* rule has been absorbed into the force's own structure: the bonds and relationships forged by training encode the rule's interface and replacement pattern, so that the force can execute the transformation autonomously. The Force-Maker's doctrine, in this light, is a library of rewrite rules that her training establishment has installed across the force. Posen (1984) argued that the content of military doctrine—which rules get installed, and which are neglected—is determined as much by organizational politics and civil-military bargaining as by strategic rationality. In our language: the rule library is not optimized; it is *negotiated*, and the Force-Maker may find that the rules her institution has installed are not the ones she most needs.[7]

The catalytic structure is recursive. The drill instructor who trains hoplites was herself trained by someone—a senior instructor, perhaps, or the institutional memory encoded in a training manual. That higher-level training is a different rewrite rule, with the training institution as its catalyst. And the training institution was itself established by a doctrinal decision—another rewrite rule, catalyzed by a staff college or a reforming general. The recursive structure of Chapter 1 reappears at the level of transformations: catalysts are trained by catalysts, all the way up.

Frederick's insight. No historical figure illustrates the primacy of rule installation more clearly than Frederick the Great. Frederick inherited a Prussian army that was, in configurational terms, unremarkable: muskets, bayonets, cannon, cavalry horses—roughly the same molecular inventory as every other European army of the mid-eighteenth century. What he built was a training regime of extraordinary intensity and precision, and the result was not a differently equipped force but a force with a radically larger library of installed rewrite rules.

The oblique order—Frederick's signature tactical innovation—is a rewrite rule. Its left-hand side L is a conventional line of battle; its interface K preserves the battalions and their internal cohesion; its right-hand side R concentrates strength on one wing while the refused wing holds in place, the entire line advancing *en échelon*. As a span it is elegant. As a battlefield maneuver it is nearly impossible unless every battalion can wheel, halt, change front, and advance at a fixed oblique angle without losing cohesion—and this is exactly what Frederick's relentless drill on the fields at Potsdam was designed to install. The drill did not merely teach the rule; it made execution autonomous. At Leuthen in 1757, the Prussian army executed the oblique order against an Austrian force nearly twice its size, and the maneuver worked because the battalions did not need Frederick to micromanage each wheel: the rule had been installed so deeply that the units carried it in their own structure.

[7] The distinction between catalyzed and installed rules maps roughly onto the military distinction between "supervised" and "unsupervised" training. In supervised training, the catalyst (instructor, observer-controller) is present and the rule fires under her guidance. In unsupervised training and in combat, the unit executes the rule from its own internalized repertoire. The progression from supervised to unsupervised is the process of installation. Huntington (1957) argued that professionalization—the creation of a self-regulating military profession—is precisely the institutionalization of this process.

Frederick understood something that many of his imitators did not: a rewrite rule that exists only in the mind of the commander—uninstalled, uncatalyzed, unpracticed—is not a capability but a fantasy. The value of a rule is zero until it is installed, and the cost of installation is the cost of training. Frederick's genius was not primarily in scheme-invention (the oblique order had precedents in Epaminondas and Gustavus Adolphus) or in match-finding (his tactical reads were sometimes disastrous, as at Kunersdorf in 1759). It was in *installation*: the obsessive, unglamorous, repetitive work of drilling rules into the force until they became part of its graph. Citino (2005) traces the entire tradition of German operational art back to this Frederician insight—that the decisive advantage lies not in having better plans but in having forces that can execute complex plans without centralized direction, because the plans have been installed as rewrite rules through training.

2.4 Reorganization

The command hierarchy $(V, \leqslant_{cc})$ and the non-command relations R of a structured force can also change. Historically, these organizational transformations have often been more consequential than any change in equipment. They are also where the theory is genuinely open.

Configurational rewrites operate on labeled hypergraphs—and DPO is the canonical algebraic theory for rewriting labeled hypergraphs. The category $\mathcal{G}$ is adhesive; the gluing condition is clean; parallel independence is a theorem. The organizational layer is different. The command hierarchy is a partial order, not a general graph, and the category of finite partial orders does not inherit adhesivity.[8] This means that whatever account we give of organizational transformation, it cannot simply be DPO on the command hierarchy.

Rather than force the organizational layer into a framework whose conditions it does not satisfy, we take the honest path: describe what organizational transformation *is*, show what it looks like historically, and name the open question it leaves.

What changes in an organizational transformation is the partial order $(V, \leqslant_{cc})$ and the non-command relations R—while, in the archetypal case, the configu-

[8] Pushouts along monomorphisms in the category of finite partial orders need not be van Kampen squares, which is the condition adhesive categories require. The DPO theory does not transfer wholesale.

rational content is held constant. Units may dissolve, be created, be merged, or be split. Command relationships that existed before may no longer exist after, and new ones may appear. The assignment function assign : $V \to \mathbb{M}_L^\star$ must adjust accordingly. This much is clear. What is not clear is: what conditions make an organizational transformation *valid*? What is the organizational analogue of the gluing condition? There is no satisfactory answer to this question, and the present work does not pretend to supply one. In place of a formal validity condition, we examine the cases where organizational transformation has been historically decisive—cases that illuminate the phenomenon and constrain whatever answer eventually emerges.

The comparative politics literature has produced partial answers that, though they fall short of a formal validity condition, constrain the search space for one. Cohen (2002) argues that civilian control produces more capable armed forces when it is intrusive but not directive: civilian leaders who press their military professionals with demanding questions and insist on strategic outcomes—without prescribing operational methods—produce organizations more capable of adaptation than those who either defer entirely or micromanage. In our language, this suggests that valid organizational transformation requires the civilian principal to constrain the destination in the convertibility preorder (which force should the reorganization reach?) without prescribing the morphism sequence used to get there. Avant (1994) argues from a complementary direction: organizational form shapes combat behavior in ways that are systematically tied to the institutional environment—the distribution of information, the incentive structures of officers, the principal–agent relationships that constitute military authority. Together, these arguments suggest that a validity condition for organizational transformation may need to be *relational* rather than intrinsic: not a property of the new partial order $(V', \leqslant'_{cc})$ in isolation, but a property of its fit with the informational and institutional environment in which it must operate. Formalizing that relational condition is the open problem.

The Roman system. The Greek phalanx and the Roman legion fielded broadly similar raw materials: armored infantry equipped with shields, swords, and spears. The transformation from one to the other was overwhelmingly organizational. Where the phalanx was a single deep formation—a monolithic molecule, in our language—the legion decomposed into centuries, maniples, and cohorts, each capable of semi-independent maneuver. The partial order $(V, \leqslant_{cc})$ changed from a flat hierarchy (one general commands one phalanx) to a deep, articulated tree. The molecular configurations barely changed; the wiring diagram was revolutionized.

The military consequences were decisive. A phalanx is devastatingly effective on flat ground and useless on broken terrain; the legion's articulated structure let its subordinate units maneuver independently, flow around obstacles, and exploit gaps. This is what van Creveld means when he argues that the organizational layer is not merely an administrative convenience but a *cognitive* structure: the command hierarchy determines how information flows, and hence which matches the Force-Maker and her subordinates can even *see* (van Creveld, 1985). A flat hierarchy with few intermediate levels pushes all match-finding to the top; a deep hierarchy with capable intermediate commanders distributes it. The transformation from phalanx to legion was therefore not just a rearrangement of boxes on a chart but a reconfiguration of the force's epistemic capacity.

The Napoleonic corps. The invention of the corps d'armée in the late eighteenth century was perhaps the purest organizational innovation in modern military history (Chandler, 1966). The corps was a miniature army: infantry, cavalry, and artillery combined under a single commander, capable of marching independently and fighting unsupported for a day. Before the corps system, armies marched as a single column; after it, they marched in parallel columns that could converge on the battlefield. Again, the configurational content—muskets, cannon, horses—was essentially unchanged. What changed was the partial order: a new level was inserted into the command hierarchy, and the non-command relations R were enriched with lateral coordination bonds between corps that had no analog in the old structure.

The insertion of an organizational level is exactly the kind of change that is hard to capture cleanly with either elementary edits (what is the edit that inserts a level?) or DPO (the gluing condition for partial orders is unclear). It is a transformation of the *shape* of the force—of how many levels of command exist, and how authority flows between them. The before and after do not merely differ in content; they differ in structure.

The Zulu ιμπονδο ζενκομο. The examples above are institutional transformations: they take months or years and are executed by political authorities. But organizational restructuring also happens on the battlefield, in minutes, when a force executes a pre-drilled maneuver.

The Zulu *impondo zenkomo*—the "beast's horns"—is perhaps the purest example (Knight, 1995). Shaka kaSenzangakhona developed the formation in the early nineteenth century, and it remained the standard Zulu tactical doctrine through the Anglo-Zulu War of 1879. The maneuver has four elements: the *isifuba* (chest), which pins the enemy with a frontal assault; the left and right *izimpondo* (horns), which sweep wide to envelop; and the *umuva* (loins), a reserve force that sits behind the chest, often facing away from the battle to prevent premature engagement.

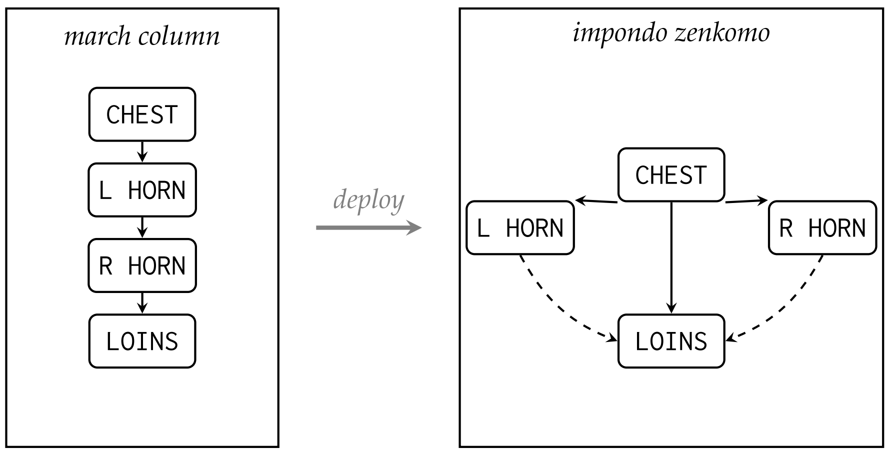

Figure 2.3: *An organizational transformation executed on the battlefield. The four units—chest, left horn, right horn, loins—are the same warriors with the same weapons. Only the command structure (solid arrows) and coordination bonds (dashed arrows) change. The transformation is a pre-installed rewrite rule: drilled until the impi can execute it autonomously on the battlefield.*

The configurational content of each element is unchanged: the same warriors carry the same *iklwa* (stabbing spear) and *isihlangu* (cowhide shield) in the march column and in the deployed formation. What changes is entirely organizational: the partial order $(V, \leqslant_{cc})$ is restructured from a linear chain to a branching tree, and new non-command relations R—lateral bonds of coordination between the horns and the reserve—are added.

The buffalo horns is an organizational rewrite rule, and it illustrates two things. First, organizational transformations can be *tactical*, not just institutional—executed in minutes on the battlefield, not over months in a ministry. Second, it demonstrates the concept of installed rules from the preceding section. The *impondo zenkomo* was drilled relentlessly into every impi. The training—catalyzed by experienced indunas—installed the rewrite rule so deeply that the units could execute it without further instruction. At the Battle of Isandlwana in 1879, the Zulu force executed the maneuver against a British column despite the death of the commanding induna early in the engagement: the rule had been installed in the units themselves, not held by a single catalyst (Knight, 1995).

Goldwater–Nichols. The Goldwater–Nichols Department of Defense Reorganization Act of 1986 rewired the command hierarchy of the United States military from service-centric to joint, and represents the most consequential organizational transformation of a major power's military in the postwar era. Before 1986, the chain of command ran through the service chiefs, who simultaneously managed their services and exercised operational authority over deployed forces; theater commanders had limited authority over forces assigned from other services, and joint operations required coordination among officials with no formal obligation to defer to one another. After 1986, operational command ran through regional and functional combatant commanders assigned forces from all services and reporting directly to the civilian leadership; the service chiefs were removed from the operational chain of command entirely, their role confined to organizing, training, and equipping. The same soldiers, the same equipment, the same bases—a fundamentally different organizational graph, with the command partial order $(V, \leqslant_{\mathrm{cc}})$ rewired from a set of parallel service trees to cross-service functional structures. van Creveld's sequence-sensitivity argument illuminates why the transformation took three decades after Vietnam: the new organizational topology required prior catalytic investments—joint professional military education, interservice exercises, years of personnel rotations—that had to be in place for the new command relations to function; the same rewrite applied in 1955 would have produced a structurally different force. The first Gulf War, five years after the act, provided the test under combat conditions, and the joint operations it enabled were decisive (Rosen, 1991): Desert Storm was not merely the first war fought under the new command structure, but the demonstration that the catalytic investments had been sufficient to activate it.

These examples span a spectrum from battlefield maneuver (the buffalo horns, minutes) to peacetime legislation (Goldwater–Nichols, years). They illustrate a structural fact: a structured force has two layers that can change independently. And they reveal what makes organizational transformation hard to formalize.

In the configurational case, the gluing condition does real work: it tells you, at the moment a rewrite rule is applied, whether the application is coherent. The surrounding structure either protects itself (the condition holds) or refuses the transformation (the condition fails). The validity is algebraic.

In the organizational case, no analogous condition is known. The question of when a reorganization is "valid"—when it produces a coherent structured force, when its interaction with the configurational layer is well-defined—depends on institutional, political, and cognitive constraints that resist algebraic capture. van Creveld (1985) argued at length that the sequence in which command structures evolve matters enormously: the Prussian general staff system could not have been installed after the introduction of radio communications in the same way as before, because the institutional context had changed. Sequence sensitivity is not noise; it is signal. But capturing it formally requires a validity condition we do not yet have.

What the category **Force** *can* say, even without a complete formalism for organizational transformation, is this: whatever these morphisms are, they compose. The Zulu deployment composes with subsequent configurational rewrites (the warriors still fight). The Goldwater–Nichols reorganization composes with subsequent training investments. The category axioms are secure even where the internal structure of the morphisms is not. This is the value of the categorical frame: it tells us what structure we are committed to, without requiring us to have solved every implementation question.

2.5 The Category of Force

Having examined four accounts of what force transformations might be—elementary edits, configurational rewrites, catalytic transformations, organizational restructuring—we now step back from the accounts. What follows are properties of **Force** that hold regardless of which specific account of the morphisms one adopts; the building phase is complete, and we forget the scaffolding.

We have the objects: $\mathbb{F}^\star(\mathbb{M}_L^\star)$, the structured forces of Chapter 1. We have the morphisms: force transformations, built from configurational rewrites and organizational restructurings as described above. We verify that these assemble into a category—and then examine what the categorical structure delivers.

2.13 Definition (The Category Force)

The category of force, **Force**, *is defined as follows:*

1. ***Objects.*** *The members of* $\mathbb{F}^\star(\mathbb{M}_L^\star)$*—all structured forces.*
2. ***Morphisms.*** *For each pair of structured forces* F_1, F_2*, the morphisms* **Force**(F_1, F_2) *are the force transformations from* F_1 *to* F_2*: finite sequences*

$$F_1 = G_0 \xrightarrow{\phi_1} G_1 \xrightarrow{\phi_2} \cdots \xrightarrow{\phi_n} G_n = F_2,$$

where each ϕ_i is either a configurational rewrite step or an organizational transformation, and each intermediate $G_j \in \mathbb{F}^\star(\mathbb{M}_L^\star)$.

3. ***Identity.*** *For each $F \in \mathbb{F}^\star(\mathbb{M}_L^\star)$, the identity morphism id_F is the empty sequence.*

4. ***Composition.*** *Given $\phi \in \mathbf{Force}(F_1, F_2)$ and $\psi \in \mathbf{Force}(F_2, F_3)$, the composite $\psi \circ \phi \in \mathbf{Force}(F_1, F_3)$ is the concatenation of the two sequences.*

2.14 Proposition

Force *is a category.*

Proof. Associativity: concatenation of finite sequences is associative. Unitality: concatenating the empty sequence with any sequence ϕ yields ϕ. ∎

The category **Force** inherits a monoidal structure from the disjoint union operation of Chapter 1: forces can be combined in parallel, and the combined force can be transformed by transforming each component independently.

2.15 Definition (Monoidal Product)

The monoidal product on **Force** *is defined by:*

1. ***On objects.*** *$F_1 \oplus F_2$ is the disjoint union of structured forces.*

2. ***On morphisms.*** *Given $\phi : F_1 \to F_1'$ and $\psi : F_2 \to F_2'$, the product $\phi \oplus \psi : F_1 \oplus F_2 \to F_1' \oplus F_2'$ applies ϕ to the F_1 component and ψ to the F_2 component independently.*

3. ***Unit.*** *The void force $\nvdash_{\mathbb{M}_L}$ is the monoidal unit.*

2.16 Proposition

$(\mathbf{Force}, \oplus, \nvdash_{\mathbb{M}_L})$ *is a symmetric monoidal category.*

Proof. Symmetry: $F_1 \oplus F_2 \cong F_2 \oplus F_1$ by relabeling the disjoint components. Associativity: $(F_1 \oplus F_2) \oplus F_3 \cong F_1 \oplus (F_2 \oplus F_3)$ by reassociation of the disjoint union. The monoidal product is well-defined on morphisms because configurational rewrites and organizational transformations applied to one disjoint component do not affect the other. Coherence follows from the strictness of the disjoint union on finite structures. ■

Force is a symmetric monoidal category: a resource theory in the sense of Coecke, Fritz and Spekkens (2016). Its objects are the resources (structured forces), its morphisms are the transformations between them, and its monoidal product is the parallel composition of independent forces.

Chapter 1 gave the Force-Maker an ontology of what force *is*—the space $\mathbb{F}^\star(\mathbb{M}_L^\star)$: a photograph album, complete but still. The category **Force** is a *road map* superimposed on the album. Every structured force is a point in the space; every morphism is a directed path between points. The Force-Maker can now ask, and answer, questions that the static ontology could not: *Can my force become that one?* (Does a morphism exist?) *Can I transform both divisions simultaneously?* (Does the monoidal product apply?) *If I do A before B, do I get the same result as B before A?* (Are the two steps parallel independent?) *Is this transformation structurally impossible, or merely expensive?* (Does no morphism exist, or does one exist but at high cost?)

The restricted category. Fix a rule library $\mathcal{R}$—the set of transformation steps the Force-Maker's doctrine, training, and supply chain support.

2.17 Definition (Rule Library and Restricted Category)

A rule library $\mathcal{R}$ *is a set of rewrite rules (spans in* $\mathcal{G}$*) and organizational transformation templates. The* restricted category $\mathbf{Force}_{\mathcal{R}}$ *has the same objects as* **Force** *and those morphisms whose every transformation step uses a rule from* $\mathcal{R}$*.*

The universal library $\mathcal{R}_{\text{all}}$ contains every possible rule. Under $\mathcal{R}_{\text{all}}$, no transformation is excluded. A bounded library $\mathcal{R} \subsetneq \mathcal{R}_{\text{all}}$ reflects real constraints: the rules the Force-Maker's institutions have developed, the doctrines her training infrastructure has installed, the transformations her supply chain can support.

When Frederick drills the oblique order into the Prussian army, he is not only changing the force F—he is changing the rule library: $\mathcal{R} \to \mathcal{R} \cup \{p_{\text{oblique}}\}$.

Paths that did not exist under the old library now exist under the new one. The Force-Maker's planning space has expanded.[9]

Doctrine, in this framing, is at least in part the choice of $\mathcal{R}$: a judgment about which transformations the force should be prepared to execute, institutionalized as rules that the training establishment installs and the supply chain supports.

The Force-Maker's reachability cone. Fix a structured force F. The *coslice category* $F/\mathbf{Force}$ has objects (G, ϕ) where $\phi : F \to G$ is a specific transformation path from F to G, and morphisms $(G, \phi) \to (H, \psi)$ are continuations $\alpha : G \to H$ with $\psi = \alpha \circ \phi$. This is the Force-Maker's *reachability cone*: every force she can build from what she has, together with every path to get there.[10]

The dual construction, the *slice category* $\mathbf{Force}/F$, has objects (G, ϕ) where $\phi : G \to F$—forces from which F is reachable—and is the Force-Maker's *ancestry cone*. Together, the coslice and slice at F describe her complete strategic position: where she came from, where she is, and where she can go.

2.18 Proposition (Rule Monotonicity of the Coslice)

Let $\mathcal{R} \subseteq \mathcal{R}'$ be rule libraries. For any structured force F, the inclusion functor $\mathbf{Force}_{\mathcal{R}} \hookrightarrow \mathbf{Force}_{\mathcal{R}'}$ induces a faithful functor $F/\mathbf{Force}_{\mathcal{R}} \to F/\mathbf{Force}_{\mathcal{R}'}$. In particular, every force reachable from F under $\mathcal{R}$ is reachable under $\mathcal{R}'$.

Proof. An object (G, ϕ) of $F/\mathbf{Force}_{\mathcal{R}}$ has $\phi : F \to G$ a morphism in $\mathbf{Force}_{\mathcal{R}}$. Since $\mathcal{R} \subseteq \mathcal{R}'$, every $\mathcal{R}$-step is an $\mathcal{R}'$-step, so ϕ is also a morphism in $\mathbf{Force}_{\mathcal{R}'}$. Faithfulness: distinct morphisms in $F/\mathbf{Force}_{\mathcal{R}}$ remain distinct, since they are the same underlying morphisms in $\mathbf{Force}$. ∎

2.19 Corollary (Reinforcement Monotonicity of the Coslice)

For any structured forces F, H and rule library $\mathcal{R}$, the monoidal product $- \oplus H$ induces a faithful functor $F/\mathbf{Force}_{\mathcal{R}} \to (F \oplus H)/\mathbf{Force}_{\mathcal{R}}$ sending (G, ϕ) to $(G \oplus H, \phi \oplus \mathrm{id}_H)$.

[9] Strictly: the training that installs the rule does change F, because training forges bonds in the configuration graph. But the conceptual point is that the salient consequence of training is the expansion of $\mathcal{R}$. A soldier with trained bonds but no authorization to use them has been modified without expanding $\mathcal{R}$. A soldier authorized to use a maneuver he has not trained has $\mathcal{R}$ expanded without modifying F. The full picture requires both.

[10] The coslice category is standard in category theory; see Aluffi (2009, Chapter I) for a textbook treatment and Coecke, Fritz and Spekkens (2016, Section 3.2) for its use in resource theories.

Proof. By Proposition 2.22, $\phi \oplus \mathrm{id}_H : F \oplus H \to G \oplus H$ is a morphism in $\mathbf{Force}_{\mathcal{R}}$, so $(G \oplus H, \phi \oplus \mathrm{id}_H)$ is a well-defined coslice object. Faithfulness holds because $\alpha_1 \neq \alpha_2$ implies $\alpha_1 \oplus \mathrm{id}_H \neq \alpha_2 \oplus \mathrm{id}_H$. ■

Proposition 2.18 is the formal content of Frederick's investment in doctrinal innovation: installing the oblique order transitioned his rule library from $\mathcal{R}$ to $\mathcal{R} \cup \{p_{\text{oblique}}\}$, and the proposition guarantees that his coslice grew weakly—and grew strictly whenever the new rule opened transformation paths that genuinely did not exist before. Proposition 2.18 also encodes Posen's insight from the other direction: a doctrine that forbids certain rules contracts the coslice. Doctrinal rigidity is not merely a cultural observation; it is a provable reduction in the Force-Maker's reachable set (Posen, 1984).

2.20 Remark (Factorization through Elementary Edits)

The four elementary edits of Section 2.1—+Node, −Node, +Bond, −Bond—are not merely the most conservative account of force transformations. They are universal: every morphism in **Force***, under any of the four accounts examined in this chapter, can be expressed as a finite composition of elementary edits. A configurational rewrite decomposes into deletions of the left-hand pattern and additions of the right-hand pattern; a catalytic transformation decomposes into the same sequence while preserving the catalyst; an organizational transformation decomposes into deletions and additions of command-structure bonds. The elementary edits are a generating set for the morphisms of* **Force**.[11]

This factorization has a structural consequence for Chapter 3. If every morphism factors through elementary edits, then a cost assigned to elementary edits extends to every morphism by composition: the cost of a morphism is the infimum over all factorizations into elementary edits of the sum of the costs of those edits. The infimum is taken because some paths are cheaper than others, and the Force-Maker always prefers the least expensive route. The conservative first account of Section 2.1 thus turns out to be the foundational one: it is the level at which cost can be introduced, and all higher accounts inherit their costs from it.

[11] More precisely: the subcategory generated by the four elementary-edit morphisms is all of **Force**. The factorization is not unique—there are typically many sequences of elementary edits realizing the same net transformation—but existence is guaranteed by the adhesive structure of $\mathcal{G}$.

2.6 The Convertibility Preorder and Military Genius

2.6.1 The Convertibility Preorder

The category **Force** gives the Force-Maker a map of what is possible. If there exists a morphism $F_1 \to F_2$ in **Force**, then F_1 can be transformed into F_2. This induces a natural ordering.

2.21 Definition (Convertibility Preorder)

For $F_1, F_2 \in \mathbb{F}^\star(\mathbb{M}_L^\star)$, write $F_1 \precsim F_2$ if there exists a morphism $F_1 \to F_2$ in **Force**.

The relation $\precsim$ is a preorder: it is reflexive (the identity morphism witnesses $F \precsim F$) and transitive (composition of morphisms witnesses the chain). It is not in general antisymmetric: two forces may be mutually convertible ($F_1 \precsim F_2$ and $F_2 \precsim F_1$) without being identical. The Force-Maker who can reorganize her division into a brigade and back again has mutual convertibility, but the two organizational forms are distinct members of $\mathbb{F}^\star(\mathbb{M}_L^\star)$.

The preorder interacts with the monoidal product in a way that reflects a basic feature of parallel forces: adding an independent element to both sides of a comparison does not disturb it. The following proposition records this formally.

2.22 Proposition (Monoidal Monotonicity)

If $F_1 \precsim F_2$, then $F_1 \oplus G \precsim F_2 \oplus G$ for any structured force G.

Proof. Let $\phi : F_1 \to F_2$ be a morphism in **Force**. Then $\phi \oplus \mathrm{id}_G : F_1 \oplus G \to F_2 \oplus G$ is a morphism in **Force**: apply ϕ to the F_1 component and do nothing to the G component. ∎

Convertibility is robust to augmentation: if your infantry division can be transformed into a mechanized division, then your infantry division *plus an allied air wing* can be transformed into a mechanized division *plus the same allied air wing*. Adding resources to the picture does not block transformations that were already available.

2.23 Definition (Resource Monotone)

A function $\mu : \mathbb{F}^\star(\mathbb{M}_L^\star) \to \mathbb{R}$ is a resource monotone *if $F_1 \precsim F_2$ implies $\mu(F_1) \leq \mu(F_2)$.*

The connected case: unrestricted rules. Under the unrestricted category **Force** = $\mathbf{Force}_{\mathcal{R}_{\text{all}}}$, the void force $\nvdash_{\mathbb{M}_L}$ is a *universal waypoint*: any force F can be dismantled to $\nvdash_{\mathbb{M}_L}$ by deleting its elements one at a time, and any target force G can be assembled from $\nvdash_{\mathbb{M}_L}$ by creating its elements one at a time. Composing, the convertibility preorder is trivially total: $F \precsim G$ and $G \precsim F$ for all F, G.

The totality is the problem—because the preorder, being total, is trivial. It tells the Force-Maker that everything is reachable from everything, which is true and useless. An infantry army can become an air force, yes—via the destruction of everything it is, followed by the creation of everything it is not. The interesting question was never whether the path exists but what the path costs, and that is the subject of Chapter 3.

The restricted case: bounded rule libraries. The more substantively interesting case arises in $\mathbf{Force}_{\mathcal{R}}$ for a proper restriction $\mathcal{R} \subsetneq \mathcal{R}_{\text{all}}$. Real militaries cannot manufacture fighter jets from infantrymen, cannot create trained pilots from nothing, cannot delete carrier groups and spontaneously generate armored divisions.

2.24 Proposition (Non-Totality Under Restricted Rules)
There exist rule libraries $\mathcal{R}$ and structured forces $F_1, F_2 \in \mathbb{F}^\star(\mathbb{M}_L^\star)$ such that $F_1 \not\precsim_{\mathcal{R}} F_2$ and $F_2 \not\precsim_{\mathcal{R}} F_1$.

Proof. By construction. Let $L = \{\alpha, \beta\}$ with $\alpha \neq \beta$, and let $\mathcal{R}$ contain only rules that preserve node labels—no rule relabels a node from α to β or vice versa, and no rule creates nodes *ex nihilo* or deletes them entirely. Let F_1 be a structured force whose configuration consists entirely of α-labeled nodes, and let F_2 be a structured force whose configuration consists entirely of β-labeled nodes. No finite sequence of rules from $\mathcal{R}$ can transform F_1 into F_2: every intermediate step preserves node labels, so the α-nodes are never replaced by β-nodes. Hence no morphism $F_1 \to F_2$ exists in $\mathbf{Force}_{\mathcal{R}}$, and symmetrically, no morphism $F_2 \to F_1$ exists. ∎

In the restricted case, non-totality has a direct consequence for the measurement of force.

2.25 Proposition (Monotones Separate Incomparable Forces)
Let F_1, F_2 be incomparable ($F_1 \not\precsim F_2$ and $F_2 \not\precsim F_1$). Then there is no resource monotone μ that is both faithful *(reflecting the preorder: $\mu(A) \leq \mu(B)$ implies $A \precsim B$) and total.*

In particular, for any two monotones μ, ν that together detect the incomparability, one has $\mu(F_1) > \mu(F_2)$ and $\nu(F_1) < \nu(F_2)$ (or vice versa).

Proof. If μ were both faithful and total, then $\mu(F_1) \leq \mu(F_2)$ would imply $F_1 \precsim F_2$, and $\mu(F_2) \leq \mu(F_1)$ would imply $F_2 \precsim F_1$. Since exactly one of $\mu(F_1) \leq \mu(F_2)$ and $\mu(F_2) \leq \mu(F_1)$ must hold (totality of $\leq$ on $\mathbb{R}$), we would have $F_1 \precsim F_2$ or $F_2 \precsim F_1$—contradicting incomparability. For the second claim: if both $\mu(F_1) \leq \mu(F_2)$ and $\nu(F_1) \leq \nu(F_2)$, neither detects the incomparability. To capture the structure of the preorder at (F_1, F_2), we need at least two monotones that disagree on the ranking. ∎

This is the formal reason that force cannot be measured by a single number. A force composed entirely of infantry and a force composed entirely of aircraft are incomparable under any realistic rule library. Any single number that ranks them contradicts the convertibility structure.

This connects directly to a long-standing debate in the international relations literature. Mearsheimer's (2001) army-centric approach to measuring power and Biddle's (2004) force-employment thesis are each, in our language, proposing different monotones—different real-valued compressions of the same underlying preorder. Proposition 2.25 says they will necessarily disagree on incomparable forces—and that the disagreement is structural, not a defect of measurement. The right response is not to search for the "correct" single monotone but to accept that the convertibility preorder is the fundamental object.

2.26 Example (Three Concrete Monotones)

Define:

1. $\mu_{\text{node}}(F) = |V(F)|$*: the total number of element nodes across all configurations allocated to units of* F*.*

2. $\mu_{\text{bond}}(F) = |E(F)|$*: the total number of bonds across all configurations allocated to units of* F*.*

3. $\mu_{\text{depth}}(F)$*: the length of the longest chain in the command hierarchy* $(V, \leqslant_{\text{cc}})$ *of* F*.*

These monotones separate: a freshly conscripted army of ten thousand men with almost no training bonds has high μ_{node} *and low* μ_{bond}*; a veteran special-operations unit of* *two hundred operators bonded by years of shared training has low μ_{node} and high μ_{bond}. Neither monotone is "wrong"—they are measuring genuinely different dimensions of the same space.*

2.6.2 Free and Costly Transformations

Every resource theory distinguishes *free* operations—those available without cost—from *costly* ones (Coecke, Fritz and Spekkens, 2016). In our setting, the natural candidate for the free subcategory is *reversible* transformations: morphisms $\phi : F_1 \to F_2$ for which there exists a morphism $\psi : F_2 \to F_1$ such that $\psi \circ \phi$ and $\phi \circ \psi$ are both identity-equivalent.

The Zulu *impondo zenkomo* is a free operation in this sense: the march column becomes the buffalo horns, and the buffalo horns can become the march column again, with no material consumed. A division reorganized into brigades can be reconstituted as a division. The reversible morphisms of **Force** form a subcategory with a special property.

2.27 Proposition (The Rearrangement Groupoid)

The reversible morphisms of **Force***—those $\phi : F_1 \to F_2$ for which there exists $\psi : F_2 \to F_1$ with $\psi \circ \phi = \mathrm{id}_{F_1}$ and $\phi \circ \psi = \mathrm{id}_{F_2}$—form a groupoid* $\mathbf{Force}^{\times} \subseteq \mathbf{Force}$. *We call* $\mathbf{Force}^{\times}$ *the* rearrangement groupoid *of force.*

Proof. Closure: the composition of reversible morphisms is reversible (compose their inverses in opposite order). Identities: id_F is reversible (its own inverse). Inverses: by definition of reversibility. ■

The rearrangement groupoid captures the transformations that shuffle material and organizational structure without net creation or destruction of resources: the reorganizations, deployments, and formations that cost effort but leave the force's inventory unchanged.

The rearrangement groupoid of a fixed force F—the set of all reversible morphisms with domain or codomain F—is the formal expression of that force's *tactical freedom*: the range of configurations it can adopt and un-adopt without expenditure. Every orbit under $\mathbf{Force}^{\times}$ is a collection of structurally equivalent forces, each a different organizational expression of the same inventory.

Consider the Zulu impi that Shaka trained. The march column and the *impondo zenkomo* are two points in the same orbit: the same warriors carry the same *iklwa* and *isihlangu* in both configurations, and moving between them consumes nothing. No warrior is lost, no weapon expended, no bond permanently severed. The transformation is free precisely because the rewrite rule was installed—drilled until the units could execute it without further instruction. An impi that had never trained the buffalo horns could not access this orbit; installation is what puts a transformation in $\mathbf{Force}^{\times}$.

This separates two dimensions of military capability that formal models typically conflate. The convertibility preorder $\precsim$ measures where a force stands in terms of *costly* reachability: what can it become, by expending resources? The rearrangement groupoid measures what the force can express for free: how many different configurations can it adopt without expenditure? A mass conscript army and a small professional force may occupy the same position in the preorder—each can, at high cost, become something like the other—while having very different groupoids. The conscripts' groupoid may be thin: few installed formation rules, limited tactical flexibility, little freedom of maneuver within a fixed inventory. The professionals' groupoid may be rich: a repertoire of practiced maneuvers, each freely reversible, each executable without the commanding officer present. This is the structural content of what the military effectiveness literature calls the distinction between quantity and quality: quality is, in part, a property of the rearrangement groupoid. Rosen's (1991) observation that military innovation is driven by officers willing to invest in new operational concepts—not just new equipment—is, in our language, an observation about which forces are willing to expand their rearrangement groupoids.

Transformations outside $\mathbf{Force}^{\times}$ are irreversible, and irreversibility is the structural reason that cost theory must be asymmetric. The US Navy's commitment in the 1920s to developing carrier aviation—converting yard capacity, officer career paths, and doctrinal energy to a platform with no proven combat record—is an example: the industrial base and institutional investment built around carrier operations could not simply be reconverted to battleship production once the commitment had been made. Rosen (1991) traces this commitment and notes that it required officers willing to build careers around a vision of warfare that might never materialize. The cost was asymmetric in a precise sense: committing to the carrier program was cheaper than reversing the commitment after the institutional investment had accumulated. Firing ammunition is similarly irreversible—the `AMMO` node is consumed and cannot be spontaneously recreated—while building a new armored division is irreversible in the opposite direction. It is precisely these non-free morphisms, and the asymmetry of their costs, that Chapter 3 must price.

2.6.3 Three Kinds of Military Genius

The categorical frame makes precise a distinction that has always lived in the military theory literature, usually under vaguer names.

A rewrite rule $p = (L \xleftarrow{\ell} K \xrightarrow{r} R)$ does nothing until someone finds a match $m : L \to G$. The Force-Maker who looks at her current force G and *sees* the pattern L embedded in it—who recognizes that the current situation admits a known transformation—is exercising one kind of genius. Call it *match-finding*: the ability to perceive opportunities for transformation in the fog of the present configuration. This is Clausewitz's *coup d'oeil*, the "inward eye" that takes in a situation at a glance and knows what can be done with it.

The Force-Maker who conceives a *new* rewrite rule—who envisions a span $L \xleftarrow{\ell} K \xrightarrow{r} R$ that nobody has written down—is exercising a second, rarer faculty. Call it *scheme-invention*: the ability to design the interface through which the current force can be rewritten into a configuration that does not yet exist.

Rosen's (1991) military innovators—the architects of carrier aviation, armored warfare, submarine operations—are scheme-inventors. They did not merely find matches for existing rules; they wrote new rules. The British officers who conceived the tank in 1915 saw a new span: L = infantry pinned in trenches, K = the soldiers themselves (preserved), R = soldiers mounted in armored vehicles with caterpillar tracks. The interface preserved what mattered (trained men) and replaced what did not work (exposure to machine-gun fire) with something that had never existed before.

The categorical framework gives each faculty a precise home: match-finding is a problem of graph homomorphism (finding $m : L \to G$), while scheme-invention is a problem of span design (authoring $L \xleftarrow{\ell} K \xrightarrow{r} R$). The two faculties are genuinely distinct, and the literature on generalship has often conflated them; the categorical frame makes the structural difference explicit.

But Frederick's example from Section 2.3 suggests a third faculty that is neither match-finding nor scheme-invention but may matter more than either: the judgment of which rules to *install*. A Force-Maker with a hundred brilliant schemes that her force cannot execute is no better off than one with no schemes at all. The deepest strategic question is not "what transformation is possible?" or "what transformation can I imagine?" but "what transformations should my force be able to execute without me?" This is the question of *preparedness*—of investment in catalytic depth, training infrastructure, and the institutional recursion that sustains them—and it is the question that connects the grammar of this chapter to the costs of the next.

2.7 Looking Ahead

The category **Force** characterizes which structured forces are transformable into which, and what compositional structure those transformations obey. It is silent on the question of cost: morphisms exist or they do not, but the category assigns no price to them.

McNeill (1982) traced the interaction of technological change and military organization over a millennium. In our language, his narrative is a sequence of configurational and organizational rewrites applied to the global force graph over centuries, each one shifting the convertibility preorder of every force in the system. The stirrup is a configurational rewrite that installs the heavy cavalryman as a new node type—and the command hierarchy had to change to exploit it: the mounted knight required a different logistical and social structure than the Roman legionary. Gunpowder is a +Node rewrite at the level of weapon technology that dissolves the advantage of fortification and armor, forcing organizational rewrites that eventually produce the massed infantry battalion. The railroad is a +Bond rewrite on a continental scale, connecting interior forces to frontier deployment zones in ways that would not close for fifty years without it, and that demanded new staff functions—the railway officer, the mobilization schedule—to manage the new bond's complexity. The internal combustion engine is a configurational rewrite whose full organizational consequence—combined arms, air-ground coordination, the operational art that the Germans formalized as *Auftragstaktik*—took decades to install as rules into the relevant forces. In each case, the convertibility preorder shifted: forces that had been adequate no longer were, and forces that could apply the new rule expanded their reachability cones while those that could not fell in the ordering. This is the formal content of what the "military revolution" literature calls obsolescence: a single configurational rule whose adoption forces a preorder shift so large that non-adopters cannot compensate by any sequence of rules remaining in their libraries (Parker, 1988). The category **Force** gives this narrative a formal home.

Chapter 3 assigns a cost to every morphism in **Force**: a real-valued expenditure of time, materiel, institutional capital, and political will that the Force-Maker must spend to effect a transformation. The cost structure respects composition—a longer path is at least as expensive as a shorter one—and it is asymmetric: building capability typically costs more than destroying it, though the reverse holds in contexts of rapid technological obsolescence. The infimum cost over all morphisms connecting two structured forces defines a natural distance between force postures, and it is that distance—not any geometric notion—that captures what strategists mean when they ask how far apart two forces are. A force posture that is formally reachable from the current one may be so expensive to reach that it is practically unreachable; cost theory makes that distinction precise.

By Remark 2.20, this cost can always be computed by decomposing morphisms into elementary edits and taking the infimum over all such decompositions: the cheapest route from one force to another is the cheapest sequence of node and bond additions and deletions that effects the net transformation. The Lawvere metric generated by this infimum cost gives $\mathbb{F}^\star(\mathbb{M}_L^\star)$ a topology—the topology of transformation cost—in which two forces are close when the cheapest transformation between them is cheap. It is in this topology that continuity of the force function m will finally mean something: not merely that the function exists (it does, by countability) but that it respects the structure it compresses—the structure we just built.

Chapter 3

Force is Costly

This retribution, which has a geometrical rigor, which operates automatically to penalize the abuse of force, was the main subject of Greek thought. It is the soul of the epic. Under the name of Nemesis, it functions as the mainspring of Aeschylus's tragedies. To the Pythagoreans, to Socrates and Plato, it was the jumping-off point of speculation upon the nature of man and the universe. Wherever Hellenism has penetrated, we find the idea of it familiar. In Oriental countries which are steeped in Buddhism, it is perhaps this Greek idea that has lived on under the name of Kharma. The Occident, however, has lost it, and no longer even has a word to express it in any of its languages: conceptions of limit, measure, equilibrium, which ought to determine the conduct of life are, in the West, restricted to a servile function in the vocabulary of technics. We are only geometricians of matter; the Greeks were, first of all, geometricians in their apprenticeship to virtue.

Simone Weil, "The Iliad, or the Poem of Force"

In October 1906, HMS *Dreadnought* was commissioned into the Royal Navy, rendering every capital ship afloat—including most of Britain's own fleet—obsolete overnight. The naval race that followed was, on the surface, a story about morphisms: both Britain and Germany possessed the rules necessary to build dreadnought-class battleships, and from 1906 to 1914 both sides exercised those rules continuously. Chapter 2 analyzed the grammar of that competition—the DPO rewrites, the organizational rules, the catalytic infrastructure of shipyard and design bureau. But what the category **Force** said nothing about was why Britain won. The convertibility preorder (Section 2.4) placed both navies in the same morphism-connected world; it could tell you which ships each side *could* build, not which side could build them *cheaply*. The answer lay in what

Brewer called the fiscal-military state (Brewer, 1990): a century of institutional development that gave Britain a fundamentally lower cost function—deeper credit markets, more efficient tax collection, and superior naval administration—so that identical transformation rules cost Germany structurally more to execute. The convertibility preorder said both sides could build the fleet; what it could not say was that Britain could do it at a fraction of the cost. The Lawvere metric says that, and this chapter builds it.

Chapter 1 built the space $\mathbb{F}^\star(\mathbb{M}_L^\star)$—the set of all structured forces. Chapter 2 built the category **Force**—the map of valid transformations between them. The present chapter builds the geometry: it assigns to every morphism in **Force** a *cost* valued in a mathematical structure rich enough to capture the asymmetry, non-additivity, and multidimensionality that military cost displays in practice. From these costs we derive a *distance* between forces—the cheapest path through the category—and from that distance a *topology* on the force space, in which "nearby" means "cheaply reachable" and "far" means "expensive or impossible." That topology, and the second-countability theorem that its construction delivers, is the prerequisite for everything in Chapter 4: it is the condition Debreu's representation theorem requires before a continuous real-valued force function can be guaranteed to exist.

Three Stylized Facts

Before constructing the mathematics, it is worth anchoring the inquiry in the empirical regularities that the mathematics must accommodate. Three facts about military cost recur persistently across historical cases, and they collectively rule out the structures that economic models most often reach for.

Asymmetry. At the end of the Second World War, the United States Army stood at 8.3 million soldiers. Within two years it had demobilized to roughly 591,000—a reduction of more than ninety percent accomplished at near-zero marginal cost per departing soldier, because the infrastructure of conscription, separation centers, and discharge processing had been built precisely for this purpose. The Korean War reversed the direction: rearmament from 591,000 to a wartime footing took years and billions of dollars in recruitment, training, equipment procurement, and institutional reconstitution. The cost of $F \to G$ was not the cost of $G \to F$. Federle, Meier, Müller, Mutschler, and Schularick's (2026) analysis of 150 years and 60 countries provides the most systematic quantification of this asymmetry yet assembled. A war of average intensity—casualties equal to roughly 2 percent of the war-site population—produces an output decline of nearly 10 percent relative to prewar trend, with a trough approximately four years after onset. Complete output recovery takes roughly twelve years. The mean war in their sample spans three years. The asymmetry is structural and severe: a force-maker moving a war-site economy toward a devastated war footing moves fast and cheaply relative to the return journey, even before accounting for the capital stock destruction, TFP decline, and institutional erosion that recovery must undo. Any cost language that assigns symmetric values to forward and backward transformations will misrepresent both the speed of demobilization and the expense of reconstitution—treating a structural asymmetry as if it were a bookkeeping symmetry.

Non-additivity. Van Creveld's study of military logistics (Van Creveld, 2004) reveals that Napoleon's Russian campaign was destroyed not by Russian resistance but by the compounding of march costs, supply costs, and operational costs in a way that no additive model could have predicted from the individual components. The cost of a 600-mile march in summer and the cost of campaign operations were each manageable in isolation; their sequential composition created a logistical crisis whose severity exceeded the sum of the parts, because the march consumed the draught animals and supply reserves on which the operational phase depended. Kennedy's argument about great-power overstretch (Kennedy, 1989) is the same observation at the strategic scale: the costs of maintaining a global military position compound rather than sum, because earlier commitments consume the fiscal and institutional resources on which later ones depend. An additive cost language—one where the cost of a sequence of transformations is always the sum of the individual costs—cannot represent compounding, and compounding is the dominant mode of military cost accumulation under sustained exertion.

Multidimensionality. Cappella Zielinski's analysis of war finance (Cappella Zielinski, 2016) identifies four qualitatively distinct mechanisms by which states pay for wars: taxation, borrowing, monetary creation, and asset liquidation. These are not commensurable entries in a single ledger; they impose different burdens on different parts of the political economy, they exhaust different resources, and they leave different long-run institutional residues. Stiglitz and Bilmes (Stiglitz and Bilmes, 2008) estimated the fiscal cost of the Iraq War at three trillion dollars, and immediately acknowledged that even this figure omitted institutional credibility, geopolitical standing, and human costs that resisted reduction to dollars. The problem was not that their estimate was imprecise; it was that the quantale they worked in—the real numbers under addition—could not represent a multidimensional cost as multidimensional. Federle et al. (2026) document the same point with systematic evidence: output, consumer prices, the capital stock, total factor productivity, long-term interest rates, equity returns, military spending, and military personnel all respond to war onset—at different magnitudes, with different timing, and with different exposure patterns across war sites, belligerent countries, and third parties. TFP collapses immediately and sharply at the war site but remains flat in belligerent countries that fight abroad; capital stock declines persistently everywhere but more severely on the war site; interest rates remain elevated for more than a decade after the fighting ends. These are not commensurable scalar entries that trade off at a fixed rate: a framework that adds them up destroys the information in their separate dynamics, and it is precisely that information—which costs recover quickly, which do not, and which propagate across borders via trade linkages—that strategy requires. The compression to a scalar was not a harmless simplification; it was a substantive theoretical claim that all dimensions of cost trade off at a fixed rate, a claim the evidence does not support.

These are not philosophical objections to scalar models; they are empirical regularities that any cost language adequate to military affairs must accommodate.

3.1 The Cost of a Morphism

What mathematical structure should a cost language have? The real numbers are the default answer in economic models of military spending—and for the same reasons that Chapter 0 showed they were the wrong home for force itself, they are the wrong home for its cost. The real line enforces symmetry (the cost of $A \to B$ is the same type of object as the cost of $B \to A$), additivity (costs compose by addition), and total order (any two costs can be compared)—none of which the three stylized facts above will permit. The right structure is a quantale.[1]

[1] The connection between quantales and generalized metric spaces was established by Lawvere (1973), who showed that a metric space is a category enriched over $([0, \infty], \geqslant, +)$. The accessible

Four requirements lead directly to the definition. The Force-Maker needs *comparability*: she must be able to say which of two costs is larger, even if some pairs are genuinely incomparable. She needs *composition*: if morphism ϕ costs a and morphism ψ costs b, the composite $\psi \circ \phi$ has a cost $a \otimes b$, where $\otimes$ need not be addition but must be associative. She needs *a zero*: the identity morphism—doing nothing—costs nothing, and this zero is the unit for $\otimes$. And she needs *a top*: some transformations are impossible or effectively so, and an absorbing element $\top$ must represent infinite cost so that a plan passing through an impossible step is itself impossibly expensive. These four requirements pick out a single algebraic structure.

3.1 Definition (Quantale)

A quantale *is a triple* $(Q, \leqslant, \otimes)$ *where:*

1. $(Q, \leqslant)$ *is a complete lattice: every subset of* Q *has a greatest lower bound (meet,* $\bigwedge$*) and a least upper bound (join,* $\bigvee$*), with bottom element* $\bot = \bigwedge Q$ *and top element* $\top = \bigvee Q$*.*
2. $\otimes \colon Q \times Q \to Q$ *is an associative binary operation.*
3. $\otimes$ *distributes over arbitrary joins: for all* $a \in Q$ *and* $S \subseteq Q$*,*

$$a \otimes \left(\bigvee S\right) = \bigvee_{s \in S} (a \otimes s) \qquad \text{and} \qquad \left(\bigvee S\right) \otimes a = \bigvee_{s \in S} (s \otimes a).$$

4. *The bottom element* $\bot$ *is the unit for* $\otimes$*:* $a \otimes \bot = a = \bot \otimes a$ *for all* $a \in Q$*.*

We interpret $\bot$ *as zero cost and* $\top$ *as infinite or impossible cost. The ordering* $\leqslant$ *reads "is at most as costly as."*

The definition is compact; the intuition lives in the examples, each of which corresponds to a distinct tradition of reasoning about military cost.

treatment in Fong and Spivak (2019, Chapter 2) makes the enrichment perspective available without heavy prerequisite; the full theory is developed in Kelly (1982).

The textbook quantale. The extended nonnegative reals $(\mathbb{R}_{\geqslant 0} \cup \{\infty\}, \leqslant, +)$ form a quantale. Costs are nonnegative real numbers or infinity; composition is addition; $\perp = 0$ and $\top = \infty$. This is the quantale in which Stiglitz and Bilmes (2008) worked when they estimated the cost of the Iraq War, and it is the quantale Powell (1993) implicitly inhabits when he writes the cost of military spending as a linear budget constraint. It is a perfectly good quantale—the simplest nontrivial one—but it is commutative and totally ordered, and neither property holds in general. Stiglitz and Bilmes themselves acknowledged the inadequacy, noting that their dollar figure could not capture institutional credibility or geopolitical standing; the textbook quantale made that inadequacy invisible by forcing everything onto a single axis.

The probability quantale. The unit interval $([0, 1], \geqslant, \times)$ forms a quantale, with the order reversed: $a \leqslant b$ when $a \geqslant b$ as real numbers, so that higher probability is cheaper. Composition is multiplication; $\perp = 1$ (certain success costs nothing in probability) and $\top = 0$ (impossible). This quantale is natural for sequential attrition models, where each engagement has a survival probability and the probability of surviving a sequence of engagements is the product. The Lanchester-style models that price transformations by attrition rates are implicitly working in this quantale whenever they compute survival across sequential engagements. Cold War nuclear strike planning made the structure vivid: the probability that a delivery vehicle would survive the journey from launch to target was calculated as the product of survival probabilities across successive threat rings—early warning radars, interceptor aircraft, surface-to-air missiles, terminal point defenses. The composition law exposed why layered defenses are qualitatively more effective than concentrated ones: two independent defensive layers each with fifty-percent effectiveness combine to seventy-five percent, not one hundred percent as linear addition would suggest. Each additional layer multiplies rather than adds to the attacker's cost, and the planners who sized strike packages to saturate defensive capacity were solving, in effect, an optimization problem in the probability quantale.

The time-money quantale. Let $Q = (\mathbb{R}_{\geqslant 0} \cup \{\infty\})^2$, ordered componentwise: $(t_1, m_1) \leqslant (t_2, m_2)$ iff $t_1 \leqslant t_2$ and $m_1 \leqslant m_2$. Composition is componentwise addition; $\perp = (0, 0)$ and $\top = (\infty, \infty)$. A destroyer delivered in five years for thirty billion dollars and one delivered in ten years for fifteen billion dollars satisfy $(5, 30) \nleqslant (10, 15)$ and $(10, 15) \nleqslant (5, 30)$: they are genuinely incomparable.

Kennedy's argument about great-power decline (Kennedy, 1989) is naturally housed here: the time dimension accumulates as economic base shrinks, and the time-money cost of maintaining global commitments eventually exceeds the composite resource available to the state. Brewer's fiscal-military state (Brewer, 1990) gave Britain a lower time-money cost function than its continental rivals—the same capability targets reached faster and more cheaply—because its institutional infrastructure compressed both dimensions of the cost functor simultaneously.

The resource-set quantale. Let $S = \{\text{money}, \text{time}, \text{political capital}, \text{institutional capacity}\}$ and form the power set $(2^S, \subseteq, \cup)$. The cost of a transformation is the set of resource types it consumes; composition takes the union. The unit is $\bot = \varnothing$ (free transformations consume nothing) and $\top = S$ (transformations consuming all resources are maximally costly). The ordering is by inclusion, so $\{\text{money}\}$ and $\{\text{time}\}$ are incomparable—correctly representing genuine incommensurability. This is the quantale Brewer implicitly needs: Britain's fiscal-military state enabled it to sustain the simultaneous consumption of credit and taxation without either crowding the other out, which is a claim about the structure of cost composition in the resource-set quantale, not a scalar one. Cappella Zielinski's taxonomy of war-finance mechanisms (Cappella Zielinski, 2016) maps directly onto distinct elements of S: taxation, borrowing, monetary creation, and asset liquidation are qualitatively different resource types, and her finding that they impose different political burdens is the finding that they correspond to distinct and incomparable elements of the power set.

The Boolean quantale. The two-element chain $(\{\bot, \top\}, \leqslant, \vee)$—"possible" and "impossible," with composition as join—is the degenerate quantale in which cost carries no information beyond reachability. A morphism costs $\bot$ if it exists and $\top$ if it does not; the join of two costs is $\top$ iff either step is impossible, so a sequence is possible only when every step is. Enriching **Force** over this quantale assigns the same cost $\bot$ to every morphism that exists and $\top$ to every pair of forces with no morphism between them; the cost structure records only which forces are connected, and this is exactly the convertibility preorder of Chapter 2. Chapter 2 was, in this sense, a study of the Boolean quantale—asking which transformations are achievable without asking what they cost. Force planners operating at the level of strategic net assessment often work implicitly in the Boolean quantale: the question "can this adversary threaten our carrier groups?" is Boolean, and the answer "yes, with anti-ship missiles" is a statement about the convertibility preorder. The question "what would it cost them?"—Marshall's question—requires a richer quantale, because the Boolean structure collapses all achievable transformations onto a single cost value and thereby forfeits the information that competitive strategy requires.

The choice of quantale is not a technical convenience but a theoretical decision. Powell worked in the textbook quantale; Brewer needed the resource-set quantale; Kennedy needed the time-money quantale. The quantale makes this choice explicit and consequential—the same transformation priced in different quantales yields different geometries on the same force space, and those different geometries support different strategic conclusions.

The Force-Maker's choice of quantale. The quantale is part of the Force-Maker's model, not a feature of the world. Two analysts studying the same dreadnought race may assign its costs to different quantales—one working in pounds sterling, the other in a resource-set that includes naval administrative capacity—and the geometries they derive will be genuinely distinct. This is the cost analogue of the epistemic stance from Chapter 1: just as the Force-Maker's scope and discrimination determine which forces she can see, her choice of quantale determines which dimensions of cost she can represent and compare. The construction that follows is parametric in Q; all results hold for any quantale satisfying Definition 3.1, and the examples specialize to concrete cases when illustration is needed.

3.2 The Geometry of Force

3.2.1 Costing Morphisms

The category **Force** (Definition 2.13) has objects and morphisms; the quantale Q has costs. The task is to connect them: to assign a cost in Q to every morphism in **Force** in a way that respects the compositional structure of the category.

3.2 Definition (Cost Functor)

A cost functor *on* **Force** *with values in a quantale* $(Q, \leq, \otimes)$ *is a function* $c\colon \mathrm{Mor}(\mathbf{Force}) \to Q$ *satisfying:*

1. Identity. $c(\mathrm{id}_F) = \bot$ *for all* $F \in \mathbb{F}^\star(\mathbb{M}_L^\star)$. *Doing nothing costs nothing.*

2. Subadditivity. $c(\psi \circ \phi) \leq c(\phi) \otimes c(\psi)$ *for all composable morphisms* ϕ, ψ. *The cost of a composite is at most the composed costs of its parts.*

The subadditivity condition deserves emphasis. We require $\leq$, not $=$: the cost of the composite may be *strictly less* than the composition of the individual costs, which is the formal expression of synergy. When the first transformation prepares the ground for the second—when training a soldier also partially equips her, or when an organizational restructuring creates slots that the next recruitment drive fills cheaply—the composite is cheaper than the sum of its parts. Biddle's (2004) analysis of modern force employment provides systematic evidence for this at the tactical level: the cost of achieving a given combat objective falls dramatically when armor, artillery, and infantry are employed in coordinated sequences rather than independently, because each element's transformation of the local situation—suppressing fire, fixing the defender, masking the assault—reduces the cost of the subsequent element's action. The composite transformation achieves the objective at a fraction of what any single arm would pay alone, and the reduction is structural rather than marginal: it reflects the sequential preparation of ground that the subadditivity inequality captures formally. The inequality also accommodates friction in the opposite direction: if chaining transformations introduces interference, the composite cost rises toward (but can never exceed) the composed individual costs. The subadditivity condition is not an assumption of rationality; it is a constraint on how the cost function behaves, and it holds regardless of whether the Force-Maker prices correctly.

The cost functor turns **Force** into a *Q-enriched category*: rather than hom-sets of morphisms, it has hom-objects valued in Q.[2]

3.3 Definition (Enriched Hom-Object)

Given a cost functor c on **Force**, *the* enriched hom-object *from F_1 to F_2 is:*

$$\mathbf{Force}_c(F_1, F_2) = \bigwedge_{\phi \in \mathbf{Force}(F_1, F_2)} c(\phi),$$

the infimum of the costs of all morphisms from F_1 to F_2. If $\mathbf{Force}(F_1, F_2) = \varnothing$ *(no morphism exists under the current rule library), we set* $\mathbf{Force}_c(F_1, F_2) = \top$.

[2] The insight that metric spaces are enriched categories over $([0, \infty], \geq, +)$ is due to Lawvere (1973). In that setting, the "hom-object" from x to y is the distance $d(x, y)$, and the triangle inequality is the composition law. Fong and Spivak (2019, Chapter 2) give a clear account of enriched categories in the spirit used here; the full abstract theory is in Kelly (1982).

The enriched hom-object answers the Force-Maker's most basic planning question: what is the cheapest way to transform F_1 into F_2? It collapses the entire hom-set into a single value in Q, representing the cost of the best available path and assigning infinite cost when no path exists.

This is, in its essence, the calculation that military staff work performs. Before a major operation, a staff does not merely ask whether a given force transformation is achievable—that is the preorder question of Chapter 2. It asks which path through the achievable transformation space is cheapest: which combination of existing units, capability acquisitions, and organizational adjustments reaches the required force posture at least cost in time, money, institutional capital, and political will. The German planning process that produced the Manstein Plan for the 1940 campaign in France was precisely such a search. The standard planning assumptions—a frontal assault through Belgium, the route every German staff college exercise had used since Schlieffen—represented one path through the morphism space. Manstein's proposal substituted armored penetration through the Ardennes, a route judged operationally near-impossible by the defenders and initially by most German planners as well. What it offered was a path whose composite cost, measured in strategic surprise and time to operational decision, was dramatically lower than any alternative: if it worked, the French forces in Belgium would be encircled before they could respond, at a fraction of the attrition cost of forcing the prepared Belgian and Dutch defenses. The enriched hom-object is what the planning staff computes: not a list of all possible transformations, but the infimum of their costs—the best achievable plan, after every path has been priced and the cheapest one identified.

3.4 Proposition (Enrichment Axioms)

The enriched hom-objects satisfy:

1. $\mathbf{Force}_c(F, F) = \bot$ *for all* F.
2. $\mathbf{Force}_c(F_1, F_2) \otimes \mathbf{Force}_c(F_2, F_3) \geq \mathbf{Force}_c(F_1, F_3)$ *for all* F_1, F_2, F_3.

Proof. For (1): $\mathrm{id}_F \in \mathbf{Force}(F, F)$ and $c(\mathrm{id}_F) = \bot$, so the infimum is at most $\bot$; since $\bot$ is the bottom of Q, equality holds.

For (2): for any $\phi \in \mathbf{Force}(F_1, F_2)$ and $\psi \in \mathbf{Force}(F_2, F_3)$, the composite $\psi \circ \phi \in \mathbf{Force}(F_1, F_3)$, so

$$\mathbf{Force}_c(F_1, F_3) \leq c(\psi \circ \phi) \leq c(\phi) \otimes c(\psi).$$

Taking the infimum over all ϕ and then all ψ, and using distributivity of $\otimes$ over meets in Q, yields $\mathbf{Force}_c(F_1, F_3) \leq \mathbf{Force}_c(F_1, F_2) \otimes \mathbf{Force}_c(F_2, F_3)$. ■

The second axiom is the triangle inequality for costs: going through a waypoint F_2 cannot be cheaper than the direct route from F_1 to F_3. Its derivation is important: the triangle inequality is not an assumption about the Force-Maker's rationality or planning competence. It is a theorem, following from the cost functor axioms and the completeness of the quantale. A Force-Maker who prices transformations inconsistently can still have an enriched hom-object that satisfies the triangle inequality, because the infimum construction absorbs local inconsistencies into a globally coherent distance.

Rule-library relativity. The Force-Maker does not inhabit the universal category **Force**; she inhabits the restricted category $\mathbf{Force}_{\mathcal{R}}$ for her rule library $\mathcal{R}$ (Definition 2.17). Since $\mathbf{Force}_{\mathcal{R}}(F_1, F_2) \subseteq \mathbf{Force}(F_1, F_2)$, the infimum over a smaller set is weakly larger:

$$\mathbf{Force}_{c,\mathcal{R}}(F_1, F_2) = \bigwedge_{\phi \in \mathbf{Force}_{\mathcal{R}}(F_1, F_2)} c(\phi) \geq \mathbf{Force}_c(F_1, F_2).$$

A restricted rule library means higher costs: every transformation available in the universal library but absent from $\mathcal{R}$ is a shortcut the Force-Maker cannot take, and its absence drives up the infimum. Different rule libraries yield different enrichments on the same objects, so different Force-Makers in different institutional and doctrinal contexts see different cost geometries on the same underlying space. Two NATO allies facing the same Warsaw Pact adversary with different doctrines do not merely have different preferences about how to fight—they inhabit different Lawvere metrics on the same force space. The US Army in the 1980s, with its AirLand Battle doctrine emphasizing deep attack and operational maneuver, selected a rule library that compressed the metric to specific combined-arms configurations involving deep fires and interdiction. The Bundeswehr, shaped by geography and political constraints to emphasize forward defense, selected a different rule library that compressed the metric to very different configurations—high infantry density, interlocking defensive fires, deliberate counterattacks. Both were within the same alliance and nominally facing the same adversary, but the metrics they inhabited were genuinely different, and interoperability problems between them were, in part, a consequence of attempting to compose morphisms from two different enriched categories into a coherent joint campaign plan.

3.2.2 The Lawvere Metric

The enriched hom-objects define a generalized distance on the force space.

3.5 Definition (Lawvere Metric on Force)
Given a cost functor c on $\mathbf{Force}_{\mathcal{R}}$ *valued in* $(Q, \leq, \otimes)$, *the* Lawvere metric *on* $\mathbb{F}^{\star}(\mathbb{M}_L^{\star})$ *is the function*

$$d(F_1, F_2) = \bigwedge_{\phi \in \mathbf{Force}_{\mathcal{R}}(F_1, F_2)} c(\phi),$$

with $d(F_1, F_2) = \top$ *when* $\mathbf{Force}_{\mathcal{R}}(F_1, F_2) = \varnothing$.

3.6 Theorem (Properties of the Lawvere Metric)
The function d satisfies:

1. Reflexivity. $d(F, F) = \perp$ *for all* F.
2. Triangle inequality. $d(F_1, F_3) \leq d(F_1, F_2) \otimes d(F_2, F_3)$ *for all* F_1, F_2, F_3.

The function d is asymmetric *in general:* $d(F_1, F_2) \neq d(F_2, F_1)$ *is the typical case. It is* non-separating *in general:* $d(F_1, F_2) = \perp$ *does not imply* $F_1 = F_2$.

Proof. Reflexivity follows from Proposition 3.4(1).

The triangle inequality is Proposition 3.4(2).

For asymmetry: consider the post-WWII demobilization example above. Let F_1 be the U.S. Army at 8.3 million and F_2 be the same army at 591,000. The cost of disbanding ($F_1 \to F_2$) is the cost of separation processing: low, accomplished in under two years by an administrative apparatus built precisely for that purpose. The cost of reconstitution ($F_2 \to F_1$) is the cost of recruitment, training, equipment procurement, and the reactivation of bases, commands, and the institutional memory that connects them: years and billions of dollars, with no shortcut available. In any reasonable cost functor, $d(F_1, F_2) \neq d(F_2, F_1)$. Federle et al.'s (2026) systematic evidence confirms the asymmetry at scale: in their 150-year, 60-country sample, the transformation from effective to war-devastated takes years of combat with a trough four years after onset, while the return transformation from devastated to effective takes roughly twelve years of recovery—a reconstruction cost that is structural, persistent, and consistent across every context in which it has been measured.

For non-separation: the Zulu *impondo zenkomo* (buffalo-horns formation) and the march column from which it deploys (Proposition 2.27) are distinct elements of $\mathbb{F}^\star(\mathbb{M}_L^\star)$—they differ in their command hierarchy and relational bonds—but the transformation between them costs nothing because it was installed through drill and fires automatically at command. Hence $d(F_{\text{column}}, F_{\text{horns}}) = \bot$ and $d(F_{\text{horns}}, F_{\text{column}}) = \bot$, yet $F_{\text{column}} \neq F_{\text{horns}}$. ■

The Lawvere metric is a *generalized* metric in the sense of Lawvere (1973): it satisfies reflexivity and the triangle inequality but does not impose symmetry or separation. Classical metric spaces are the special case where $Q = (\mathbb{R}_{\geqslant 0} \cup \{\infty\}, \leqslant, +)$, symmetry holds, and $d(x, y) = 0$ implies $x = y$. The generalization is not a weakening for its own sake; it faithfully represents two structural features of military transformation that a classical metric would suppress: the asymmetry between building and dismantling, and the nontrivial free neighborhood generated by the rearrangement groupoid $\mathbf{Force}^\times$ (Proposition 2.27).

The asymmetry of the Lawvere metric is the formal expression of a fact that military strategists have always known: it is cheaper to destroy than to build. This asymmetry has at least two strategic consequences that the metric makes precise. The first is the logic of deterrence. The threat of destruction has credibility because the attacker's cost $d(F_{\text{current}}, F_{\text{adversary-destroyed}})$ is low—destruction is cheap, especially when the adversary is a sitting target—while the defender knows that $d(F_{\text{destroyed}}, F_{\text{restored}})$ is very high: the cost of recovery, once imposed, runs to years and fractions of national income. Both sides can compute these asymmetric distances, and the gap between them is part of what gives the threat its force. The metric asymmetry does not guarantee that deterrence works—that depends on credibility, communication, and resolve—but it specifies the cost structure that deterrence exploits. The second consequence concerns the offense-defense balance. Biddle's (2004) analysis of combined-arms effectiveness implies a claim about metric asymmetry: the cost of *degrading* a combined-arms force is low (a successful attack on the coordinating headquarters suffices), while the cost of *restoring* one is very high. What is lost is not equipment—equipment can be replaced—but the trained relationships between arms that are installed by catalysis (Section 2.3), not by procurement. Experienced NCOs, combined-arms training cadres, and joint doctrine writers are the catalytic nodes that make the combined-arms formation run; destroying the nodes destroys the paths that lead to effectiveness, and reconstructing those nodes—which must happen before any path to combined-arms capability exists—costs exactly what Proposition 2.12 says: at minimum the reconstruction cost of the catalyst itself. Federle et al.'s finding that total factor productivity is the slowest-recovering war indicator confirms the pattern at the macroeconomic level: TFP encodes the organizational and institutional knowledge that the Lawvere metric identifies as the most expensive component of any restoration path, and it behaves accordingly.

A third consequence follows from the first two and extends deterrence theory to alliance commitments. Extended deterrence—the US commitment to defend allies under the nuclear umbrella or NATO Article 5—depends on the adversary's belief that the guarantee is credible. The metric perspective identifies where that credibility is grounded. The attacker's distance to the destruction of the protected ally is low: concentrated force applied to a smaller state produces rapid degradation at relatively low cost. The alliance's distance to restoring the ally's security after conquest, however, is very high: it requires not just military reconquest but the reconstruction of governance, administration, and the catalytic military infrastructure that was destroyed. This is the sense in which Schelling's (1966) distinction between deterrence by denial (preventing the attack) and deterrence by punishment (making the attack costly) maps onto the Lawvere metric. Deterrence by denial requires that the defender's distance to stopping the attack be smaller than the attacker's distance to completing it—a comparison of forward distances. Deterrence by punishment requires that the attacker's distance from the world-after-attack to a tolerable outcome be very high—a claim about the restoration cost the punishment imposes. In both cases, the credibility of the guarantee depends not on resolve or communication alone but on the underlying metric structure: whether the asymmetries in the Force-Makers' cost functions make the attacker's conquest cheap and the alliance's recovery expensive, or the reverse. Metric-aware deterrence theory is therefore not a refinement of existing accounts but a specification of the cost-structural conditions on which those accounts implicitly rely.

The Lawvere metric partitions the force space into three zones relative to any given force F.

The free neighborhood. The free neighborhood of F consists of all forces G with $d(F, G) = \bot$: the forces reachable from F by transformations that cost nothing. Formally, it is the orbit of F under $\mathbf{Force}^{\times}$, the rearrangement groupoid of Proposition 2.27. Non-separation of the Lawvere metric is precisely the statement that this orbit can be nontrivial: distinct forces can stand at distance zero from each other because the rules connecting them carry no cost to execute. The Zulu *impondo zenkomo* formation illustrates the point. The march column and the buffalo-horns attack formation are distinct structured forces—different relational bonds, different command hierarchies—but Shaka's decade of drilling the transformation into the impi's muscle memory placed both configurations inside a single free neighborhood. The cost of deploying from column to horns was $\bot$, and the cost of reforming from horns to column was also $\bot$; the free neighborhood of any Zulu field force therefore contained not just the march column but every tactical configuration in the drill library, all reachable without expenditure. This is the central point about the free neighborhood: it does not consist of forces reachable by "doing nothing." It consists of forces reachable by transformations whose cost was prepaid through prior investment in drill, catalytic infrastructure, and organizational design. The size of a force's free neighborhood is the return on that prior investment, made visible by the metric. Frederick the Great's oblique-order campaign at Leuthen in 1757 rested on the same logic. A decade of Prussian parade-ground drill had placed the oblique-order formation inside the Prussian free neighborhood, reachable at $\bot$ cost from any march formation at a moment's command. No Austrian force could enter that neighborhood at any finite cost, because the Austrian rule library contained no path to the oblique order's execution speed and geometrical precision. The Prussian free neighborhood was simultaneously Austrian unreachable space—and that disparity, not the numbers on the field, decided the battle. A generation later, the Imperial Japanese Navy engineered the same structure in a different domain. The carrier air attack on Pearl Harbor in December 1941 was not improvised; it was the culmination of a decade of rehearsal in which Japanese naval aviators drilled shallow-water torpedo runs, coordinated dive-bombing and torpedo sequences, and practiced navigating the specific approach routes over open ocean. By 7 December 1941, the attack formation was inside the Japanese free neighborhood: every element of the strike—the approach, the attack sequence, the withdrawal—cost $\bot$ to execute, because every step had been installed through catalytic training. The free neighborhood of the Kido Butai at the moment of attack was a precisely engineered structure, not an accident of organizational design.

The costly-reachable zone. The costly-reachable zone of F consists of all forces G with $\bot < d(F, G) < \top$: forces that F can become, but only by expending resources. The smaller $d(F, G)$, the cheaper the transformation; the larger, the more it costs. The bulk of strategic planning lives in this zone. The 1906–1914 dreadnought race is its clearest illustration. At every point in the race, both Britain and Germany stood in each other's costly-reachable zone for fleet superiority: no path was $\top$, because both sides possessed the industrial capacity and organizational rules to build dreadnought-class battleships. The question was never *whether* but *at what cost*, and Brewer's fiscal-military state (1990) gave Britain a compressed metric to every fleet-size target. Germany's higher Lawvere distance to the same targets was not a capability gap in the preorder sense—Germany was never *unable* to build dreadnoughts—but a cost gap that became a strategic liability as the race extended over years. The German navy could see the required force configurations in the preorder, could reach them, but faced systematically higher costs in every dimension: shipyard capacity, credit markets, naval administrative experience, revenue extraction efficiency. That is why the race was exhausting for Germany and sustainable for Britain even as both sides maintained continuous production. The costly-reachable zone is where arms races, military buildups, and cost-imposing strategies play out. Competitive strategy, precisely defined, is the art of forcing the adversary to spend more in her metric to reach a given target than you spend in yours to reach the counter-target: a structural exploitation of asymmetries in how each side's costly-reachable zone is priced.

The unreachable zone. The unreachable zone of F consists of all forces G with $d(F, G) = \top$: forces that F cannot become under the current rule library at any cost. They are structurally inaccessible, not merely expensive, and the distinction matters. Very expensive paths have finite cost and remain in the costly-reachable zone; $\top$ is reserved for the absence of any path, regardless of resources. China's position in 1895, after the First Sino-Japanese War, makes the distinction vivid. The Qing military's inability to build a modern battle fleet comparable to Japan's was not primarily a budgetary problem—China's economy was larger than Japan's, and the Qing state had, in the 1880s, actually purchased foreign warships. It was a rule-library problem. The Meiji state had spent three decades systematically installing the transformation rules for precision optics, breech-loading ordnance manufacture, modern naval gunnery, and Western staff doctrine; the Qing state had not, having pursued a strategy of acquiring Western *hardware* without acquiring the doctrinal and organizational rules on which that hardware's effectiveness depended. The result was that Chinese forces were metrically separated from Japanese forces not by a price gap but by a connected-component gap: no morphism in the Chinese rule library led to the configuration the Japanese navy occupied. Reformers like Zhang Zhidong understood this, even without the vocabulary of the present framework: what the First Sino-Japanese War demonstrated was that the Qing military did not occupy an expensive position within the Japanese force component but a different component entirely—one in which the modern fleet configuration did not appear as a node at any finite cost. The formal diagnosis is precise: the self-strengthening movement tried to add nodes (Western warships, rifles, artillery) without adding the bonds (trained gunnery officers, logistics chains, combined-arms doctrine) on which those nodes' effectiveness depended. In the DPO language of Chapter 2, they applied the node-addition operation without the bond-addition operations that would have made the added nodes load-bearing. The result was a force in which the hardware components existed as isolated nodes in the configuration graph, connected to the rest of the force by no trained relationship—precisely the condition that makes the hardware ineffective and the distance to effectiveness $\top$. The lesson is that unreachability is the condition the metric records as $\top$, not a large but finite number. Confusing the two—treating genuine rule-library absence as if it were merely a high price—produces the strategic error of believing that sufficient expenditure can substitute for the organizational and doctrinal transformations that the expenditure cannot purchase. The self-strengthening reformers were not wrong that China needed more resources; they were wrong that resources were the binding constraint.

The three-zone partition makes a claim that deserves formal statement. The boundary between the costly-reachable zone and the unreachable zone is precisely the convertibility preorder of Chapter 2: a force G is unreachable from F—that is, $d(F, G) = \top$—if and only if $F \not\leq_{\mathcal{R}} G$ (Proposition 3.7). The Lawvere metric does not discard any of the information that Chapter 2 assembled about the structure of valid transformations: it retains all of it and adds cost information on top. This is not obvious in principle. A cost function could in principle assign cost $\top$ to morphisms that do exist, treating achievable transformations as if they were impossible, and thereby introducing a gap between the preorder and the metric. The rule-library structure prevents this: the rule library $\mathcal{R}$ already determines which morphisms exist, and the cost functor assigns finite costs within that library—infinite cost arising only from the *absence* of a morphism, not from the presence of a prohibitive price. Under this structure, the following holds.

3.7 Proposition (Preorder-Metric Bridge)
Let c be a cost functor on $\mathbf{Force}_{\mathcal{R}}$ *with* $c(\phi) < \top$ *for all* $\phi \in \mathrm{Mor}(\mathbf{Force}_{\mathcal{R}})$*. Then*

$$F_1 \precsim_{\mathcal{R}} F_2 \iff d_{c,\mathcal{R}}(F_1, F_2) < \top.$$

Proof. If $F_1 \precsim_{\mathcal{R}} F_2$, there exists $\phi \in \mathbf{Force}_{\mathcal{R}}(F_1, F_2)$. Since $c(\phi) < \top$ by hypothesis, $d_{c,\mathcal{R}}(F_1, F_2) \leq c(\phi) < \top$. Conversely, if $d_{c,\mathcal{R}}(F_1, F_2) < \top$, the infimum over $\mathbf{Force}_{\mathcal{R}}(F_1, F_2)$ is strictly below $\top$; since the infimum over the empty set is $\top$ by convention, the hom-set must be nonempty, so $F_1 \precsim_{\mathcal{R}} F_2$. ■

The proposition closes a loop that has been open since Chapter 2. When $Q = (\{\bot, \top\}, \leq, \vee)$—the Boolean quantale of Section 3.1—the cost functor assigns $\bot$ to every morphism, and the Lawvere metric assigns $\bot$ to every pair for which a morphism exists and $\top$ to every pair for which none does. The metric contains exactly the information of the preorder, no more and no less. The Boolean quantale is the information-minimal cost structure: it is what you get when you price all executable transformation rules as free and let the only cost be impossibility itself.

The metric upgrade that richer quantales provide is therefore not a replacement of the preorder but a refinement of it. Planners who ask only "can we get there?"—the Boolean question—are working in the preorder. Planners who ask "what will it cost us, and what will it cost them?"—Marshall's question—are working in the Lawvere metric. The metric strictly contains the preorder as its reachability shadow.

The previous proposition answers a question about what the Lawvere metric preserves. A different question concerns what it *is*: among all the distance functions one might conceivably define on the force space, why is the Lawvere metric the right one? A Force-Maker who has assigned costs to individual morphisms might imagine that she retains some discretion in how to aggregate those assessments into a global distance—that different aggregation rules could yield different but equally defensible metrics. The following theorem establishes that she does not. Once she accepts three minimal standards of rational planning, her distance function is determined for her.

3.8 Proposition (The Metric Is Forced)
The Lawvere metric $d_{c,\mathcal{R}}$ *is the unique function* $d\colon \mathbb{F}^\star \times \mathbb{F}^\star \to Q$ *that satisfies all three of the following conditions and is maximal among all such functions in the pointwise order on Q-valued functions:*

1. Reflexivity. $d(F, F) = \bot$ *for all* F.

2. Triangle inequality. $d(F_1, F_3) \leq d(F_1, F_2) \otimes d(F_2, F_3)$ *for all* F_1, F_2, F_3.

3. Compatibility. $d(F_1, F_2) \leq c(\phi)$ *for every* $\phi \in \mathbf{Force}_{\mathcal{R}}(F_1, F_2)$.

Proof. Let d' be any function satisfying (1)–(3). By condition (3), $d'(F_1, F_2) \leq c(\phi)$ for every $\phi \in \mathbf{Force}_{\mathcal{R}}(F_1, F_2)$. Taking the meet over all such ϕ gives $d'(F_1, F_2) \leq \bigwedge_\phi c(\phi) = d_{c,\mathcal{R}}(F_1, F_2)$. Hence $d_{c,\mathcal{R}}$ dominates every competitor pointwise. Since $d_{c,\mathcal{R}}$ itself satisfies (1)–(3) by Proposition 3.4 and Definition 3.5, it is the maximum. ■

The proof is short; the content is not. Condition (1) says that doing nothing costs nothing—a minimal standard of consistency. Condition (2) is the triangle inequality: routing through an intermediate force F_2 cannot be cheaper than the direct distance from F_1 to F_3. If it could, the Force-Maker would be claiming that a detour saves money—which contradicts the infimum construction. Condition (3) is the substantive commitment: the Force-Maker's global distance assessment respects her individual morphism costs. She cannot believe that getting from F_1 to F_2 costs q if every morphism between them costs more than q.

Together, the three conditions leave no room for discretion. The Force-Maker who satisfies them is committed to the Lawvere metric and nothing else. Any other distance function that respected these three standards would assign *lower* values to some pair of forces—claiming that pair is closer than any actual morphism between them warrants. That is the planning error Proposition 3.8 rules out: wishful pricing, the belief that a transformation is cheaper than the cheapest path through the rule library permits, violates compatibility and is corrected by the infimum.

The metric is, in this sense, the rationality constraint on strategic cost assessment. A force planner who assigns independent costs to transformation rules and then derives a global distance by any rule other than the infimum is either being inconsistent or leaving information on the table. The Lawvere metric is what consistency requires.

The compatibility condition (condition 3 of Proposition 3.8) has a specific failure mode that Cold War intelligence practice made historically visible. During the 1960s and 1970s, the CIA estimated Soviet defense spending by converting Soviet military expenditures to dollars: pricing Soviet military transformation rules in the American cost functor rather than the Soviet one. This violated compatibility in the precise sense of the proposition—the CIA was assigning its own cost function c_{US} to Soviet morphisms, rather than the Soviet cost function c_{USSR}. The result was systematic distortion in both directions. Soviet manpower was cheap in rubles and expensive in dollars, because US labor costs far exceeded Soviet labor costs for identical conscript manpower; dollar-cost pricing therefore overestimated Soviet defense burdens as a share of national output. Soviet heavy industrial production, by contrast, had lower dollar-equivalent costs than US procurement for equivalent physical quantities of steel, aluminum, and propellants; pricing in dollars therefore underestimated the physical output of Soviet military industry for a given apparent expenditure. The Lawvere metric says these errors are structurally related: once you substitute c_{US} for c_{USSR}, the resulting distance function fails compatibility for both actors. The Team B exercise in 1976—which concluded that the standard CIA methodology simultaneously underestimated Soviet military capability and overestimated economic burden—was, in the language of the present framework, an argument that the CIA had been computing a function that satisfied compatibility for neither cost functor, and was therefore the Lawvere metric for no Force-Maker in the competition. The uniqueness theorem explains why this matters: compatibility is the condition that identifies which distance function the Force-Maker is committed to, and violating it removes all uniqueness guarantees while leaving the analyst with a distance measure that understates the adversary's capabilities in some dimensions and overstates them in others, without a principled way to tell which is which.

3.2.3 The Topology of Transformation Cost

The Lawvere metric generates a topology on $\mathbb{F}^{\star}(\mathbb{M}_L^{\star})$.

A topology formalizes the notion of "nearness" without specifying exact distances: it partitions neighborhoods from non-neighborhoods, continuous changes from discontinuous ones, limit points from isolated ones. For a military planner, the relevant notion of nearness is not Euclidean proximity of organizational charts but cost proximity in the metric: two forces are near if the transformation from one to the other is cheap. The Lawvere topology makes this precise. It gives the force space a shape—connected components, open neighborhoods, convergent sequences—derived from the cost structure rather than imposed on it. Why topology rather than just the metric? Because topology is the level of structure Debreu's representation theorem requires: not that distances be measured precisely, but that open sets and convergence be well-defined, so that the Force-Maker's ranking can be continuous relative to the natural notion of proximity in the space she inhabits. The metric generates the topology; the topology enables the representation; the representation gives the Force-Maker her force function. The chain is complete, and this subsection builds the first link.

3.9 Definition (Lawvere Topology)
For $F \in \mathbb{F}^\star(\mathbb{M}_L^\star)$ and $\epsilon \in Q$ with $\epsilon > \perp$, the forward open ball *of radius ϵ centered at F is:*

$$B_\epsilon^+(F) = \{G \in \mathbb{F}^\star(\mathbb{M}_L^\star) : d(F,G) < \epsilon\}.$$

The Lawvere topology τ_d *on $\mathbb{F}^\star(\mathbb{M}_L^\star)$ is the topology generated by all forward open balls.*[3]

Two forces are topologically close in τ_d if and only if the cheapest transformation from one to the other is cheap. Forces that differ by a single low-cost DPO rewrite step are near neighbors; forces connected only through costly restructuring are distant; forces in different unreachable zones lie in different connected components. The topology makes precise the informal notion that "similar force" means "cheaply reachable force." But "cheaply reachable" is not the same as "structurally similar," and the distinction is strategically consequential.

Two forces can share the same organizational chart, the same equipment inventory, and the same doctrine on paper while being metrically far—and therefore topologically distant—because the rules that actually connect them carry high cost. Structural similarity is a property of objects; metric distance is a property of paths between them. The US Army and the Wehrmacht in 1940 illustrate the gap. Both were organized around divisions with combined arms; both operated tanks, artillery, and motorized infantry in roughly similar proportions. On any structural measure of organizational form—order of battle, equipment densities, divisional structure—they were near each other. But the Wehrmacht's operational effectiveness reflected six years of deliberate doctrinal development: the institutionalization of *Auftragstaktik* (mission-type orders

[3] Because d is asymmetric, forward balls $B_\epsilon^+(F) = \{G : d(F,G) < \epsilon\}$ and backward balls $B_\epsilon^-(F) = \{G : d(G,F) < \epsilon\}$ generate different topologies. The forward topology captures the Force-Maker's planning perspective: which forces can I cheaply become? The backward topology captures the intelligence perspective: which forces could have cheaply become mine? We work throughout with the forward topology.

that allowed junior commanders to act without explicit authorization), the incorporation of lessons from the Spanish Civil War into training programs, and the rehearsal of combined arms at the corps level through the 1939 Poland campaign. The US Army in 1940 had none of this. The path from the 1940 US Army to Wehrmacht operational effectiveness would require years of combat learning, officer rotation, doctrine revision, and the painful experience that the North Africa campaign provided. The two forces were structurally near but metrically far: topologically distant in the Lawvere topology despite their surface resemblance. A force analysis that monitors only structural similarity—who has tanks, who has divisions—reads the topology of the objects and misses the topology of the metric, which is where the operationally decisive information resides.

The asymmetry of the metric generates not one topology on the force space but two, and they answer different strategic questions. The forward topology, generated by the forward balls $B^+_\epsilon(F) = \{G : d(F, G) < \epsilon\}$, is the planner's topology: which forces can I cheaply become? The backward topology, generated by the backward balls $B^-_\epsilon(F) = \{G : d(G, F) < \epsilon\}$, is the intelligence analyst's topology: which forces could have cheaply become mine, or could cheaply become a threat? These generate genuinely different neighborhoods, and conflating them produces genuine analytical errors. After the 1973 Yom Kippur War, American and Israeli intelligence analysts faced both questions simultaneously but found it useful to separate them. The forward question—what can each belligerent cheaply become next?—concerned the military balance and directed attention to each side's rule library and cost functor. The backward question—which regional forces could cheaply adopt Egyptian-style combined arms and anti-tank doctrine?—was an intelligence question about threat emergence: identify all F such that $d(F, G^\star_{\text{Egyptian-doctrine}})$ is small. The two questions have different answers because the forward and backward topologies have different neighborhoods. A state that maintains situational awareness only in the forward topology—tracking what adversaries are building—is blind to forces that are currently dissimilar but metrically near in the backward topology, the forces that could cheaply become a threat under the right catalytic conditions. The Syrian military in 1973 illustrates the point. Its order of battle—Soviet-equipped, divisionally organized, with surface-to-air missile coverage that constrained Israeli air operations—was dissimilar to the Egyptian combined-arms formation in organizational culture and doctrine. But in the backward topology, Syria was close to the Egyptian configuration: it possessed the hardware (T-62 tanks, SAM batteries), the organizational skeleton (combined arms division structure), and, crucially, Soviet advisors who served as the catalytic nodes for integrating the elements. The metric distance $d(F_{\text{Syria}}, G^{\star}_{\text{Egyptian-doctrine}})$ was relatively small because the rule-library additions required were modest—primarily doctrinal and coordinative rather than industrial. Israel had monitored the forward topology carefully (what could Syria and Egypt become from their current positions?) but the backward topology revealed a different threat: that a pre-positioned combined-arms assault coordinated between both fronts could cheaply reach a configuration that neither front alone could approach. The October 1973 attack exploited precisely this metric structure: two forces individually in Israel's costly-reachable zone combined into a configuration that briefly placed Israel itself in the costly-reachable zone of its adversaries.

The topology's connected components are the unreachable zones: forces in different components cannot reach each other at any finite cost under the current rule library. This is the topology's coarsest information, corresponding to the partition of the force space by the convertibility preorder. But within a connected component, the metric gives a gradient: some directions of movement are cheap and some are expensive, and the Lawvere topology makes that gradient visible in its open sets. Strategic planning, in the topological language, is the problem of navigating toward a target $G^{\star}$ by moving through cheap neighborhoods, each step a low-cost transformation that places the force in a new open ball closer to the objective. The force planner who proceeds by identifying the next cheapest transformation—chaining together DPO rewrites, organizational adjustments, and catalytic installations in order of cost—is doing what the topology formalizes as moving along a path toward a limit point. This is not a metaphor: the Lawvere topology's open sets are precisely the sets of forces reachable below a given cost threshold, and a planning sequence that moves from open ball to open ball—from "reachable below ϵ_1" to "reachable below $\epsilon_2 < \epsilon_1$"—is a sequence that converges to the target in the topology. The intelligence cycle plays a specific role in this framework: it is the process by which the planner identifies which direction within the current open ball is cheapest—which next transformation places her in the largest or most useful subsequent neighborhood. The key structural fact that guarantees this process is well-behaved is second-countability. A topology without a countable base can have capability targets that are not approachable by any *sequence* of intermediate forces—targets one can reach only by a single discontinuous jump or not at all, because every neighborhood of the target fails to contain any force reachable from the current position. This pathology—a strategic cliff with no gradient—would mean that some capability improvements are impossible to approach incrementally; they must be achieved all at once or not at all. Second-countability eliminates this pathology. Every capability target in the Lawvere topology, under the countability conditions of Proposition 3.10, can be approached by a sequence of intermediate forces of increasing capability and increasing cost. Planning can proceed incrementally: no capability target requires an impossible leap, and every capability gap has a sequence of partial improvements that approach it from below.

3.10 Proposition (Second Countability)

If Q contains a countable order-dense subset $D \subseteq Q$ (meaning: for all $a < b$ in Q, there exists $\delta \in D$ with $a \leqslant \delta < b$), then the Lawvere topology τ_d on $\mathbb{F}^\star(\mathbb{M}_L^\star)$ is second-countable.

Proof. The force space $\mathbb{F}^\star(\mathbb{M}_L^\star)$ is countable (Proposition 1.17). The collection

$$\mathcal{B} = \left\{B_\delta^+(F) : F \in \mathbb{F}^\star(\mathbb{M}_L^\star), \delta \in D, \delta > \bot\right\}$$

is a countable collection of open sets (countable centers times countable radii). We claim $\mathcal{B}$ is a base for τ_d. Let $U \in \tau_d$ and $G \in U$. By definition of the generated topology, there exist finitely many centers $F_1, \ldots, F_k$ and radii $\epsilon_1, \ldots, \epsilon_k$ with $G \in B_{\epsilon_1}^+(F_1) \cap \cdots \cap B_{\epsilon_k}^+(F_k) \subseteq U$. For each i, $d(F_i, G) < \epsilon_i$; by order-density of D, there exists $\delta_i \in D$ with $d(F_i, G) \leqslant \delta_i < \epsilon_i$. Taking $F = G$ as center and using the triangle inequality, we find a ball in $\mathcal{B}$ centered at G contained in each $B_{\epsilon_i}^+(F_i)$, hence contained in U. Thus $\mathcal{B}$ is a countable base. ∎

The density condition is satisfied by all four quantale examples: the textbook quantale has $\mathbb{Q}_{\geqslant 0}$ as a dense subset; the time-money quantale has $\mathbb{Q}_{\geqslant 0}^2$; the probability quantale has $\mathbb{Q} \cap [0, 1]$; the resource-set quantale, being finite, is trivially countable. Second countability is not assumed here; it is earned, from the countability of the force space (Proposition 1.17) and the structure of the quantale. It is also the technical prerequisite for Debreu's representation theorem: it is the topological condition that guarantees a continuous real-valued function representing the Force-Maker's ranking of forces can exist (Chapter 4). The result is a consequence of three chapters of prior construction: the force space is countable because of the recursive construction of Chapter 1; the morphisms compose in a category because of the DPO and organizational machinery of Chapter 2; the costs define a quantale with a countable dense subset by assumption on Q. Second-countability follows from all three, not from any one alone, and it is not assumed as a regularity condition but derived as a theorem. This matters for the status of Debreu's representation result in Chapter 4: the topological hypothesis of that theorem is not a background assumption about the Force-Maker's strategic landscape—it is a proved consequence of the mathematical structure she already operates within.

3.3 Doctrine and the Cost of Restriction

The Lawvere metric is rule-library-relative: different institutions, inhabiting different doctrinal traditions, face different metrics on the same space. This is not merely a theoretical observation. It is the formal content of a large empirical literature on how military doctrine shapes organizational behavior, and the present framework makes that content precise in a way the narrative literature does not.

"Doctrine," in the sense Posen (1984) uses the term, is the set of preferred concepts for the use of force: which formations to adopt, which operations to privilege, which force compositions to optimize. In the present framework, doctrine maps directly onto the rule library $\mathcal{R}$: doctrine specifies which transformation rules the organization endorses, which it discourages, and which it forbids. The rule library is not a complete enumeration of all physically possible transformations; it is the organization's *selection* from that larger space—the rules it has codified, trained, and equipped for. Different doctrines select different rule libraries and thereby induce different metrics on the same underlying force space. An organization with a maneuver-warfare doctrine selects rules that compress the metric to rapid positional changes and deep penetrations; an organization with an attrition doctrine selects rules that compress the metric to sustained firepower generation and logistics endurance. The same set of forces, in the same strategic environment, appears geometrically very different from inside each doctrine—not because the forces themselves differ but because the rule selection determines the landscape of cheap and expensive paths. This is the formalization of Posen's core empirical claim: organizational doctrine shapes not just what a military prefers to do but what it can cheaply do, and these two things are not always the same.

3.11 Proposition (Doctrinal Restriction Raises Cost)
If $\mathcal{R}' \subseteq \mathcal{R}$*, then* $d_{\mathcal{R}'}(F_1, F_2) \geq d_{\mathcal{R}}(F_1, F_2)$ *for all* $F_1, F_2 \in \mathbb{F}^\star(\mathbb{M}_L^\star)$.

Proof. $\mathbf{Force}_{\mathcal{R}'}(F_1, F_2) \subseteq \mathbf{Force}_{\mathcal{R}}(F_1, F_2)$, so the infimum over the smaller set is weakly larger. ■

Proposition 3.11 is the quantitative companion to Proposition 2.18 from Chapter 2, which established that doctrinal restriction shrinks the convertibility preorder. The present result goes further: restriction does not merely eliminate paths, it raises costs. A doctrine that forbids certain transformation rules forces the Force-Maker onto longer, more expensive routes to the same destinations.

Posen's (1984) argument that military organizations develop doctrine that reflects their strategic environment and organizational interests has a natural reading in this framework: doctrine determines $\mathcal{R}$, and the choice of $\mathcal{R}$ determines the Lawvere metric. An organization whose doctrine is well-matched to its strategic environment faces low Lawvere distances to the forces it needs; an organization whose doctrine is poorly matched faces high distances, because the rules it permits are not the ones that lead cheaply to the required capability. Doctrinal rigidity is not merely a cultural observation; it is an enforced cost premium.

Proposition 3.11 has a positive direction as well: an *expansion* of the rule library—$\mathcal{R} \subseteq \mathcal{R}'$—compresses the Lawvere metric to every reachable target, because additional rules create additional paths and the infimum over a larger set is weakly smaller. The military revolution literature is, in this reading, a sequence of metric compressions: each revolution installs rules that dramatically reduce the cost of reaching capabilities that previously required far more expensive paths. The seventeenth-century drill revolution—Gustavus Adolphus refining Swedish massed-volley tactics, Louvois institutionalizing drill in the French army, Frederick systematizing it into the Prussian machine—expanded rule libraries precisely in the direction that compressed the metric to coordinated combined-arms fire. Before systematic drill, the transformation from "assembled troops" to "sustained disciplined musket volleys" required experienced officers who physically positioned and cued each rank; after drill, that transformation was installed catalytically through training and executed at command, dropping its cost to near $\bot$. The metric from any trained unit to any other trained formation in the library compressed simultaneously, because the same catalytic investment that installed one formation installed the grammar by which formations could be chained. This is the formal content of what military historians mean when they say that a revolution changes the "terms of competition." The revolution does not merely give adopters new capabilities in the preorder sense—it restructures the Lawvere metric so that adopters face much lower costs to decisive capabilities than laggards do. Laggards face high metric distances to the same capabilities not because they are too poor to pay the market price but because their rule libraries lack the relevant transformations, and no amount of spending can substitute for the organizational and doctrinal investment that installs them. The competitive pressure to adopt is therefore structural: the Lawvere metric penalizes doctrinal lag with compounding cost disadvantages, in exactly the same way that a restricted rule library raises every distance and eliminates every shortcut.

The coup-proofing literature offers a particularly precise illustration. Talmadge's (2015) analysis of authoritarian military organizations shows that dictators who fear military coups systematically purge or rotate the officers—drill instructors, staff planners, experienced commanders—who serve as catalytic nodes in the force transformation network. In the language of Section 2.3, these purges destroy the catalytic subgraph C on which the training and combined-arms rules depend (Proposition 2.12). Without the catalyst present in the host graph, every path to a trained, effective force must first recreate the catalyst from the void, at a cost of at least $d(\nvdash, C)$ where $\nvdash$ is the void force. The purge itself is cheap; the dictator removes officers at low cost. The reconstruction is expensive: it takes years, requires experienced mentors who may not exist, and involves building back precisely the institutional structure the purge was designed to destroy. The asymmetry of the Lawvere metric captures this with precision: $d(\text{effective}, \text{coup-proofed})$ is low—the dictator moves easily in the direction of suppressing military effectiveness—while $d(\text{coup-proofed}, \text{effective})$ is very high, because every path through the metric to military effectiveness runs through catalyst reconstruction.

This is, in the vocabulary of Section 2.3, a systematic study of the cost consequences of catalytic destruction. Talmadge's coup-proofed armies do not merely face a smaller rule library (the qualitative claim from Chapter 2); they face a higher Lawvere distance to any force requiring the destroyed catalysts, by at least the cost of catalyst reconstruction. The Chapter 2 analysis said what doctrinal restriction removes; the present chapter says what that removal costs.

3.4 Net Assessment as Metric Comparison

The Lawvere metric is a planning tool for a single Force-Maker. The most consequential applications of cost reasoning in military affairs have been comparative: what does it cost us versus what does it cost them, and where do the differences lie? Andrew Marshall directed the Pentagon's Office of Net Assessment for four decades with exactly this question as his organizing commitment.[4] Marshall's central insight was not "who has more force" but "who is overpaying for the force they have, and who can exploit that overpayment." In the language of the present framework, Marshall was comparing Lawvere metrics.

Formal setup. Consider two Force-Makers FM_1 and FM_2 with rule libraries $\mathcal{R}_1$ and $\mathcal{R}_2$ and cost functors c_1 and c_2 valued (after mapping to a common quantale) in the same Q. Each sees a different Lawvere metric on the same force space:

$$d_1(F,G) = \bigwedge_{\phi \in \mathbf{Force}_{\mathcal{R}_1}(F,G)} c_1(\phi), \qquad d_2(F,G) = \bigwedge_{\phi \in \mathbf{Force}_{\mathcal{R}_2}(F,G)} c_2(\phi).$$

Net assessment is the systematic comparison of d_1 and d_2. The differences between them arise from structural asymmetries in rule libraries (what each side knows how to do), industrial bases (what each side can produce), and institutional costs (what it costs each side to execute a given transformation)—not from error or irrationality on either side. The assumption that both force-makers value costs in a common quantale Q is the working assumption of comparative net assessment: that there exists a shared scale on which to compare costs, even if neither actor explicitly computes in that scale. In practice, this means projecting each actor's cost assessments onto a common measurement framework—exactly what Marshall's office attempted to do by developing commensurable metrics for US and Soviet military capability. When no common quantale exists—when the two actors' cost dimensions are genuinely incommensurable—net assessment can still proceed by comparing the metrics pointwise on specific target pairs rather than globally, asking not "who has the better metric overall" but "who is cheaper for this specific target?" The pointwise comparison is sufficient for competitive strategy purposes: a cost-imposing strategy requires only that there exist a target $G^\star$ such that $d_2(F_2, G^\star) \gg d_1(F_1, G^\star_{\text{counter}})$, not that the full metrics be comparable everywhere.

[4] Marshall directed the Office of Net Assessment from 1973 until his retirement in 2015. His approach—comparative analysis of the military balance, with sustained emphasis on competitive strategies and cost-imposing approaches—is most accessible through Krepinevich and Watts (2015); the connection to the transformation debate of the 1990s is developed in Krepinevich (1997).

Competitive advantage as metric asymmetry. Fix a target capability $G^\star$—a specific force configuration representing, say, the ability to project power across a given strait. The cost to FM_1 of reaching $G^\star$ is $d_1(F_1, G^\star)$; the cost to FM_2 is $d_2(F_2, G^\star)$. Three cases arise:

- $d_1(F_1, G^\star) < d_2(F_2, G^\star)$: FM_1 holds a cost advantage at this target.
- $d_1(F_1, G^\star) > d_2(F_2, G^\star)$: FM_2 holds the advantage.
- The distances are incomparable in Q: neither side has a clear advantage, and the competition turns on which dimension of cost becomes binding first.

The power of this framework is that different targets yield different comparisons. FM_1 may be cheaper for naval supremacy ($d_1(F_1, G^\star_{\text{naval}}) < d_2(F_2, G^\star_{\text{naval}})$) while FM_2 is cheaper for ground dominance ($d_2(F_2, G^\star_{\text{ground}}) < d_1(F_1, G^\star_{\text{ground}})$). This is the quantification of what Proposition 2.25 established qualitatively: that monotones separate incomparable forces. Different metric projections rank the same pair of forces differently, and the question of which ranking matters depends on which target capability is strategically salient.

Burden sharing as metric divergence. The two-metric framework is also the formal apparatus that the alliance burden-sharing literature has long needed. Olson and Zeckhauser's (1966) foundational analysis of collective defense showed that NATO allies systematically underprovide defense relative to the socially optimal level: larger, wealthier members bear a disproportionate share of the alliance burden, while smaller members free-ride. Their explanation invoked differential marginal costs of defense provision without a formal language for representing those differences. The Lawvere metric supplies that language. Each ally faces its own Lawvere metric on the shared force space—its own $d_i(F_i, G^\star_{\text{deterrence}})$—reflecting its rule library, industrial base, and institutional cost structure. The larger ally's metric to the collective deterrence target is lower not merely because it is wealthier in the abstract but because its cost functor is compressed across the relevant dimensions: deeper military-industrial infrastructure, a more experienced professional officer corps, established logistics networks that make each individual transformation cheaper to execute. The smaller ally's higher metric means that reaching the same deterrence target costs it more, in the same sense that the Ardennes route cost Germany less than the Belgian frontal assault. Free-riding is therefore not a failure of alliance solidarity but a predictable consequence of metric asymmetry: rational allies contribute up to the point where their marginal cost of additional provision equals their marginal benefit, and allies with higher metrics reach that point sooner. The burden-sharing problem, on this reading, is the problem of designing alliance structures—commitment mechanisms, specialization agreements, capability requirements—that account for the structural metric differences between members rather than pretending all allies face the same cost geometry.

The Soviet anti-ship missile. Marshall's most celebrated application of metric asymmetry concerned the contest between American carrier strike groups and Soviet naval strategy. Carrier groups cost billions of dollars, required years to construct and assemble, and demanded thousands of trained personnel: $d_1(F_1, G^\star_{\text{carrier-defense}})$ was very high. The Soviet Union, confronting the question of how to neutralize this capability without building competing carrier groups of its own (for which $d_2(F_2, G^\star_{\text{carrier-match}}) \approx \top$), discovered a structural shortcut. Anti-ship cruise missiles—a known class of transformation in the Soviet rule library, applied to a new strategic context—could deny American carrier supremacy at a small fraction of the cost:

$$d_2\left(F_2, G^\star_{\text{carrier-denial}}\right) \ll d_1\left(F_1, G^\star_{\text{carrier-defense}}\right).$$

The cost asymmetry was structural, not accidental: the Soviet rule library contained rules for precision missiles that the American library did not contain for an equivalent counter-response, and the industrial and institutional costs of missile production were a fraction of carrier group maintenance. This asymmetry was visible only when transformations were priced—when someone asked not "who has more platforms" but "what does it cost each side to reach strategic parity at sea"—and its discovery was foundational to the competitive-strategies approach to defense planning (Krepinevich and Watts, 2015).

A cost-imposing strategy, defined precisely, is a strategy that exploits asymmetries in the Lawvere metric: forcing the adversary to spend more in her metric to achieve or maintain a given capability than you spend in yours to deny it. The asymmetry lies in structure—different rule libraries, different industrial bases, different institutional costs—not in error, and it is therefore durable.

Cost genius. Chapter 2 identified two kinds of military genius: match-finding (recognizing that a known transformation applies to the current situation) and scheme-invention (designing a transformation that does not yet exist in any rule library). The Lawvere metric reveals a third kind: *cost genius*—the ability to find a path through the metric that achieves the strategic objective at a fraction of the expected cost. The Soviet missile designers were cost geniuses in this sense: they did not invent a new class of weapon (anti-ship missiles had conceptual precedents in torpedo doctrine) nor did they find an unexploited application of a familiar rule (the strategic novelty was real). What they found was a shortcut in the Lawvere metric, a known class of transformations that achieved a strategic objective at a structural cost advantage that the convertibility preorder alone could not have revealed. The preorder is a map; the metric is a map with distances; and the distances are where strategy lives.

Arms races, overstretch, and hegemonic transition. Kennedy's (1989) argument that great powers decline when the costs of military commitments outstrip the economic base that supports them is a claim about Lawvere metric dynamics: as the economic base shrinks, the cost functor changes, and distances to maintenance targets grow. Forces that were once cheaply held become expensive to maintain; positions once occupied at low cost become costly to defend; and the Force-Maker drifts away from the capabilities that once defined her power. Gilpin's (1981) account of hegemonic transition gives this dynamic a specific comparative structure. A rising challenger does not simply accumulate resources; its cost functor changes composition as industrialization, military learning, and institutional development compress the metric to precisely the capability targets that define great-power competition. The transition point, in the language of the present framework, is when $d_{\text{challenger}}(F, G^{\star})$ falls below $d_{\text{hegemon}}(F, G^{\star})$ for a sufficient set of strategically decisive targets: when the challenger can reach the capabilities that define hegemonic status more cheaply than the incumbent can maintain them. At that point the existing distribution of benefits and costs no longer reflects the underlying distribution of Lawvere distances, and the pressure toward political renegotiation—peaceful or otherwise—becomes structural. Brewer's (1990) fiscal-military state is, on this reading, precisely what delayed that transition for Britain: the institutional infrastructure that compressed Britain's cost functor—functioning credit markets, efficient revenue extraction, professional naval administration—gave it a durable metric advantage even as Germany's industrial capacity caught and eventually surpassed it in the aggregate. The dreadnought race compressed Gilpin's transition dynamic into eight years: the convertibility preorder said both Britain and Germany could build dreadnought fleets, but the Lawvere metric said Britain could build them cheaply and Germany could not, and that asymmetry compounded with every year of the race until the fiscal strain on Germany became a strategic liability of its own.

3.5 Looking Ahead

Three layers of structure are now complete. Chapter 1 built the objects: the set $\mathbb{F}^\star(\mathbb{M}_L^\star)$ of all structured forces, constructed recursively from elements to molecules to configurations to organizations, each force a structured graph carrying command hierarchy and tactical bonds. Chapter 2 built the morphisms: the category **Force**, equipped with DPO rewriting, catalytic transformation, and organizational restructuring, encoding which transformations are valid and how they compose—and demonstrating that the convertibility preorder, the rearrangement groupoid, and the rule-monotonicity theorem all follow from this categorical structure. The present chapter built the geometry: a cost functor valued in a quantale, inducing a Lawvere metric on the force space and, from that metric, a topology in which proximity means cheapness, connected components mean reachability, and second-countability means that every capability target can be approached by a sequence of cheaper intermediate forces.

What is still missing is the Force-Maker's judgment. She does not merely navigate the force space; she ranks it. She has a doctrine—a preorder on structured forces, a judgment that this posture is at least as capable as that one—and she acts on that ranking. The contest model's force function m_i, Dupuy's Combat Power Potential, the intelligence analyst's net assessment index: all are attempts to compress the structured force into a scalar that carries the Force-Maker's ranking. Chapter 0 asked whether this compression is faithful, and the question could not then be answered because the topology had not been built. It can be answered now.

The answer requires a classical result from mathematical economics. Debreu (1954) showed that a complete preference ordering on a second-countable topological space can always be represented by a continuous real-valued function. Proposition 3.10 delivers the topological half of that condition: the Lawvere topology is second-countable, not by assumption but as a theorem. The remaining half—whether the Force-Maker's doctrine constitutes a complete preorder, ranking any two forces against each other—is a claim about the doctrine itself, not about the geometry. That is the question Chapter 4 takes up.

Return, finally, to the epigraph. Weil wrote that the retribution visited upon abusers of force has a geometrical rigor that operates automatically. The Lawvere metric is the formalization of that rigor: it does not punish ambition, it prices it, and the price is computable. The Force-Maker who understands the metric sees not merely what she can reach but what it costs her adversary to respond—and that asymmetry, visible only when transformations are priced, is the decisive advantage that Marshall spent four decades seeking and the fiscal-military state gave Britain a century before him.

Chapter 4

Force is Subjective

Only he who has measured the dominion of force, and knows how not to respect it, is capable of love and justice.

Simone Weil, "The Iliad, or the Poem of Force"

Alfred Thayer Mahan's *The Influence of Sea Power upon History*, published in 1890, argued that national power rested on naval supremacy—on the capital ship, the fleet-in-being, and control of the sea lanes through which commerce flowed and armies were supplied. A doctrine that ranked force configurations by their contribution to sea control placed battleship fleets at the top, coastal defense forces in a supporting role, and ground armies as instruments of secondary importance, useful for holding what the fleet had won. Giulio Douhet's *The Command of the Air*, published in 1921, offered a different ranking from different premises. The decisive instrument was the strategic bomber; armies and navies were subordinate forces, capable of occupying territory but unable to compel capitulation; and any force configuration whose primary contribution was ground-holding occupied a lower position in the ranking than a configuration built around long-range airpower. Soviet military doctrine, developed through the interwar period and formalized in the concept of the *sootnoshenie sil*—the correlation of forces—was more ambitious still: a comprehensive ranking of military, economic, and political capability in a single index, aggregating across force types, alliance structures, and what Soviet planners called the ideological balance of forces. Each of these three doctrines was a ranking of the same structured force space from a different vantage point, and each implied different production priorities, different alliance commitments, and different employment concepts.

The interwar debates between battleship advocates and carrier advocates, between strategic bombing theorists and ground-force commanders, between capital-intensive and labor-intensive force structures, were debates about whose ranking was right. The British admirals who defended the Washington Naval Treaty ratios were protecting a Mahanian preorder in which capital ship tonnage was the relevant measure; the airpower theorists who argued that capital ships were obsolete were proposing a Douhetian preorder in which those same ships occupied a lower position. The disagreement was not primarily empirical—neither side lacked information about ships or aircraft. It was doctrinal: a disagreement about which ranking was the correct representation of what force could accomplish.

The preceding chapters have built the mathematical structure needed to make this observation precise. Chapter 1 built the objects: the set $\mathbb{F}^\star(\mathbb{M}_L^\star)$ of all structured forces, constructed recursively from elements to molecules to configurations to organizations. Chapter 2 built the morphisms: the category **Force**, equipped with DPO rewriting, catalytic transformation, and organizational restructuring. Chapter 3 built the geometry: a cost functor valued in a quantale, inducing a Lawvere metric on the force space and, from that metric, a second-countable topology in which proximity means cheapness and connected components mean reachability. What those chapters did not provide is the Force-Maker's judgment—her ranking of the positions she can occupy in that geometry. That ranking is the subject of the present chapter. The central result is that the ranking can be faithfully compressed into a continuous real-valued function when—and only when—the doctrine satisfies conditions that are substantive claims about strategy, not consequences of any mathematical construction. The force function is not a fact of the world; it is a representation of a doctrine. Force is subjective in this precise sense.

4.1 Doctrine as Preorder

The doctrines described above—Mahanian, Douhetian, Soviet—share a common mathematical structure. Each is a binary relation on the structured force space $\mathbb{F}^\star(\mathbb{M}_L^\star)$ satisfying two conditions: reflexivity, since every force is at least as capable as itself, and transitivity, since if F outranks G and G outranks H then F outranks H. A relation with these two properties is a preorder, and the present chapter calls a preorder on the force space a doctrine.

4.1 Definition (Doctrine)
A doctrine *is a binary relation* $\succsim$ *on* $\mathbb{F}^\star(\mathbb{M}_L^\star)$ *that is:*

1. Reflexive*:* $F \succsim F$ *for all* $F \in \mathbb{F}^\star(\mathbb{M}_L^\star)$*; and*
2. Transitive*: if* $F \succsim G$ *and* $G \succsim H$*, then* $F \succsim H$*.*

Write $F \succ G$ *(*strictly preferred*) when* $F \succsim G$ *but not* $G \succsim F$*; write* $F \sim G$ *(*doctrinally equivalent*) when* $F \succsim G$ *and* $G \succsim F$*.*

The definition is deliberately spare. A preorder does not require that every pair of forces be comparable: it is entirely possible that $F \not\succsim G$ and $G \not\succsim F$ simultaneously, meaning the doctrine has no verdict about the comparison between F and G. Whether a doctrine is *complete*—whether it ranks every pair—is a separate assumption, and a substantive one.

Completeness. A doctrine $\succsim$ is *complete* if for all $F_1, F_2 \in \mathbb{F}^\star(\mathbb{M}_L^\star)$, either $F_1 \succsim F_2$ or $F_2 \succsim F_1$ (or both). Completeness is the claim that the Force-Maker's doctrine covers every comparison in the force space—that no two forces fall outside the scope of her ranking.

Whether completeness holds is a doctrinal question, not a mathematical one, and the historical record exhibits both cases with clarity. The Wehrmacht's operational assessment of Polish cavalry against German panzer divisions in 1939 is a clear case for completeness within a relevant subdomain: firepower, mobility, and protection all pointed in the same direction, and the doctrine produced an unambiguous ranking without effort. The United States Army's doctrine in 2001 presents the opposite case. That doctrine was fully articulated and internally consistent over the full range of conventional combined-arms configurations—it could rank any two mechanized formations against each other, compute required force ratios for offensive operations, and specify logistics requirements for sustained maneuver. But it had no developed ranking over counterinsurgency-configured forces, which occupied an essentially empty region of the force space from the standpoint of the extant preorder. The doctrine was incomplete not because of a logical error but because it had simply never been required to cover configurations the Army did not plan to field. Completeness is an achievement, not a default.

The debate between Mearsheimer (2001) and Biddle (2004) is, at its core, a debate about whether the Force-Maker's doctrine is complete. Mearsheimer's theory effectively assumes completeness: every force can be ranked against every other by a common index of lethality-weighted military assets, and that ranking determines which states can threaten which. Biddle's critique is that this assumption fails because force comparisons are mission-dependent—a combined-arms formation configured for maneuver warfare cannot be coherently ranked against a counterinsurgency-configured light infantry unit without first specifying the operational environment that determines which ranking is relevant. To force the comparison without that specification is not to achieve completeness but to suppress the incompleteness beneath a homogenizing formula. The formula produces numbers; but numbers are not doctrine, and the preorder those numbers purport to represent does not exist in the form the formula implies.

Continuity. Completeness is a condition on the coverage of the doctrine. Continuity is a condition on its smoothness: the doctrine should not assign dramatically different capability assessments to forces that are close in the Lawvere metric—forces cheaply reachable from each other.

4.2 Proposition (Continuity of Doctrine)

A doctrine $\succsim$ is continuous with respect to the Lawvere topology τ_d if and only if for all $F \in \mathbb{F}^\star(\mathbb{M}_L^\star)$:

1. *the upper contour set $\{G : G \succsim F\}$ is closed in τ_d; and*
2. *the lower contour set $\{G : F \succsim G\}$ is closed in τ_d.*

Proof. By definition of the Lawvere topology τ_d (Proposition 3.10), a set is open if it is a union of forward open balls $B_\epsilon^+(F) = \{G : d(F, G) < \epsilon\}$. A preorder $\succsim$ is continuous with respect to a topology if and only if its upper and lower contour sets are closed in that topology—equivalently, that their complements are open. This is precisely the standard topological condition for continuity of a preorder; see Debreu (1954) for the general statement applied to a second-countable space. ∎

The substantive content of continuity is a prohibition on doctrinal cliffs. A doctrine that is discontinuous in the Lawvere topology assigns dramatically different capability assessments to forces connected by a single low-cost transformation—calling one formation decisive and the next, differing only by the addition of an artillery brigade reachable at modest Lawvere distance, strategically marginal. Such a cliff cannot be grounded in the force space's geometry, because the geometry sees those two formations as close; the doctrine would have to be attributing the cliff to some feature the metric does not capture. Continuity rules this out not because doctrinal cliffs are logically impossible but because they signal an inconsistency between the doctrine and the geometry it purports to evaluate. A coherent doctrine respects the metric.

4.2 The Representation Theorem

The preceding section identified completeness and continuity as the conditions under which a doctrine is well-behaved. The present section states the central result of the chapter: when both conditions hold on the force space built in Chapter 3, the doctrine can be faithfully compressed into a single continuous real-valued function—a force function. This is Debreu's 1954 representation theorem, and the second-countable topology established in Proposition 3.10 is precisely the topological condition it requires.

4.3 Theorem (Debreu Representation)

Let $\succsim$ be a complete, continuous doctrine on $(\mathbb{F}^\star(\mathbb{M}_L^\star), \tau_d)$. Then there exists a continuous function $m : \mathbb{F}^\star(\mathbb{M}_L^\star) \to \mathbb{R}$ such that

$$F_1 \succsim F_2 \iff m(F_1) \geq m(F_2).$$

The function m is unique up to strictly increasing transformation: if m' also represents $\succsim$ continuously, then $m' = \phi \circ m$ for some strictly increasing $\phi : \mathbb{R} \to \mathbb{R}$.

Proof. This is Debreu's 1954 theorem. The construction runs as follows. Proposition 3.10 established that $(\mathbb{F}^\star, \tau_d)$ is second-countable: it has a countable base, which implies separability, so there exists a countable dense subset $\{F_n\}_{n\in\mathbb{N}} \subset \mathbb{F}^\star$. Since $\succsim$ is complete and transitive, its restriction to $\{F_n\}$ is a complete preorder on a countable set, which admits an order-isomorphic embedding into $(\mathbb{Q}, \geq)$ by a standard diagonalization argument. The resulting function on $\{F_n\}$ extends to a continuous function m on all of $\mathbb{F}^\star$ by the density of $\{F_n\}$ and the continuity of $\succsim$. Uniqueness up to strictly increasing transformation follows immediately from the order-isomorphism. ∎

The theorem delivers three things simultaneously, each with strategic content.

Existence. The Force-Maker's doctrine is not too complex to be compressed into a scalar—provided it is complete and continuous. This is not obvious. The force space $\mathbb{F}^\star$ carries the full complexity of the structured objects built in Chapter 1: nodes, bonds, command hierarchies, catalytic infrastructure, tactical bonds, and rule libraries. The doctrine ranks all of this. The theorem says that if the ranking is complete and continuous, the underlying complexity collapses to a number. The richness of the force space does not prevent compression; what prevents compression is an incoherent or incomplete doctrine, not a large space.

Uniqueness up to monotone transformation. There is no canonical force function, only a family of order-equivalent ones. If m represents $\succsim$, then so does any $\phi \circ m$ for any strictly increasing ϕ, and no other function represents $\succsim$ continuously. This is the formal statement of what practitioners know informally: that the choice between Lanchester equations, Dupuy's indices, and contest-model win probabilities is not a choice between different facts about the world but a choice between different conventional scalings of the same underlying ranking. Different scalings produce different absolute numbers; the doctrine determines the ordering, not the numbers. The sense in which force is subjective is precise: the ranking is determined by the doctrine, and the scaling is conventional.

Continuity in the Lawvere topology. The force function respects the metric geometry built in Chapter 3. Forces that are cheaply reachable from each other have close m-values; forces separated by high Lawvere distance have very different m-values. The scalar compression does not distort the strategic landscape; it is faithful to the geometry. This faithfulness follows from the continuity of $\succsim$ established in Proposition 4.2: a force function that were discontinuous in the Lawvere topology would represent an incoherent doctrine, and the theorem excludes it.

The index number problem appears at the seam between existence and uniqueness. The existence of m is a mathematical fact given the doctrine; the uniqueness up to monotone transformation means that the choice of scaling is conventional, not factual. In defense analysis, this is the problem that bedevils burden-sharing calculations, NATO capability benchmarks, and net assessment indices: different monotone rescalings produce different absolute numbers while preserving the same ordinal ranking, and arguments about which number is "right" are arguments about convention—about which features of the force space the chosen scaling should emphasize—not about the underlying doctrine. Mearsheimer's lethality-weighted index, Dupuy's Combat Power Potential, and the ONA's theater-specific metrics are valid representations of their respective doctrines, assuming completeness and continuity hold; they are not competing claims about a single strategic fact.

Properness and compact arming. The Debreu representation delivers a continuous m but says nothing about whether the preimage of a bounded interval in $\mathbb{R}$ corresponds to a well-behaved set of force configurations. The question for the strategic arming literature runs in the other direction: when the game theorist writes $f_i \in [0, \hat{f}]$, what must the force space look like for that compact interval to have an adequate foundation in the structured force space? The answer is that m must be *proper*—the preimage of every compact set must be compact.

A force function m is *coercive* if there exists a baseline force $\mathcal{S}_0$ such that for every $c \in \mathbb{R}$, the sublevel set $\{\mathcal{F} : m(\mathcal{F}) \leq c\}$ is contained in some budget ball $\mathcal{D}_{\mathcal{S}_0}(\beta)$ with $\beta < \top$. Coercivity means that the doctrine assigns unboundedly higher scores to forces farther from the baseline: there is no ceiling beyond which more capable forces are ranked indifferently. The opposing failure is a saturating doctrine, which places an effective upper bound on what any force can achieve, so that $\{m \leq \hat{f}\}$ extends arbitrarily far from any baseline and the game theorist's compact interval has no force-space anchor.

Budget balls are compact whenever R and L are both finite. With a finite rule library and a finite element type set, the number of distinct DPO rewrites applicable to any given force is finite; each rewrite that changes force structure (as opposed to rearrangements, which are isomorphisms) has strictly positive cost; and the rearrangement groupoid acts on finitely many isomorphism classes of molecules within any bounded region. Any path from $\mathcal{S}$ within total cost β therefore has bounded length, and there are finitely many such paths, making $\mathcal{D}_{\mathcal{S}}(\beta)$ a finite set—and a finite set is compact in any topology.

4.4 Proposition (Properness of the Force Function)
Let $m : (\mathbb{F}^\star(\mathbb{M}_L^\star), \tau_d) \to \mathbb{R}$ be a continuous force function representing $\succsim$. If the force space is locally finite (R and L both finite) and m is coercive, then m is proper: for every compact $K \subset \mathbb{R}$, the preimage $m^{-1}(K)$ is compact in τ_d.

Proof. Let $K \subset \mathbb{R}$ be compact; then K is bounded, so $K \subset (-\infty, c]$ for some c. By coercivity, $m^{-1}((-\infty, c]) \subset \mathcal{D}_{\mathcal{S}_0}(\beta(c))$ for some finite $\beta(c)$, and $\mathcal{D}_{\mathcal{S}_0}(\beta(c))$

is compact by local finiteness. Since m is continuous and K is closed in $\mathbb{R}$, the preimage $m^{-1}(K)$ is closed in τ_d; a closed subset of a compact set is compact.■

The strategic content of Proposition 4.4 is a reading of the compact-strategy assumption that appears throughout the strategic arming literature. The contest models of Hirshleifer (1989) and Skaperdas (1996) restrict each player's arming choice to a compact interval $[0, \widehat{f}] \subset \mathbb{R}$, and equilibrium existence follows by standard fixed-point arguments: compact convex strategy sets and continuous quasi-concave payoffs suffice. In the present framework, the game theorist's compact interval $[0, \widehat{f}]$ corresponds to the preimage $m^{-1}([0, \widehat{f}])$—a compact set of actual force configurations when m is proper. The Force-Maker who writes $f_i \in [0, \widehat{f}]$ is implicitly committing to three structural claims: (i) the doctrine is complete and continuous, so m exists by Theorem 4.3; (ii) the force space is locally finite and m is coercive, so m is proper by Proposition 4.4; and (iii) the preimage $m^{-1}([0, \widehat{f}])$ is therefore a compact set of genuine force configurations, not a scalar interval with no force-space interpretation.

What the scalar reduction loses is the structure *within* each level set $m^{-1}(c)$. Multiple distinct force configurations may share the same m-value—they are doctrinally equivalent but metrically distinct, reachable from each other at positive Lawvere distance. The contest-model equilibrium identifies an optimal force *level* but provides no guidance for selecting among the configurations that achieve it. The force-space structure visible at the level of $\mathbb{F}^{\star}$, collapsed by the projection to $\mathbb{R}$, is precisely where doctrine-relative questions about combined-arms composition, organizational form, and force employment live. The scalar arming model answers *how much*; the structured force space answers *what kind*.

4.3 Force Functions in Practice

The representation theorem gives a precise account of what any force function is doing: it is compressing a complete, continuous preorder on the force space into a scalar. The force functions that appear in the IR and military operations literature are not arbitrary choices—each implicitly encodes a specific doctrine, with specific assumptions about which features of force structure are strategically relevant and which can be ignored. This section reads four of the most prominent force functions as representations of their underlying doctrines, making explicit what completeness and continuity assumptions each requires and what each doctrine's preorder is silent about.

4.3.1 Lanchester's Laws

Lanchester's 1916 laws distinguish two modes of combat, each generating a different force function. The *linear law* governs close or hand-to-hand combat, in which each fighter engages exactly one opponent at a time: side A loses combatants at rate $-b \cdot A$ and side B at rate $-a \cdot B$, so combat power is proportional to quality times numbers and the decision-point condition is $a \cdot A = b \cdot B$. The *square law* governs modern ranged combat, in which each weapon can engage any member of the opposing force: $dA/dt = -b \cdot B$ and $dB/dt = -a \cdot A$, so the steady-state condition becomes $a \cdot A^2 = b \cdot B^2$ and the implicit force function is $m(F) = q(F) \cdot n(F)^2$, where $q(F)$ is a weapons-quality coefficient and $n(F)$ is the number of weapons in F.

The two laws encode two different doctrines about the geometry of combat. Under the linear law, doubling force doubles combat power: force scales linearly, and there is no structural advantage to concentration beyond what quality provides. Under the square law, doubling force quadruples combat power, because a larger force both inflicts casualties faster and absorbs them more slowly: the premium on mass is enormous, and the doctrine selects forces on a single dimension—quality times numbers squared—with no diminishing returns to concentration. The square law is therefore a specific and substantive claim about tactical geometry, not a neutral analytical convenience.

Both preorders are complete by construction: the formula is always defined, and any two forces with identified weapons counts and quality indices can be compared. Both are continuous in the Lawvere topology: adding one weapon changes the score by a computable increment without discontinuity. Both silence everything that Chs. 1–2 labored to represent: command hierarchy, tactical bonds, logistical depth, rule libraries, catalytic infrastructure, and the methods by which forces are employed. The Lanchester doctrine compresses the full structured force space of $\mathbb{F}^\star$ to a pair $(q, n) \in \mathbb{R}_{>0} \times \mathbb{N}$ and ranks all forces by a scalar function of that pair, treating every organizational distinction that does not show up in weapons quality or count as doctrinally irrelevant.

Biddle's 2004 critique of Lanchester-style force counting is a critique of the underlying preorder, not of the mathematics. The Lanchester doctrine ranks forces that are identical in weapons quality and numbers as equally capable regardless of how they are employed—massed or dispersed, in prepared positions or in the open, with combined-arms coordination or without it. Dupuy (1979) tested Lanchester's equations against historical data from sixty-odd engagements and found that neither law describes actual attrition reliably; casualty exchange rates depend on tactical posture, terrain, and leadership in ways the Lanchester preorder cannot represent. The disparity between German and Soviet casualties in 1941, between Israeli and Arab casualties in 1967, between Coalition and Iraqi casualties in 1991, cannot be recovered from weapons counts and quality indices alone. A doctrine that ranks by the Lanchester formula is not making a mathematical error; it is committing to a preorder that is systematically silent on the dimensions that, empirically, matter most. The representation theorem does not say the preorder is correct; it says only that a complete and continuous preorder can always be represented as a scalar, and Lanchester's preorder satisfies both conditions. The question is whether the preorder is strategically informative.

4.3.2 Dupuy's Combat Power Potential

Trevor Dupuy's 1979 Quantified Judgment Model begins where Lanchester ends: with the recognition that weapons quantity and quality alone cannot explain combat outcomes. The construction of Dupuy's core quantity—the Operational Lethality Index (OLI)— involves two steps. First, a Theoretical Lethality Index (TLI) is computed for each weapons system from engineering data: rate of fire, effective range, accuracy, and the probability and severity of incapacitation per hit. The TLI is a proving-ground measure of lethal capacity under idealized conditions of target density. Second, the TLI is converted to an OLI by dividing by the *Dispersion Factor*— the area, in square kilometers, over which one hundred thousand soldiers are typically deployed in a given era. As modern weapons have made massed formations suicidal, forces have dispersed: from roughly one square kilometer in ancient battles to several thousand in the 1973 Arab-Israeli War. The OLI of a weapons system therefore depends not only on its engineering properties but on the tactical deployment patterns of the era being modeled, making OLI values era-specific in a way TLI values are not.

The full QJM is not a weighted additive combination but a product of an additive base and multiplicative modifiers. The base is the sum of OLI values across all weapons systems in the force, $\sum_j \mathrm{OLI}_j(F)$, which gives aggregate firepower. This is then multiplied by categorical factors for terrain and weather, tactical posture (attack or defense, prepared or hasty), mobility, and vulnerability—and by an intangible factor $C(F)$ capturing leadership, training, morale, and logistics:

$$m(F) = \left(\sum_j \mathrm{OLI}_j(F) \right) \times V(F) \times M(F) \times E(F) \times C(F).$$

The multiplicative structure carries doctrinal content that a purely additive formula obscures: firepower and positional advantage are not substitutes. Doubling a force's aggregate OLI does not compensate for operating in terrain that halves the mobility factor; leadership and force employment act as proportional scalings of raw firepower, not as independent additive contributions. The preorder encoded by the QJM is not simply "more is better across all dimensions" but a specific judgment about the relative importance of each factor and the functional form of their interaction.

Dupuy calibrated the model to sixty engagements from the Italian theater of the Second World War—the Salerno, Volturno, Anzio, and Rome campaigns—using casualty data from the Historical Evaluation and Research Organization database, vetted by military historians drawn from the Army's branch schools. The Dispersion Factors, vulnerability tables, posture multipliers, and factor ranges were all set to make the model's predictions consistent with those sixty outcomes. This is the construction that the Debreu theorem illuminates: the sixty engagements are the doctrine. They specify, in quantitative terms, which features of force structure are relevant and in what proportion—how much prepared defense is worth against a hasty attack, how to weight firepower against maneuver, how far the Italian theater's experience generalizes. A different calibration—to the Pacific theater, to counterinsurgency in Vietnam, to Cold War armored warfare in Central Europe—would produce a different set of factor values and a different preorder on the same force space, even within the same QJM framework.

The index number problem appears most acutely in the intangible factor $C(F)$. Leadership, morale, training, and logistics resist the engineering-data approach that grounds the TLI; Dupuy acknowledged that the intangible factor can be estimated objectively only when prior historical data provide a calibration basis. In practice, $C(F)$ is determined post-hoc from observed casualty ratios that the measurable factors cannot explain: it is what is left over after OLI, terrain, posture, and mobility account for what they can. The choice of how to apportion unexplained variance between genuine leadership effects and measurement error is itself a doctrinal judgment. Different choices produce different values of $C(F)$, different force function outputs, and different preorders on the same force configurations—all valid as representations of their respective doctrines, and none uniquely correct.

4.3.3 Contest Success Functions

The contest model in political science and economics (Hirshleifer, 1989; Skaperdas, 1996; Tullock, 1980) uses a force function of the form $m_i(F) = F_i^\alpha / \sum_j F_j^\alpha$—the probability that side i wins the contest—or the ratio form $m(F) = F_A^\alpha / F_B^\alpha$. The contest is defined entirely in terms of aggregate resource levels, and the force function assigns higher win probabilities to larger resource concentrations. The preorder is complete—more resources is always better, and any two forces can be ranked—and continuous in any reasonable topology on the resource space.

The contest model's force space is a drastic reduction of the structured force space built in Chapter 1. The full space carries nodes, bonds, command hierarchies, catalytic infrastructure, tactical bonds, and rule libraries; the contest model reduces this to a single number per side. The preorder implicit in the contest model treats all of that structure as strategically irrelevant—every feature of force organization that the preceding chapters labored to represent drops out of the ranking. This is a substantive doctrinal commitment, and not one that the military operations literature would endorse for conventional combat between organized forces. It is, however, defensible for certain classes of political competition—lobbying, electoral resource allocation, resource extraction disputes—where the organizational structure of the parties is genuinely less important than the resources they bring to bear.

The contest model's force function is tractable and generates useful predictions about equilibrium deterrence, war onset, and bargaining breakdown. Its limitations are the limitations of its doctrine, not of the representation theory. A theorist who adopts the contest model is not making an error about mathematics; she is committing to a preorder that treats the full organizational complexity of force as irrelevant to the ranking, and the Debreu theorem shows that such a preorder is representable—complete and continuous—whatever its strategic merits.

The opening section of this book named the contest model as the canonical instance of a model with tight deductive links and a slack semantic link: m_i did its mathematical work, but what sort of object it was—soldiers, spending, capacity, power—was left undetermined, and different applications pushed different things through the same symbol without any account of what the symbol required (Section 0.1). The program announced in response was to *earn* the projection from structured force to scalar rather than to assume it. Four chapters later, the theorem we just proved makes it possible to say, precisely, what the contest model's m_i is and what doctrine it represents.

Let $\pi : \mathbb{F}^\star(\mathbb{M}_L^\star) \to \mathbb{R}_{\geq 0}$ be the *resource-extraction map* that sends each structured force to its aggregate resource level—the total quantity of personnel, platforms, and materiel the force represents, measured in the scalar units the contest model employs—discarding all information about how those resources are organized within the force's command hierarchy, rule library, and catalytic infrastructure. The pullback of the total order on $\mathbb{R}_{\geq 0}$ through π defines a preorder $\succsim_\pi$ on $\mathbb{F}^\star$: $F_1 \succsim_\pi F_2$ if and only if $\pi(F_1) \geq \pi(F_2)$. The contest success function $m_i(F) = \pi(F_i)^\alpha / \sum_j \pi(F_j)^\alpha$ is then a continuous representation of $\succsim_\pi$ in the sense of Theorem 4.3: unique up to strictly increasing transformation, with α playing the role of one such transformation. What varies across specifications of the contest form—ratio versions, additive forms, different values of α and λ—is not the underlying doctrine but the scaling of its representation. The doctrine is always $\succsim_\pi$.

The theorem also specifies, with equal precision, what $\succsim_\pi$ cannot see. Its equivalence classes are the level sets $\pi^{-1}(c)$: every structured force with aggregate resource level c ranks equally with every other, regardless of how those resources are organized. A cohesive combined-arms formation and a disorganized crowd of the same resource level, a force with a refined rule library and a force with none, a formation under unified command and a collection of autonomous units—all are doctrinally equivalent under $\succsim_\pi$. Biddle's finding that force employment explains most of the variance in combat outcomes is, in the language of the present framework, the observation that the Force-Maker's actual doctrine separates forces within $\pi^{-1}(c)$ that $\succsim_\pi$ cannot: the level sets of the resource-extraction map are too coarse to capture the distinctions that determine who wins. The contest model is entitled to those level sets; it simply cannot, after the Debreu theorem, pretend not to know what it is committing to.

The domain identification problem that Section 0.1 raised but could not answer—what makes a particular instance of the contest form a model of *force* rather than lobbying, litigation, or rent-seeking—finds its resolution in the same place. The functional form is identical across all of those domains and cannot distinguish them; only the domain of π does. When π maps actual structured forces—elements bonded into molecules, molecules organized under command hierarchies, configurations equipped with rule libraries—to aggregate resource levels, the contest game represents a competition of organized physical coercion, and the semantic link between m_i and force is secured by the construction of the preceding chapters. When π instead maps lobbying expenditures or litigation budgets to resource levels, the same mathematics operates on a domain where organized physical coercion is absent: the contest form applies, but to objects that were never structured forces in the sense of Chapter 1. The projection is the same; what differs is whether the object being projected was ever a structured force at all. Answering that question requires a theory of the domain, and Chapters 1 to 3 are that theory. The semantic link is no longer slack.

4.3.4 Net Assessment as Partial Representation

Marshall's Office of Net Assessment took a structurally different approach. Rather than specifying a single complete preorder over all forces and all mission types, the ONA produced comparative assessments targeted at specific capability pairs: submarine warfare in the North Atlantic, theater nuclear balance in Central Europe, ground force quality for a Warsaw Pact armored offensive, strategic nuclear exchange outcomes. Each assessment produced metrics relevant to that pairing, without aggregating them into a single force function applicable across all pairs and all mission domains. The ONA's product was not a scalar but a collection of scalars, each valid in its mission domain and none intended to be combined with the others into a comprehensive ranking.

In the language of the present framework, the ONA approach reflects the judgment that the Force-Maker's preorder over all force configurations is not complete. Different mission contexts generate different partial preorders—different doctrines, in the sense of Definition 4.1—each admitting a local representation by the Debreu theorem, but no single global m that ranks all pairs across all mission types. The ONA's refusal to aggregate was not methodological timidity but a principled doctrinal stance: the judgment that different mission types generate genuinely incomparable force configurations, and that forcing them into a single number destroys the strategically relevant information about what those forces can and cannot do against each other. The ONA was producing multiple local representations of multiple local doctrines, and declining to impose the completeness assumption that a global force function would require.

4.4 Emulation and Calibration

In August 1885, Chilean President Domingo Santa María hired a Prussian officer named Emil Körner to reform the Chilean military. The Chilean army had just won the War of the Pacific against Peru and Bolivia and was widely regarded as the most powerful in Latin America—the hire was not a response to weakness but to a specific aspiration: Chile wanted its force structure to look like the Prussian one (Sater and Herwig, 1999). The same dynamic was playing out in Tokyo. Körner's first choice for the Chilean position had been Jakob Meckel, who was already on loan to the Japanese government on the recommendation of the Chief of the German General Staff, Helmuth von Moltke; the Japanese had decided, after the Franco-Prussian war, that the Prussian army was the appropriate model and the previously-favored French were not (Dogauchi, 2008; Martin, 1995). Resende-Santos (2007, p. 3) records twelve emulation projects by South American states from 1870 to 1930, five targeting Germany and four targeting France. It was common in the same period for states to emulate different major powers for different force types: the German army and the British navy were separate models, each dominant in its domain, and states seeking comprehensive military reform had to navigate both.

The reforms Körner and Meckel implemented covered the full range of what Chs. 1–2 identify as the components of a structured force. They introduced the divisional and corps structure (the org chart); they replaced French and Belgian small arms with German equipment (the force molecules); they expanded the officer corps and general staff and introduced a system of military education organized around the Prussian model (the rule library and catalytic infrastructure). The Lawvere metric of Chapter 3 gives a precise meaning to the underlying aspiration: Körner and Meckel were navigating the force space, trying to move the Chilean and Japanese forces closer to the Prussian force in the Lawvere distance while staying within the political and material budget available to them.

4.5 Definition (Emulation Problem)
Let $\mathcal{S}, \mathcal{T} \in \mathbb{F}^\star(\mathbb{M}_L^\star)$ be the emulator's status quo *force and* target *force respectively, and let β be a budget element satisfying $\bot < \beta < \top$ in the quantale. The* budget ball *around $\mathcal{S}$ at radius β is*

$$\mathcal{D}_\mathcal{S}(\beta) = \{\mathcal{F} \in \mathbb{F}^\star(\mathbb{M}_L^\star) : d(\mathcal{F}, \mathcal{S}) \le \beta\}.$$

The emulation problem *is to find the force in the budget ball closest to the target:*

$$\min_{\mathcal{F} \in \mathcal{D}_\mathcal{S}(\beta)} d(\mathcal{F}, \mathcal{T}).$$

The problem is well-posed: under the local finiteness assumption that R and L are both finite, the budget ball $\mathcal{D}_{\mathcal{S}}(\beta)$ is a finite—hence compact—set, by the argument given in Section 4.2. The Lawvere metric $d(\cdot, \mathcal{T})$ is continuous in its first argument by the triangle inequality and the topology's construction, so it attains its infimum on the compact budget ball. Emulation is not wishful thinking; given the force space and its metric, the emulator can always find the best achievable approximation to the target within her means.

The role of the force function m is to identify the target $\mathcal{T}$ in the first place. Chile's doctrine ranked the Prussian force highly in m-value; the Chilean status quo occupied a lower position in the same ranking; and the emulation problem was the question of how much of that gap could be closed within the budget β. Without the doctrinal ranking—without a force function—there is no target to emulate toward, only an undifferentiated sea of force configurations with no preferred direction. Emulation is metrically efficient navigation of the force space; doctrine determines the compass bearing.

Calibration. The quantale determines what counts as "distance" in the emulation problem, and different quantales produce different emulation programs. Körner's budget was partly financial—the Chilean government provided appropriations for equipment purchases and officer salaries—but partly political, as Körner's proposed org charts met civilian resistance and had to be scaled back (Nunn, 1970). A scalar cost quantale that priced only financial expenditure would have set the wrong budget boundary: it would have allowed moves that were financially feasible but politically impossible, and excluded moves that were politically easy but financially significant. The appropriate quantale for the Chilean emulation problem was multidimensional, with separate dimensions for money, political capital, and time; the distance ball $\mathcal{D}_{\mathcal{S}}(\beta)$ was determined by the intersection of the budget constraints in each dimension.

Choosing the quantale is the *calibration* problem: identifying the cost language that faithfully prices the transformations the emulator actually faces. Dupuy's weights in the QJM (Section 4.3.2) are a calibration in this sense: they specify which features of force structure the cost language emphasizes and at what rate. A quantale calibrated to Second World War European theater data weights firepower and combined-arms coordination; one calibrated to counterinsurgency operations in urban terrain would weight civil-affairs capacity and population density differently. The emulation problem with a miscalibrated quantale produces a force that is metrically close to the target in the wrong metric—that looks Prussian in equipment while failing to achieve Prussian operational effectiveness, which is precisely what Sater and Herwig (1999) argue happened with the deeper aspects of Chilean military reform.

Multiple emulation targets. The framework extends naturally when the emulator faces multiple targets. It was common for nineteenth-century states to take the German army and the British navy as separate models, emulating each in its own domain while recognizing that neither represented a universal ideal. The emulator can either define an ideal target by combining the models—a force that is to the German army in its ground dimensions what the British navy is in its maritime ones—or proceed sequentially, first minimizing distance to one target and then using the resulting force as the new status quo for the second minimization (Nunn, 1983; Goldman and Eliason, 2003). Both approaches remain within the budget-ball framework and both have solutions under the same topological argument; the choice between them is a doctrinal judgment about whether the two targets can be reconciled into a single coherent force structure or whether they are best pursued independently.

4.5 Incompleteness and Its Consequences

The Debreu theorem gives a clean result when completeness holds. When completeness fails—when the doctrine is a partial preorder that does not rank every pair—the theorem simply does not apply, and no continuous function from $\mathbb{F}^\star$ to $\mathbb{R}$ can represent what is not there.

4.6 Proposition (Failure of Scalar Representation)
If $\succsim$ is a preorder on $\mathbb{F}^\star(\mathbb{M}_L^\star)$ that is not complete, then no continuous function $m : \mathbb{F}^\star(\mathbb{M}_L^\star) \to \mathbb{R}$ represents $\succsim$.

Proof. Let F_1 and F_2 be incomparable under $\succsim$: neither $F_1 \succsim F_2$ nor $F_2 \succsim F_1$. If m represented $\succsim$, then since $\mathbb{R}$ is totally ordered, either $m(F_1) \geq m(F_2)$ or $m(F_2) \geq m(F_1)$. But $m(F_1) \geq m(F_2)$ would imply $F_1 \succsim F_2$, and $m(F_2) \geq m(F_1)$ would imply $F_2 \succsim F_1$—in either case contradicting the incomparability of F_1 and F_2. ∎

The impossibility result does not leave the Force-Maker without resources. Richter (1966) and Peleg (1970) showed that a partial preorder on a second-countable topological space admits a *multi-utility representation*: a countable family of continuous functions $\{m_n\}_{n\in\mathbb{N}}$ such that

$$F_1 \succsim F_2 \Longleftrightarrow m_n(F_1) \geq m_n(F_2) \text{ for all } n \in \mathbb{N}.$$

The family $\{m_n\}$ collectively represents the doctrine—every comparison the doctrine makes is faithfully captured—without requiring a single scalar to do the work that no single scalar can do. Each m_n is a continuous real-valued function; together, they represent a preorder that no one of them could represent alone.

The connection to Chapter 2 is direct and resolves a question left open there. The monotones introduced in Section 2.6 of Chapter 2—the continuous functions from $\mathbb{F}^\star$ to $[0,\infty)$ that are non-decreasing along the convertibility preorder—are exactly such a family for the convertibility preorder $\succsim_{\text{conv}}$. The convertibility preorder is partial: incomparable configurations exist because not every force can be converted into every other, and the preorder has no verdict about pairs separated by the $d = \top$ boundary. No single scalar represents it. But the countable family of monotones does: $F_1 \succsim_{\text{conv}} F_2$ if and only if $m_n(F_1) \geq m_n(F_2)$ for every monotone m_n in the family, and when F_1 and F_2 are incomparable the family separates them by producing an m_n with $m_n(F_1) \ngeq m_n(F_2)$. The Richter-Peleg theorem is the bridge between the monotones of Chapter 2 and the doctrine of the present chapter: the monotones *are* the multi-utility representation of the convertibility preorder, and that preorder is the doctrine implicit in the category **Force** itself.

The scaling problem. The multi-utility representation solves the existence problem—an incomplete doctrine is still representable, by a family rather than a scalar—but introduces an aggregation problem. Any attempt to combine the family $\{m_n\}$ into a single number requires a weighting rule: some function $M(F) = \phi(m_1(F), m_2(F), \ldots)$ that aggregates the individual utilities into a common score. Every such weighting rule is a doctrinal choice. It specifies how to trade off gains in one dimension of capability against losses in another—how to weigh anti-submarine capability against ground-force quality, or theater nuclear balance against conventional maneuver advantage—and different Force-Makers will make that trade differently depending on what they take to be strategically decisive. The Richter-Peleg theorem shows that the individual dimensions are representable; it provides no basis for choosing among weighting rules, because the choice of weighting rule is not a mathematical question but a strategic one. The ONA's refusal to aggregate is, in this light, the judgment that no weighting rule is defensible across all strategic contexts—that the individual m_n carry information that any aggregation ϕ would destroy, and that the information-loss is strategically unacceptable.

Three doctrinal responses to incompleteness. The Force-Maker who faces an incomplete doctrine has three available strategies, each with its own costs and characteristic failure modes.

The first strategy is to extend the doctrine: to add comparison rules that cover the previously incomparable pairs, converting a partial preorder into a complete one. If the doctrine does not rank counterinsurgency-configured forces against conventional combined-arms formations, the Force-Maker can attempt to develop criteria that make the comparison. The cost is coherence: new comparison rules may conflict with existing ones, producing a doctrine that satisfies neither completeness nor transitivity—a preorder with cycles or contradictions that defeats the purpose of the extension. The post-2001 US effort to develop a doctrine of counterinsurgency, culminating in FM 3-24, was an attempt to extend the doctrine in exactly this sense, and the persistent difficulty of reconciling counterinsurgency force requirements with conventional force requirements reflects the coherence cost that extension imposes.

The second strategy is to restrict the domain: to confine each doctrine to a subdomain of force configurations where completeness holds, and apply different doctrines in different mission contexts. This is the ONA approach and the logic of mission-specific assessment indices. The cost is boundary specification: the seams between doctrinal domains must be identified, and force comparisons that span those seams—especially in coalition operations, where allies bring doctrines calibrated to different mission types—are left without a common ranking. Coalition warfare exposes seams; the inability to rank allies' force configurations in a common preorder is the doctrinal cost of the restricted-domain strategy.

The third strategy is to impose completeness artificially: to adopt a force function that is always defined—a Lanchester score, a contest-model win probability, a GDP-per-soldier ratio—and accept that it misrepresents force comparisons outside its intended domain. The cost is systematic error: the formula produces numbers for configurations it was not designed to rank, and those numbers are artifacts of the formula's algebraic form rather than representations of any genuine doctrine. The history of military assessment is largely a history of this strategy's costs: Lanchester scores applied to counterinsurgency operations, casualty-exchange ratios applied to nuclear deterrence, aggregate defense expenditure applied to alliance burden-sharing. Each produces numbers; each systematically misrepresents the configurations it was not designed to rank.

The United States military, 2003. The US military's entry into Iraq in 2003 illustrates all three dynamics in sequence, and illustrates the underlying representational problem that drives them. The existing conventional force doctrine was a complete and well-calibrated preorder over the space of conventional combined-arms configurations: it could rank any two mechanized formations, specify the force ratios required for offensive operations, and compute the logistics requirements for sustained maneuver. That doctrine did not extend to post-conflict stabilization operations or counterinsurgency: the force types required—civil affairs units, trained interpreters, police advisory teams, light infantry configured for community engagement—occupied a region of the force space that the doctrine's preorder had never been required to cover. The result was not simply poor planning but a representational failure in the precise sense of Proposition 4.6: the force function that operational commanders and Washington policymakers reached for to assess adequacy, progress, and success did not exist for the force configurations the situation required, because no complete preorder over those configurations had ever been developed. Extending the doctrine took years and was incomplete; restricting the domain produced coordination failures as conventional and counterinsurgency forces operated under incompatible preorders; and the artificial imposition of conventional force metrics on stabilization operations—area controlled, enemy killed, checkpoints established—generated systematic misrepresentation of the strategic situation. The Debreu theorem does not prescribe strategy; it diagnoses what was missing. The missing element was a doctrine that covered the relevant region of the force space.

4.6 Looking Ahead

The book's argument is now complete. Chapter 1 built the objects: the structured force space $\mathbb{F}^\star(\mathbb{M}_L^\star)$, constructed recursively from elements to molecules to configurations to organizations, each force a structured graph carrying command hierarchy and tactical bonds. Chapter 2 built the morphisms: the category **Force**, equipped with DPO rewriting, catalytic transformation, and organizational restructuring, encoding which transformations are valid and how they compose, and generating the convertibility preorder and the rearrangement groupoid as consequences of the categorical structure. Chapter 3 built the geometry: a cost functor valued in a quantale, inducing a Lawvere metric on the force space and, from that metric, a second-countable topology in which proximity means cheapness, connected components mean reachability, and second-countability means that every capability target can be approached through cheaper intermediate forces. The present chapter has built the fourth layer: the Force-Maker's doctrine, her preorder on the force space, and the conditions under which that preorder compresses into the scalar force function she uses to navigate the geometry built in Chapter 3.

The four layers are not independent. The second-countability of the Lawvere topology is what makes the Debreu theorem applicable: without it, no representation theorem holds even for complete preorders on continuous spaces. The convertibility preorder of Chapter 2 is a partial preorder on the same force space, and the Richter-Peleg multi-utility theorem shows that the monotones introduced there are not an ad hoc construction but the unique appropriate representation of a partial doctrine on a second-countable space. The force function m of the present chapter ranks the positions in the geometry; the Lawvere metric of Chapter 3 prices the paths between them; the morphisms of Chapter 2 define the paths; and the objects of Chapter 1 are the positions themselves. The four constructions form a single coherent structure, and each layer was built to support the one above it.

Return, finally, to the epigraph. Weil wrote that only he who has measured the dominion of force, and knows how not to respect it, is capable of love and justice. The force function m is the measurement: the compression of the Force-Maker's doctrine into the scalar by which she plans, acts, and judges. Knowing how not to respect it is the recognition that the measurement is doctrine-relative—that every force function encodes a preorder, that the preorder can be more or less complete, and that the claim to measure force precisely is always a claim about doctrine, not about the world. The Force-Maker who mistakes her force function for a fact—who treats Lanchester scores or contest-model probabilities as truths rather than as doctrinal representations, conventional up to a monotone rescaling—cannot see the incompleteness of her preorder, cannot extend it to configurations her doctrine does not cover, and cannot recognize when the strategic situation has moved outside the region her doctrine was built to rank.

The measurement, taken rightly, does not preclude judgment; it structures it, and the structure it imposes is the structure of force itself.

Chapter 5

Conclusion

Nothing the peoples of Europe have produced is worth the first known poem that appeared among them. Perhaps they will yet rediscover the epic genius, when they learn that there is no refuge from fate, learn not to admire force, not to hate the enemy, nor to scorn the unfortunate. How soon this will happen is another question.

Simone Weil, "The Iliad, or the Poem of Force"

The book opened with a problem: the contest model has tight deductive links but a slack semantic link, and the variable m_i slides from soldiers to spending to capacity to power without ever specifying what force is. Four chapters have built the answer. The construction is now complete; this chapter draws out what it accomplishes, states its limitations clearly, points toward the work that remains, and closes with the reflection that motivated the enterprise.

5.1 The Earned Projection

The book's central claim is not that force should be measured by a scalar, but that the scalar is earned only after four layers of structure are in place. Each layer built what the next required, and the dependency runs in only one direction.

Chapter 1 built the objects. The structured force space $\mathbb{F}^{\star}(\mathbb{M}_L^{\star})$ was constructed from the ground up: elements (nodes carrying labels from a type set L) bond into molecules (typed graphs), molecules gather into configurations, and configurations are organized under command hierarchies. Without a specified force space, there is no domain for the morphisms, no domain for the cost functor, and no domain for the Force-Maker's preorder. The recursive construction also settled the countability question: the force space is countably infinite, which will matter for the topology and, through the topology, for the representation theorem.

Chapter 2 built the morphisms. Double-pushout rewriting, catalytic transformation, and organizational restructuring together define which transformations of force are valid and how they compose. These morphisms do more than specify what the Force-Maker can do; they generate two distinct mathematical structures as consequences of the categorical architecture: the convertibility preorder (which forces are reachable from which, via valid transformation sequences) and the rearrangement groupoid (which configurations are freely expressible through pure reorganization, at zero cost). The Richter-Peleg family of monotones that characterizes the convertibility preorder is not an ad hoc construction imported from outside the theory; it is the representation of a partial doctrine on a second-countable space, as Chapter 4 established.

Chapter 3 built the geometry. A quantale-valued cost functor on **Force** assigns costs to morphisms in a way that respects identity (doing nothing costs nothing) and triangle inequality (the cost of a path is at least the cost of its shortest sub-path). The infimum over all valid paths between two forces defines the Lawvere metric on the force space, and from that metric we constructed a second-countable topology in which proximity means cheapness, connected components mean reachability, and forward balls have a countable base. The second-countability is the topological hypothesis that Debreu's representation theorem requires; without it, complete preorders on continuous spaces can fail to admit any real-valued representation at all.

Chapter 4 built the doctrine. The Force-Maker's ranking of force configurations is a preorder on the force space; Debreu's (1954) theorem says that a complete, continuous preorder on the second-countable Lawvere topology admits a continuous real-valued representation, unique up to a strictly increasing transformation. The force function $m : \mathbb{F}^{\star} \to \mathbb{R}$ is that representation: it compresses the doctrine into the scalar by which the Force-Maker plans, acts, and judges. When the doctrine is not complete, no scalar suffices, but a countable family of continuous functions represents the partial preorder (Richter, 1966; Peleg, 1970).

The projection from structured force to scalar is earned in a specific sense. We can now say exactly what the scalar $m(F)$ means—a continuous representation of a complete continuous preorder on a second-countable space—exactly what it costs to compute (the index number problem: the choice of scaling among order-equivalent representations), exactly when it fails (when the preorder is not complete), and exactly what structures it discards (all the organizational and morphological detail that the level sets of m identify as doctrinally equivalent). The contest model's m_i is a legitimate force function exactly when the resource-extraction doctrine $\succsim_\pi$ is a complete preorder on $\mathbb{F}^\star$—when the Force-Maker genuinely ranks all force configurations by aggregate resource content alone. Sometimes this is an appropriate doctrinal assumption, as in large-scale industrial mobilization where the decisive variable is indeed production capacity. More often, it is not. The construction does not prohibit the simplification; it specifies the conditions under which the simplification is faithful and the conditions under which it is not.

5.2 What the Framework Sees

The construction resolves several puzzles from the existing literature and makes several informal debates precise.

The contest model's three adequacy conditions. Section 0.1 identified three adequacy conditions for a theory of force: semantic consistency (does m_i mean the same thing across applications?), structural fitness (is $m_i \in \mathbb{R}_{\geq 0}$ an appropriate representation of force?), and domain identification (what distinguishes a contest of force from a contest of lobbying or rent-seeking?). The construction answers all three. Semantic consistency is secured by the resource-extraction map $\pi : \mathbb{F}^\star(\mathbb{M}_L^\star) \to \mathbb{R}_{\geq 0}$, which gives m_i a precise domain: it denotes the aggregate resource content of a structured force, a real object built from elements, molecules, and configurations in the sense of Chapter 1. Structural fitness is a question about the doctrine $\succsim_\pi$: whether it is complete and continuous on the force space, which the framework now allows us to assess. Domain identification is settled by the construction of π itself: the contest is a contest of force—not lobbying or rent-seeking—exactly when π maps objects from $\mathbb{F}^\star(\mathbb{M}_L^\star)$, organized means of physical coercion in the sense of Section 0.3. The semantic link is no longer slack.

The Mearsheimer-Biddle debate. Mearsheimer (2001) ranks great powers by a common measure of latent power and military assets; his theory implicitly treats the doctrine as complete over all relevant force configurations, so that any two forces can be placed in a determinate ordinal relation. Biddle (2004) argues that military effectiveness depends on force employment—how forces are used, not merely how large they are—and that forces configured for different tasks cannot be ranked without specifying the mission. The framework makes this disagreement precise: Mearsheimer and Biddle are disputing whether the Force-Maker's preorder is complete across the relevant region of the force space. This is not a definitional dispute but a doctrinal one, with empirical content. If Biddle is right, the doctrine that underlies any Mearsheimer-style assessment is incomplete over mission-type boundaries, and the force function that generals and policymakers reach for does not exist for the configurations that matter most in many real conflicts. The Debreu theorem does not resolve the debate; it specifies what kind of evidence would. Completeness is a claim about the doctrine, and the doctrine can be interrogated: does the Force-Maker have a ranking of combined-arms conventional forces against counterinsurgency-configured light infantry, and does that ranking hold consistently across contexts? If the answer is no, the doctrine is incomplete, and the framework diagnoses the failure.

The index number problem. Every force assessment that uses a scalar ranking commits to a specific representation of a specific doctrine, unique only up to a monotone rescaling. Arguments about whether NATO defense spending benchmarks, IISS military balance indices, or Dupuy's combat power scores are "right" are arguments about the scaling, not about the doctrine. The scaling is conventional, not factual. Once this is seen, the methodological question shifts from "which number is correct?" to "which representation of the same doctrine best serves the empirical purpose at hand?"—a question about information sufficiency and measurement efficiency, not about which formula captures the truth. Different scalings will give different absolute numbers but the same ordinal rankings; the ordinal rankings are the doctrine, and the doctrine is where the substantive disagreement lies.

Emulation and the Lawvere metric. The cost geometry of Chapter 3 gives the emulation problem a precise formulation. When the Force-Maker observes an adversary's force F_2 and judges it superior to her own F_1, she faces the problem of finding the cheapest sequence of valid morphisms from F_1 to a force isomorphic to F_2. The Lawvere distance $d(F_1, F_2)$ is the infimum cost of that sequence. The force function m gives her a target: she wants to reach a force whose m-value matches or exceeds F_2's. The metric gives her the cost: the Lawvere distance from her current position to that target. The two together determine the Force-Maker's planning problem—where she is ranked, where she wants to be, and what it costs to get there—and the emulation result of Chapter 4 shows that the target is always approachable through cheaper intermediate forces, provided the topology is second-countable and the force space is locally finite.

Foundations for arming models. Game-theoretic models of militarization use force functions as payoff inputs, and Nash equilibrium existence often requires that the strategy set be compact or that the payoff function satisfy compensating conditions. The properness result of Proposition 4.4 gives a sufficient condition grounded in the construction: if the force function is coercive (sublevel sets are contained in budget balls) and the force space is locally finite (R and L both finite), then m is proper—the preimage of any compact set is compact in the Lawvere topology. The game theorist's compact strategy set $[0, \widehat{f}]$ has a foundation in the structured force space: it is the image under m of the budget ball $\mathcal{D}_{\mathcal{S}_0}(\beta)$, which is compact by local finiteness. The abstraction that arming models require is not arbitrary; it is available exactly when the construction's local finiteness condition holds, and the condition is empirically interpretable: finite rule library R, finite element type set L.

5.3 Objections

Three objections deserve explicit treatment.

The abstraction objection. The construction uses category theory, graph rewriting, quantales, and point-set topology. Most international relations scholars will not find this notation accessible, and most military practitioners will not consult it. What is gained that could not have been gained more cheaply?

The reply has two parts. The first is that abstraction is not optional; it is the price of precision. Every time an IR theorist writes "force is a scalar," she is making tacit claims about completeness, continuity, and the structure of the force space—claims that the framework makes explicit and that can fail in identifiable ways. The US Army's failure in 2003 was not a failure of character or will; it was a structural failure, diagnosable in the terms of Proposition 4.6: the doctrine had no preorder over counterinsurgency-configured forces, so no force function existed for the configurations the situation required. An approach that cannot say this clearly is less precise, not simpler. The second part of the reply is that the construction interfaces with existing models. The contest model, Lanchester's laws, and Dupuy's QJM all appear in the framework as representations of specific doctrines on specific domains. The abstraction does not replace those models; it specifies when they are appropriate and when they are not, and it gives the theorist who uses them a principled basis for the choice.

The formalization objection. Military doctrine is tacit, contextual, and resistant to formalization. Practitioners develop doctrine through experience and institutional culture, not through explicit axiomatization. The preorder model assumes that doctrine has a clean mathematical structure it does not actually possess.

The reply is that the framework characterizes what a doctrine must be in order to support a force function—not what any particular doctrine is. The claim is not that US Army doctrine in 2003 was a complete preorder; Ch. 4 showed that it clearly was not. The claim is that the incompleteness has a formal diagnosis: a region of the force space that the doctrine's preorder had never been required to cover. This diagnosis generates actionable prescriptions—extend the doctrine, restrict the domain, or accept the costs of artificial completeness—and the costs of each prescription are assessable in the framework's terms. Making the structure explicit is not a claim that doctrine is simple; it is a tool for identifying precisely where doctrine is incomplete and what it would cost to extend it. The tacit knowledge that practitioners carry can be partially reconstructed from the doctrine's revealed behavior: which forces it ranks, which comparisons it refuses to make, and where its preorder runs out.

The prediction objection. The framework says nothing about who wins. International relations scholars want predictive models of war outcomes, deterrence stability, and alliance formation; the construction provides an ontology. Ontology, it may be objected, is not science.

The reply is that the objection is precisely right, and this is the point. The construction is an ontological foundation, not a predictive model—and the distinction matters because existing predictive models rest on ontological commitments that have never been made explicit or critically examined. When a Lanchester model predicts attrition rates, it assumes a specific doctrine, a specific force space reduction, and a specific continuity condition. When those conditions fail, the predictions fail, and they fail in identifiable ways that the framework can locate. The framework's contribution is not prediction but diagnosis: it tells the theorist what she is assuming when she writes down a force function, which assumptions are load-bearing, and which failure modes are in principle avoidable. Prediction built on an explicit ontological foundation is more reliable than prediction built on a tacit one, because the conditions of the prediction's validity are known and can be checked.

5.4 Extensions

The construction opens rather than closes.

Alliance and the monoidal product. The symmetric monoidal structure of **Force** already accommodates alliance. The tensor product $F_1 \otimes F_2$ represents two forces operating together under a common command interface, and the cost functor extends to the tensor product by the standard monoidal enrichment. The cost of an alliance, in the Lawvere metric, is the cost of constructing the tensor product minus the costs of each component maintained separately—a quantity that may be negative when combination achieves things neither force could achieve alone (the gains from specialization that motivate formal alliances) and positive when joint operation introduces overhead (command friction, doctrine reconciliation, communications compatibility). Whether an alliance is worth forming is a question about this difference: does the joint capability gain, as measured by the force function applied to the tensor product, exceed the transformation cost? The framework gives that question a geometry.

Arms racing as emulation. Two Force-Makers, each navigating their own Lawvere metric on the shared force space, each attempting to minimize the Lawvere distance between their current force and the adversary's current or anticipated force—this is the arms race as emulation dynamic. Stability in the race corresponds to the existence of a fixed point in the joint emulation problem: a pair of forces $(F_1^\star, F_2^\star)$ such that neither Force-Maker has a profitable unilateral deviation in the metric. Instability—the arms spiral—corresponds to the absence of such a fixed point, which can arise when the two Force-Makers' Lawvere metrics are incompatible: when each perceives the other as reachable at low cost while perceiving her own position as vulnerable at high cost. Kennedy's (1989) hegemonic overstretch is the condition in which the cost of maintaining the emulation path exceeds the sustainable resource base—the force budget ball $\mathcal{D}_{\mathcal{S}_0}(\beta)$ must grow to track the adversary, but the quantale value of that growth exceeds the economic quantale that the state can sustain. The metric framework makes the mechanism precise.

Technological change as morphism extension. A military-technological innovation adds new rules to the rule library R, extending the morphism set of **Force** and, with it, the Lawvere metric. The distinction between two types of innovation is now formally available. A *cost-collapsing* innovation makes previously expensive paths through the force space cheap: the metric changes, but the reachable frontier from any starting point is not enlarged—the same targets are now attainable at lower cost. A *frontier-expanding* innovation opens new regions of the force space that were previously unreachable under any finite budget: new elements (unmanned systems, directed energy, cyber capabilities) generate new molecules and configurations that no sequence of existing morphisms could produce. The distinction matters strategically: cost-collapsing innovations favor the offense when the defender cannot afford to maintain the now-cheaper emulation path; frontier-expanding innovations confer durable advantage when the new region of the force space is outside the adversary's doctrinal preorder entirely, leaving her without a representation of what she is facing. Much of the revolution-in-military-affairs literature has debated this distinction implicitly, without the vocabulary to make it precise.

Non-state actors and the boundaries of the theory. The construction nowhere requires the Force-Maker to be a state. Police forces, militias, private military companies, and insurgent movements are all recursively structured (they have elements, molecules, configurations, and command hierarchies), all subject to valid morphisms, all operating under cost constraints, and all guided by doctrines—explicit or implicit—that rank force configurations against each other. The state/non-state distinction appears in the construction as a difference in resource base, rule-library stability, and organizational depth, not as a categorical distinction in the mathematics. This means the framework applies to the full range of organized physical coercion, and the question of when non-state violence differs structurally from state violence is now a tractable formal question: it is a question about which rule libraries each Force-Maker has access to, what quantale prices her cost functor, and whether her doctrine is complete over the region of the force space her operations require. These are empirical questions about specific actors and specific conflicts, answerable in principle from the structure of the construction.

5.5 The Measurement

There is something humiliatingly beautiful about force—about the mathematics of it. The fractal-like structure of the force space, the elegance with which the category **Force** captures the grammar of military transformation, the way the Lawvere metric turns the strategic landscape into a geometry that can be navigated: these are genuinely lovely things. The convertibility preorder emerging from DPO rewriting without being stipulated, the Debreu theorem fitting exactly where the topology was laid, the Richter-Peleg family showing that the monotones of Chapter 2 were always already the representation of a partial doctrine—one does not work through this material without some pleasure in it.

But this is what Weil guards against. Not the pleasure of the theory—Weil was herself a mathematician—but the slide from pleasure in the structure to admiration of the thing itself. The theory is beautiful; force is not. To mistake the elegance of the representation for an endorsement of what is represented is to commit the error that the theory was designed to prevent: confusing the map for the territory, the doctrine for a fact, the force function for a feature of the world. Weil's epigraph looks forward, and the book has been an attempt to earn that forward look, or at least to prepare the ground for it.

There is no refuge from fate because there is no force configuration that is unconditionally superior. The force space is vast, the morphisms of **Force** are numerous, and what dominates under one doctrine may be irrelevant or inferior under another. A force that is maximally ranked by the resource-extraction preorder $\succsim_\pi$ may be poorly ranked under a doctrine that weights organizational depth, catalytic infrastructure, or rule-library breadth. A force that is locally dominant in conventional combined-arms operations may have no ranking at all under a doctrine designed for counterinsurgency or deterrence stability. The Force-Maker who believes she has found a permanent safe harbor in a particular configuration misreads the geometry: the Lawvere metric always poses the question of whether an adversary can close the distance, and the answer depends on the adversary's quantale, her morphisms, and her doctrine—not on any intrinsic property of the configuration itself.

Not to admire force is to resist the error of mistaking the map for the territory. A force function is a representation of a doctrine, not a discovery of nature, and the high m-value of a great power's military is a doctrinal judgment, revisable by a different preorder, rescalable by a monotone transformation, and limited in validity to the region of the force space the doctrine was built to cover. The construction has been designed precisely to foreground this: by making the doctrine explicit as a preorder, by proving that the scalar representation is unique only up to monotone rescaling, and by showing that a force function defined for one region of the force space may generate artifacts—numbers without doctrinal content—when applied to another. The admiration of force, formalized, is the error of treating the index number as a fact rather than a convention, the error of confusing the representation with the thing it represents.

Not to hate the enemy is to understand that the adversary also occupies a position in the force space, also navigates the Lawvere metric, and also operates under a doctrine that may be more or less complete. Understanding the adversary's doctrine is understanding where she is in the space and which comparisons her preorder can make: which configurations she can rank, which she cannot, and where her preorder runs out. This is not sympathy; it is analysis. The net assessment tradition understood this, which is why the ONA avoided aggregate scorecards in favor of mission-specific capability comparisons—an implicit recognition that the adversary's doctrine was a partial preorder, not a complete one, and that forcing it into a single number would destroy the strategic information that mattered.

Not to scorn the unfortunate is to recognize that the forces that appear weak under a prevailing doctrine may occupy regions of the force space that the doctrine cannot adequately rank. The incompleteness of a preorder is not an indictment of the force configurations it cannot compare; it is an indictment of the doctrine that fails to cover them. The civilian population of an occupied territory, the militia without a conventional order of battle, the state too small to appear in aggregate capability indices—these are not absent from the force space. They are present in regions that the prevailing doctrine's preorder has never been required to reach. Scorning them is the error of treating the absence of a representation as an absence of the thing itself.

This book has tried to measure the dominion of force. The measurement is partial—four layers of structure are in place, and the extensions of Section 5.4 remain to be built—but it is precise about what it measures and honest about what it leaves out. Whether the measure is good enough to know how not to respect force—that judgment belongs to the Force-Maker. She already knows the answer; we have been trying to keep up.

Appendix A

Mathematical Preliminaries

This appendix collects the mathematical background needed to follow Chapter 1–Chapter 4. The reader is assumed to be comfortable with undergraduate calculus and linear algebra, and with the set-theoretic vocabulary that appears in a first graduate methods course: sets, functions, relations, and countability.

What the reader is *not* expected to know is the language of modern pure mathematics: graph theory, category theory, lattice theory, topology, or order theory beyond simple partial orders. These are the tools that make the book's central results precise—that force configurations are objects in a symmetric monoidal category (Chapter 2), that the cost of transformation is a Lawvere metric (Chapter 3), and that a complete continuous doctrine on a second-countable space admits a continuous scalar representation (Chapter 4). The primer introduces each tool from scratch and then shows exactly how it will appear in the main text, so that when the reader encounters a definition or proposition in context, its mathematical content is already familiar.

Throughout this appendix, L denotes the element type set, $\mathbb{M}_L$ the set of force molecules, $\mathbb{F}^\star(\mathbb{M}_L^\star)$ the structured force space, **Force** the category of forces and transformations, $Q = (\mathbb{R}_{\geq 0} \cup \{\infty\}, \leq, +)$ the textbook quantale of costs, d the Lawvere metric, τ_d the Lawvere topology, $\succsim$ a doctrine, and m a force function. When a letter like X appears in a general definition (e.g., "let (X, d) be a Lawvere metric space"), the generic usage will be clear from context.

A.1 Graphs and Graph Morphisms

Force molecules—the fundamental units of force structure—will be defined in Chapter 1 as connected labeled graphs. Understanding what that definition requires is the purpose of this section.

A.1 Definition (Graph)

A graph *is a pair* $G = (V, E)$ *where* V *is a nonempty finite set of* nodes *and* $E \subseteq \binom{V}{2}$ *is a set of* edges, *each edge being an unordered pair* $\{u, v\}$ *of distinct nodes. A* directed graph *replaces unordered pairs with ordered pairs* (u, v), *distinguishing the* tail u *from the* head v.

Graphs model force by treating nodes as military elements and edges as the bonds between them: tactical bonds (flanking relationships, interlocking fields of fire), logistical bonds (supply chains), and command bonds (subordination).

A.2 Definition (Labeled graph)

A labeled graph *over a type set* L *is a triple* $G = (V, E, \ell)$ *where* (V, E) *is a graph and* $\ell : V \to L$ *is a* labeling function *assigning each node a type.*

A.3 Example (Force molecules as labeled graphs)

Chapter 1 will take L *to be the element type set: a countable index set whose members include labels such as* `SOLDIER`, `TANK`, `AIRCRAFT`, *and* `LOGISTICS_UNIT`. *The notion of a* force molecule *that Definition 1.4 will introduce is precisely a connected labeled graph* $M = (n, E, \ell : \underline{n} \to L)$, *where* n *is the number of nodes and* $\underline{n} = \{1, \ldots, n\}$ *is the node index set. A two-node graph with nodes labeled* `SOLDIER` *and* `TANK` *and a single edge between them will be the minimal combined-arms molecule: a soldier-tank bond. A three-node graph with all nodes labeled* `HOPLITE` *and all three edges present will be the hoplite phalanx molecule illustrated in the TikZ figures of Chapter 1.*

A.4 Definition (Graph isomorphism)

Two labeled graphs $G_1 = (V_1, E_1, \ell_1)$ *and* $G_2 = (V_2, E_2, \ell_2)$ *over* L *are* isomorphic, *written* $G_1 \cong G_2$, *if there exists a bijection* $\varphi : V_1 \to V_2$ *such that* $\{u, v\} \in E_1$ *iff* $\{\varphi(u), \varphi(v)\} \in E_2$, *and* $\ell_2(\varphi(v)) = \ell_1(v)$ *for all* v.

A.5 Example (Molecules up to isomorphism)

Whether the nodes of the hoplite molecule are numbered $1, 2, 3$ *or* a, b, c *is irrelevant to its military content: what matters is the type pattern and the bond structure. This is why Primitive 1.6 will define* $\mathbb{M}_L$ *as the set of all force molecules* up to isomorphism*:* $\mathbb{M}_L$ *counts structural types, not labeled instances. A subsequent countability lemma will then establish that* $\mathbb{M}_L$ *is countably infinite, because the set of connected labeled graphs on a countable type set is countable.*

A.6 Definition (Connected graph and components)

A graph $G = (V, E)$ *is* connected *if for every pair of distinct nodes* $u, v \in V$*, there exists a* path *from* u *to* v*: a sequence* $u = w_0, w_1, \ldots, w_k = v$ *with* $\{w_i, w_{i+1}\} \in E$ *for each* $i < k$*. A* connected component *of a graph is a maximal connected subgraph.*

A.7 Example (Connectivity and the molecule/configuration distinction)

Chapter 1 will require that a force molecule be connected: the elements must form a single cohesive unit. A force configuration *will be by contrast a* disconnected *labeled graph—the disjoint union of molecules. The key structural fact (Lemma 1.12) will be unique decomposition: every configuration decomposes into its connected components, each of which is a molecule. The graph union* $\uplus$ *of Definition 1.10 produces configurations by concatenating node index sets with shifted edges, ensuring that the two components acquire no spurious bonds.*

A.8 Definition (Graph homomorphism)

A graph homomorphism $f : G_1 \to G_2$ *is a function on node sets such that* $\{u, v\} \in E_1$ *implies* $\{f(u), f(v)\} \in E_2$*, and* $\ell_2(f(v)) = \ell_1(v)$ *for all* v*. An* injective *homomorphism (injective on nodes) is an* embedding*. A bijective homomorphism whose inverse is also a homomorphism is an isomorphism.*

A.9 Example (Matches as embeddings)

When the Force-Maker searches her current force graph G for a tactical pattern L, she is looking for an injective graph homomorphism $m : L \hookrightarrow G$ *— an embedding of the pattern into the host. Chapter 2 will formalize this as the* match *in a DPO rewrite step (Definition 2.4): the match* m *identifies where in G the rule's left-hand side L appears, and injectivity ensures that distinct nodes of the pattern are matched to distinct nodes of the host.*

A.10 Definition (Hypergraph)

A hypergraph *over* V *is a pair* $(V, \mathcal{E})$ *where* $\mathcal{E} \subseteq 2^V \setminus \{\varnothing\}$ *is a collection of* hyperedges*, each being a nonempty subset of* V *of any size. An ordinary graph is the special case in which every hyperedge has size exactly two.*

A.11 Example (Non-command relations)

Chapter 1 will define the non-command relations of a structured force as a collection $R = \{H_1, \ldots, H_k\}$ *of hypergraphs on the unit set* V *(Definition 1.14). A fire-support relationship involving three units simultaneously is a hyperedge of size three; it cannot be represented as a set of pairwise edges without loss of information, because the joint relationship is not reducible to any pair. Hypergraphs are needed wherever group relationships among force units are irreducibly collective.*

A.2 Categories, Functors, and Monoidal Structure

Chapter 2 will build a category **Force** whose objects are structured forces and whose morphisms are finite sequences of valid transformations. The language of categories is the language of "things and the ways of transforming them," and it is precisely the right language for a theory of force dynamics.

A.12 Definition (Category)

A category $\mathcal{C}$ *consists of: a collection of* objects*; for each pair of objects* A, B*, a set* $\mathcal{C}(A, B)$ *of* morphisms *from* A *to* B*; for each composable pair* $f \in \mathcal{C}(A, B)$ *and* $g \in \mathcal{C}(B, C)$*, a* composite $g \circ f \in \mathcal{C}(A, C)$*; and for each object* A*, an* identity $\mathrm{id}_A \in \mathcal{C}(A, A)$*. These must satisfy* associativity $(h \circ g) \circ f = h \circ (g \circ f)$ *and* unitality $\mathrm{id}_B \circ f = f = f \circ \mathrm{id}_A$*.*

Three examples orient the concept before we turn to **Force** itself.

The category Set. Objects are sets; morphisms are functions; composition is function composition; the identity at A is $\mathrm{id}_A(x) = x$. This is the most familiar category and the one in which the pushout constructions of Section A.3 are easiest to visualize.

The category $\mathcal{G}$ of labeled graphs. Objects are finite labeled graphs over L; morphisms are label-preserving graph homomorphisms; composition is function composition on node sets; the identity is the identity function. This is the category in which double-pushout rewriting will take place: the rules, matches, and rewritten graphs of Chapter 2 will all be objects and morphisms of $\mathcal{G}$.

A preorder as a category. Any preorder $(X, \precsim)$ defines a category: objects are elements of X; there is exactly one morphism from a to b when $a \precsim b$, and none otherwise; transitivity gives composition; reflexivity gives identities. The convertibility preorder $\precsim_{\mathcal{R}}$ that Chapter 2 will introduce is exactly such a category, recording reachability between forces.

A.13 Example (The category Force)

Chapter 2 will define **Force** *as exactly a category of this kind (Definition 2.13). Objects will be members of $\mathbb{F}^\star(\mathbb{M}_L^\star)$, the structured forces built in Chapter 1. A morphism from F_1 to F_2 will be a finite sequence of valid transformation steps—DPO rewrites, catalytic transformations, and organizational restructurings—that carries F_1 to F_2. Composition is concatenation of sequences; the identity at F is the empty sequence. Verifying the category axioms for this data is the content of Proposition 2.14: associativity holds because sequence concatenation is associative, and unitality holds because prepending or appending the empty sequence leaves any sequence unchanged.*

A.14 Definition (Functor)

A functor $F : \mathcal{C} \to \mathcal{D}$ *assigns to each object A an object $F(A)$, and to each morphism $f \in \mathcal{C}(A, B)$ a morphism $F(f) \in \mathcal{D}(F(A), F(B))$, satisfying $F(\mathrm{id}_A) = \mathrm{id}_{F(A)}$ and $F(g \circ f) = F(g) \circ F(f)$.*

A.15 Example (The cost functor)

Chapter 3 will introduce a cost functor $c : \mathrm{Mor}(\mathbf{Force}) \to Q$ *(Definition 3.2), a function on morphisms satisfying two axioms: $c(\mathrm{id}_F) = \bot$ (identity morphisms have zero cost) and $c(\psi \circ \phi) \leq c(\phi) \otimes c(\psi)$ (composing transformations costs at most the sum of their individual costs, with strict inequality possible when combination creates synergy). This is not quite a functor in the strict sense—it maps into the quantale Q, which is not a category in the usual sense—but it satisfies the same structural axioms, and Proposition 3.4 will prove that it makes* **Force** *into a category enriched over Q.*

A.16 Definition (Symmetric monoidal category)

A monoidal category *is a category equipped with a* tensor product *bifunctor* $\otimes : \mathcal{C} \times \mathcal{C} \to \mathcal{C}$, *a* unit object I, *and natural isomorphisms for associativity* $(A \otimes B) \otimes C \cong A \otimes (B \otimes C)$ *and unitality* $I \otimes A \cong A \cong A \otimes I$. *A* symmetric *monoidal category additionally has a natural isomorphism* $A \otimes B \cong B \otimes A$.

A.17 Example (Force as a symmetric monoidal category)

Chapter 2 will prove that $(\mathbf{Force}, \oplus, \nvdash_{\mathbb{M}_L})$ *is a symmetric monoidal category (Proposition 2.16), where* $\oplus$ *is the disjoint union of structured forces and* $\nvdash_{\mathbb{M}_L}$ *is the void force. Combining two allied forces under a coalition command is the tensor product* $F_1 \oplus F_2$*; the cost of the combined force is computed by the monoidal extension of the cost functor. The symmetric isomorphism* $F_1 \oplus F_2 \cong F_2 \oplus F_1$ *encodes the fact that the order in which two forces are placed side-by-side is immaterial to their joint structure. The extensions of Section 5.4 will use this monoidal structure to model alliance formation.*

A.18 Definition (Groupoid)

A groupoid *is a category in which every morphism has an inverse: for each* $f \in \mathcal{C}(A, B)$ *there exists* $f^{-1} \in \mathcal{C}(B, A)$ *with* $f^{-1} \circ f = \mathrm{id}_A$ *and* $f \circ f^{-1} = \mathrm{id}_B$.

A.19 Example (The rearrangement groupoid)

Chapter 2 will identify a full subcategory of **Force** *called the* rearrangement groupoid $\mathbf{Force}^{\times}$ *(Proposition 2.27): the subcategory on invertible morphisms, those transformations that can be undone at zero cost. Reshuffling a force from one formation to another and back again without spending any resources is exactly an isomorphism in* $\mathbf{Force}^{\times}$. *The orbits of* $\mathbf{Force}^{\times}$ *will partition the force space into equivalence classes of freely interconvertible configurations: forces that are "the same" up to costless reorganization. In the Lawvere metric (Section A.5), two forces in the same orbit will be at distance* $\perp$ *from each other in both directions, which is why the Lawvere metric is non-separating.*

Slice and coslice categories. The coslice category $F/\mathbf{Force}$ has objects the pairs (G, ϕ) with $\phi : F \to G$ a morphism in **Force**; it collects all forces reachable from F. The slice $\mathbf{Force}/F$ dualizes: all forces from which F is reachable. These are the Force-Maker's *forward cone* (what she can become) and *backward cone* (what could become her), and they will encode the strategic position in Chapter 2's discussion of reachability.

A.3 Pushouts and Double-Pushout Rewriting

The morphisms of **Force** will be built from double-pushout (DPO) rewriting steps, introduced by Ehrig, Pfender and Schneider (1973) and developed in Ehrig et al. (2006). Each step applies a rewrite rule to a matched subgraph of the current force: it deletes the matched portion (minus the preserved interface) and glues in the replacement. Both the deletion and the addition are instances of a single categorical construction called a *pushout*.

A.20 Definition (Pushout)

Given morphisms $f : C \to A$ and $g : C \to B$ in a category $\mathcal{C}$, a pushout *of the span (f, g) is an object P together with morphisms $i : A \to P$ and $j : B \to P$ such that $i \circ f = j \circ g$ (the square commutes), and P is* universal*: for any Q with $i' : A \to Q$ and $j' : B \to Q$ satisfying $i' \circ f = j' \circ g$, there exists a unique $u : P \to Q$ with $u \circ i = i'$ and $u \circ j = j'$.*

Informally, the pushout of $A \xleftarrow{f} C \xrightarrow{g} B$ is the object obtained by "gluing A and B together along their common sub-object C": wherever f and g identify a point of C with points of A and B, those points are merged in P; everything else remains distinct.

In **Set**, the pushout is the disjoint union $A \sqcup B$ modulo the equivalence relation generated by $f(c) \sim g(c)$ for all $c \in C$. In the category $\mathcal{G}$ of labeled graphs, the same construction applies node- and edge-wise: nodes identified by f and g are merged (provided they carry the same label), and edges are inherited.

A.21 Definition (Rewrite rule)

A rewrite rule *p in $\mathcal{G}$ is a span $p = (L \xleftarrow{\ell} K \xrightarrow{r} R)$ where L, K, R are labeled graphs and ℓ, r are injective homomorphisms. L is the* left-hand side *(the pattern to match), K is the* interface *(what is preserved), and R is the* right-hand side *(the replacement).*

A.22 Example (Rules in Force)

Chapter 2 will adopt exactly this definition of a rewrite rule (Definition 2.1). A rule that adds an artillery battery to an infantry brigade will have: L = the brigade without the battery (the pattern that must be present); K = L (the entire brigade is preserved as the interface, since nothing is deleted); and R = L plus the battery with its bonds to the existing units. A rule that replaces a throwing-spear phalanx with a stabbing-spear phalanx—Shaka's fundamental innovation—will have: L = the old molecule; K = the set of soldier nodes (preserved); and R = the new molecule with stabbing-spear bonds and the same soldier nodes re-labeled (or with bonds re-drawn). The interface K specifies what the rule cannot *change: the nodes present in both L and R, glued to the host graph in exactly the same way before and after.*

A.23 Definition (DPO rewrite step)

Given a rule $p = (L \xleftarrow{\ell} K \xrightarrow{r} R)$ *and an injective match* $m : L \hookrightarrow G$*, a* double-pushout rewrite step $G \Rightarrow_{p,m} G'$ *is defined by two pushout squares:*

$$\begin{array}{ccccc} L & \xleftarrow{\ell} & K & \xrightarrow{r} & R \\ {\scriptstyle m}\big\downarrow & & {\scriptstyle d}\big\downarrow & & \big\downarrow{\scriptstyle m'} \\ G & \xleftarrow{g} & D & \xrightarrow{g'} & G' \end{array}$$

The left square (the deletion step*) is a pushout that produces the* pushout complement *D: G with the nodes and edges of* $m(L \setminus \ell(K))$ *removed and the interface* $m(\ell(K))$ *intact. The right square (the* addition step*) is a pushout that produces G': D with R glued in along the interface, using the morphism* $d : K \to D$ *provided by the left pushout.*

A.24 Example (DPO in Force)

Chapter 2 will use exactly this DPO construction as the definition of a force transformation step (Definition 2.5). A morphism in $\mathbf{Force}_{\mathcal{R}}$ *will be a finite sequence of such steps using rules from the rule library* $\mathcal{R}$*. The Force-Maker perceives a tactical situation as a match* $m : L \hookrightarrow G$*, executes the deletion (removing what the rule transforms away), and executes the addition (installing the replacement). A result on parallel independence*

(Proposition 2.7) will show that two DPO steps whose matches do not interfere can be applied in either order with the same result, capturing the independence of simultaneous tactical transformations.

A.25 Proposition (Gluing condition)

The pushout complement D exists (uniquely up to isomorphism) if and only if the match $m : L \hookrightarrow G$ satisfies:

(i) Dangling condition*: no edge of G outside* $\mathrm{im}(m)$ *is incident to a node in* $m(L) \setminus m(\ell(K))$*; and*

(ii) Identification condition*: m is injective on $L \setminus \ell(K)$.*

A.26 Example (Validity in Force)

Chapter 2 will prove this result and apply it to the force setting (Proposition 2.6). The dangling condition will rule out rewrites that delete a unit that still has bonds to surviving units: removing a headquarters from a force while leaving its subordinate battalions bonded to the now-absent node would leave dangling edges. The identification condition rules out a match that collapses two distinct pattern nodes onto the same host node. Together, the two conditions are the structural validity check for a Force morphism: a transformation sequence is valid when every step satisfies the gluing condition, and this validity is a consequence of structure rather than an externally imposed constraint.

A.4 Complete Lattices and Quantales

The cost functor of Chapter 3 will assign costs to morphisms in a structure called a *quantale*: a complete lattice equipped with a composition operation. Quantales accommodate asymmetry (building up costs differently from tearing down), non-additivity (the cost of A then B need not equal the sum), and the distinction between zero cost and infinite cost.

A.27 Definition (Complete lattice)

A partial order *on X is a relation $\leq$ that is reflexive, antisymmetric, and transitive. A* lattice *is a partial order in which every pair x, y has a* join $x \vee y$ *(least upper bound) and a* meet $x \wedge y$ *(greatest lower bound). A* complete lattice *has joins $\bigvee S$ and meets $\bigwedge S$ for* every *subset S, not just pairs. The* bottom element *is $\bot = \bigvee \varnothing$ and the* top element *is $\top = \bigvee X$.*

In the cost setting, $\leq$ reads "is at most as costly as," $\bot$ is zero cost (the cost of doing nothing), and $\top$ is infinite cost (an impossible transformation).

A.28 Definition (Quantale)

A quantale *is a triple* $(Q, \leq, \otimes)$ *where* $(Q, \leq)$ *is a complete lattice and* $\otimes : Q \times Q \to Q$ *is an associative binary operation that:*

(i) distributes over arbitrary joins from both sides: $a \otimes (\bigvee S) = \bigvee_{s \in S}(a \otimes s)$ *and symmetrically; and*

(ii) has $\bot$ *as its two-sided unit:* $a \otimes \bot = a = \bot \otimes a$.

A.29 Example (The five quantales of Chapter 3)

Chapter 3 will state these axioms (Definition 3.1) and then present five quantale instances as the five distinct cost languages available to the Force-Maker.

The Boolean quantale $(\{\bot, \top\}, \leq, \vee)$ *is the information-minimal case:* $\otimes = \vee$ *(logical or), so costs are either possible* ($\bot$) *or impossible* ($\top$) *and composing two possible steps yields a possible step. This is the quantale in which the cost functor collapses to the convertibility preorder: a transformation either exists or it does not, and Chapter 3 will show that this Boolean specialization recovers the reachability structure of Chapter 2 as a degenerate case.*

The textbook quantale $(\mathbb{R}_{\geq 0} \cup \{\infty\}, \leq, +)$ *takes* $\otimes = +$*: costs are nonnegative real numbers that add along paths, and the Lawvere metric measures the cheapest path. This is the default quantale used throughout Chapter 3 and the one that generates the topology used in Chapter 4.*

The probability quantale $([0, 1], \geq, \times)$ *interprets costs as failure probabilities: lower probability of success corresponds to higher cost, and probabilities multiply along sequences of independent steps. Chapter 3 will use this quantale to model Cold War nuclear layered defenses, where the probability of penetrating* n *successive layers is the product of the individual penetration probabilities.*

The Force-Maker's choice of quantale is, as Chapter 3 will argue, an epistemic stance*: it determines what kind of object cost is and how costs compose along transformation sequences.*

A.5 Lawvere Metrics, Topology, and Second-Countability

The Lawvere metric of Chapter 3 will be a generalization of classical distance that permits asymmetry and non-separation. Both departures from the classical axioms are features of the force setting, not defects; and the topology the metric generates is the analytical environment in which Debreu's representation theorem operates.

A.30 Definition (Metric space)

A metric space (X, d) *has* $d : X \times X \to \mathbb{R}_{\geq 0}$ *satisfying for all* x, y, z*: (i)* $d(x, x) = 0$*; (ii)* $d(x, y) = 0 \Rightarrow x = y$*; (iii)* $d(x, y) = d(y, x)$*; and (iv)* $d(x, z) \leq d(x, y) + d(y, z)$*.*

The real line with $d(x, y) = |x - y|$ is the canonical example. The force space will violate (ii) and (iii).

A.31 Example (Why classical metrics fail)

Chapter 3 will document two structural facts about force transformation costs that require departing from the classical axioms. Asymmetry (Theorem 3.6): the demobilization of a wartime army to a peacetime garrison costs far less than the reverse mobilization—the opening of Chapter 3 will use the British demobilization of 1945 and the subsequent Korean mobilization as the paradigm case. Non-separation: two force configurations in the same orbit of the rearrangement groupoid $\mathbf{Force}^{\times}$ *are freely interconvertible at zero cost in both directions, so their distance is* $\bot = 0$ *even though they are distinct structured forces. The bridge proposition (Proposition 3.7) will make this precise: the convertibility preorder and the Lawvere metric are linked by* $F_1 \precsim_{\mathcal{R}} F_2 \Longleftrightarrow d(F_1, F_2) < \top$*, so non-separation corresponds exactly to mutual reachability at finite cost.*

A.32 Definition (Lawvere metric space)

A Lawvere metric space (X, d) *has* $d : X \times X \to Q$ *valued in a quantale* $(Q, \leq, \otimes)$*, satisfying only: (i)* $d(x, x) = \bot$*; and (ii)* $d(x, z) \leq d(x, y) \otimes d(y, z)$*. Symmetry and identity of indiscernibles are not required.*

A.33 Example (The Lawvere metric on the force space)

Chapter 3 will define the Lawvere metric on $\mathbb{F}^\star$ *(Definition 3.5) as*

$$d(F_1,F_2)=\bigwedge\nolimits_{\phi\in\mathbf{Force}_{\mathcal{R}}(F_1,F_2)} c(\phi),$$

with $d(F_1,F_2)=\top$ *when no morphism exists. Reflexivity holds because* $c(\mathrm{id}_F)=\bot$*. The triangle inequality holds because composing morphisms satisfies subadditivity, and the infimum of a composed cost is at most the quantale product of the individual infima. A maximality proposition (Proposition 3.8) will establish that this is the* unique maximal *function satisfying reflexivity, the triangle inequality, and compatibility with the cost functor: any "wishful" assignment of lower costs to transformation paths would violate one of these three conditions.*

A.34 Definition (Topological space)

A topological space (X,τ) *has a collection* $\tau\subseteq 2^X$ *of* open sets *closed under arbitrary unions and finite intersections, with* $\varnothing, X\in\tau$*. A* subbase *for* τ *is a collection* $\mathcal{S}$ *such that* τ *is the smallest topology containing* $\mathcal{S}$*: equivalently,* τ *consists of all unions of finite intersections of members of* $\mathcal{S}$*.*

A.35 Example (The Lawvere topology)

Chapter 3 will generate a topology τ_d *on* $\mathbb{F}^\star$ *from the subbase of forward open balls (Definition 3.9):*

$$B^+_\varepsilon(F)=\{G: d(F,G)<\varepsilon\}\qquad (F\in\mathbb{F}^\star,\ \varepsilon>\bot \text{ in } Q).$$

A neighborhood of F in τ_d *is a set containing all forces reachable from F below some cost threshold: proximity in* τ_d *means cheapness of transformation* from *F. The backward topology—generated by* $B^-_\varepsilon(F)=\{G: d(G,F)<\varepsilon\}$*—captures instead who can* reach *F; it is a distinct topology, not used in the representation theorem.*

A.36 Definition (Second-countable space)

A topological space (X,τ) *is* second-countable *if it has a* countable base*: a countable collection* $\{B_n\}_{n\in\mathbb{N}}$ *of open sets such that every open set in* τ *is a union of members of* $\{B_n\}$*.*

A.37 Example (Second-countability of the Lawvere topology)

Chapter 3 will prove that τ_d is second-countable under the textbook quantale (Proposition 3.10). The argument has two components: the force space $\mathbb{F}^\star(\mathbb{M}_L^\star)$ is countably infinite (established in Chapter 1); and $\mathbb{R}_{\geq 0}$ has $\mathbb{Q}_{>0}$ as a countable order-dense subset. The collection $\{B_q^+(F) : F \in \mathbb{F}^\star, q \in \mathbb{Q}_{>0}\}$ is therefore countable, and it is a base for τ_d because every ball of irrational radius ε contains a ball of rational radius $q < \varepsilon$ around the same center. This result is the bridge to Chapter 4: the second-countability of τ_d is precisely the topological hypothesis that Debreu's representation theorem requires, and the reader who reaches Chapter 4 will find that hypothesis already met.

A.38 Definition (Continuity of functions and relations)

A function $f : (X, \tau_X) \to (Y, \tau_Y)$ is continuous *if $f^{-1}(U) \in \tau_X$ for every $U \in \tau_Y$. A binary relation $\succsim$ on (X, τ) is* continuous *if for every x, the upper contour set $\{y : y \succsim x\}$ and the lower contour set $\{y : x \succsim y\}$ are both closed in τ.*

A.39 Example (Doctrinal continuity)

Chapter 4 will impose the doctrinal continuity condition in these exact terms (Proposition 4.2): a doctrine $\succsim$ is continuous with respect to τ_d if its upper and lower contour sets are closed. The strategic interpretation is that the doctrine must not assign dramatically different capability assessments to forces that are cheaply reachable from each other. If force G is reachable from F at low cost (adding one artillery brigade, say), a doctrine that ranks F as far superior to G while ranking some cheaply reachable G' as far inferior has a "ranking cliff" that no cost structure justifies; continuity rules this out.

Compactness. A topological space is *compact* if every open cover has a finite subcover. Under the local finiteness assumption (finitely many element types L and rule types in $\mathcal{R}$), every budget ball $\mathcal{D}_{\mathcal{S}}(\beta) = \{F : d(F, \mathcal{S}) \leq \beta\}$ is a finite set, because paths of bounded cost have bounded length and finitely many rules generate finitely many reachable configurations below any finite budget. A finite set is compact in any topology. Chapter 4 will use this compactness to establish properness of the force function (Proposition 4.4): the preimage of any compact set of capability levels is compact in τ_d, which is the condition game-theoretic arming models need for Nash equilibrium existence.

A.6 Preorders, Completeness, and Representation

The Force-Maker's doctrine is a preorder on the force space—a reflexive, transitive relation that need not rank every pair. The central mathematical question of Chapter 4 is: when does a preorder compress into a continuous real-valued force function?

A.40 Definition (Preorder)

A preorder *on X is a binary relation $\succsim$ that is reflexive ($x \succsim x$ for all x) and transitive (if $x \succsim y$ and $y \succsim z$ then $x \succsim z$). Write $x \succ y$ when $x \succsim y$ but not $y \succsim x$; write $x \sim y$ when both $x \succsim y$ and $y \succsim x$. The quotient of a preorder by the indifference relation $\sim$ is a partial order: the antisymmetric relation on equivalence classes $[x] = \{y : x \sim y\}$.*

A.41 Example (The two preorders on the force space)

Two distinct preorders on $\mathbb{F}^{\star}(\mathbb{M}_L^{\star})$ will appear in the book, and they should not be confused.

The convertibility preorder *$\precsim_{\mathcal{R}}$, introduced in Chapter 2 (Definition 2.21), records reachability: $F_1 \precsim_{\mathcal{R}} F_2$ when there exists a morphism from F_1 to F_2 in* $\mathbf{Force}_{\mathcal{R}}$. *Reflexivity holds because every force has an identity morphism; transitivity holds because morphisms compose. This preorder is a partial order of* strategic possibility, *not of* capability*: it says nothing about which force is better, only about which forces are reachable from which.*

The doctrine *$\succsim$, introduced in Chapter 4 (Definition 4.1), is the Force-Maker's capability ranking: $F_1 \succsim F_2$ when the Force-Maker judges F_1 to be at least as capable as F_2. It is a preorder on the same space but it records preference, not reachability. Two forces that are incomparable under $\succsim$ are not strategically equivalent; they are forces the doctrine has failed to rank.*

A.42 Definition (Completeness)

A preorder $\succsim$ on X is complete *if for every $x, y \in X$, either $x \succsim y$ or $y \succsim x$ (or both). A complete preorder is called a* total preorder.

A.43 Example (Completeness as a substantive doctrinal claim)

Chapter 4 will take up completeness as a claim about the Force-Maker's doctrine, not a mathematical convenience (Section 4.1). The Wehrmacht in 1940 held a complete ranking of Polish cavalry against German panzer divisions: every relevant dimension of combat power—firepower, mobility, protection—pointed unambiguously in the same direction. The US Army in 2001, by contrast, held a fully developed ranking over conventional combined-arms configurations but had no preorder over counterinsurgency-configured forces: its doctrine was incomplete across the mission-type boundary, and the force function its generals reached for in 2003 simply did not exist for the configurations the situation required. The Mearsheimer-Biddle debate, reconstructed in Chapter 4, will turn out to be a debate about whether the relevant doctrine is complete.

A.44 Theorem (Debreu, 1954)

Let $\succsim$ be a complete, continuous preorder on a second-countable topological space (X, τ). Then there exists a continuous function $m : X \to \mathbb{R}$ with $x \succsim y \Longleftrightarrow m(x) \geq m(y)$, and m is unique up to strictly increasing transformation.

A.45 Example (The Debreu theorem in Chapter 4)

Chapter 4 will apply this theorem to the force space under τ_d (Theorem 4.3). The theorem's three hypotheses will be met by construction: completeness is the Force-Maker's doctrinal assumption; continuity is Proposition 4.2; and second-countability is Proposition 3.10, proved in Chapter 3. Each of the force functions analyzed in Section 4.3—Lanchester's laws, Dupuy's QJM, the contest success function—will appear as a continuous function $m : \mathbb{F}^\star \to \mathbb{R}$ representing a specific complete doctrine $\succsim$ on a specific (often drastically reduced) domain. The uniqueness clause is the index number problem: the Lanchester score $q \cdot n^2$ and any strictly increasing transformation of it represent the same doctrine, and no scaling is "correct"—all are equally faithful to the preorder.

The proof of Debreu's theorem constructs m in three steps: second-countability provides a countable base; completeness and continuity allow one to assign each indifference class a rational number that respects the ordering; and the monotone closure of this assignment, extended to all of X by continuity, is the force function m. Full details are in Debreu (1954).

A.46 Theorem (Richter, 1966; Peleg, 1970)

Let $\succsim$ be a preorder (not necessarily complete) on a second-countable topological space (X, τ). Then there exists a countable family $\{m_n\}_{n\in\mathbb{N}}$ of continuous functions such that $x \succsim y \Longleftrightarrow m_n(x) \geq m_n(y)$ for all n. The family separates incomparable pairs: if $x \not\succsim y$ and $y \not\succsim x$, then some m_n ranks them in opposite orders.

A.47 Example (Richter-Peleg and the convertibility monotones)

Chapter 4 will prove the contrapositive of the Richter-Peleg theorem's necessity as a standalone result (Proposition 4.6): if $\succsim$ is incomplete—if some pair F_1, F_2 is incomparable—then no single continuous function can represent $\succsim$ on the full space. The proof is two lines: any real-valued m would have either $m(F_1) \geq m(F_2)$ or $m(F_2) \geq m(F_1)$, implying comparability; contradiction.

The Richter-Peleg theorem then gives what the single function cannot: a countable family representing the partial doctrine. The resource monotones introduced in Chapter 2 (Section 2.6)—functions $\mu : \mathbb{F}^{\star} \to \mathbb{R}$ satisfying $F_1 \precsim_{\mathcal{R}} F_2 \Rightarrow \mu(F_1) \leq \mu(F_2)$—are exactly the Richter-Peleg family for the convertibility preorder. They will be introduced in Chapter 2 as capability measures, but Chapter 4 will reveal their deeper identity: they are the unique appropriate representation of the partial doctrine $\precsim_{\mathcal{R}}$ on a second-countable space. Aggregating them into a single scalar requires a weighting rule, and every such weighting rule is itself a doctrinal choice—which is why the Office of Net Assessment's refusal to aggregate its mission-specific assessments was a principled stance, not a methodological failure.

Bibliography

Allison, Graham T. and Philip Zelikow. 1999. *Essence of Decision: Explaining the Cuban Missile Crisis*. 2nd ed. New York: Longman.

Aluffi, Paolo. 2009. *Algebra: Chapter 0*. Number v. 104 *in* "Graduate Studies in Mathematics" Providence, RI: American Mathematical Society.

Avant, Deborah D. 1994. *Political Institutions and Military Change: Lessons from Peripheral Wars*. Ithaca, NY: Cornell University Press.

Barnett, Michael and Raymond Duvall. 2005. "Power in International Politics." *International Organization* 59(1):39–75.

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

Biddle, Stephen D. 2004. *Military Power: Explaining Victory and Defeat in Modern Battle*. Princeton, NJ: Princeton University Press.

Boolos, George. 1984. "To Be Is to Be a Value of a Variable (or to Be Some Values of Some Variables)." *Journal of Philosophy* 81(8):430–449.

Brewer, John. 1990. *The Sinews of Power: War, Money, and the English State, 1688–1783*. 1st harvard university pbk. ed ed. Cambridge, MA: Harvard University Press.

Bricker, Phillip. 2016. Ontological Commitment. In *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta. Winter 2016 ed. Metaphysics Research Lab, Stanford University.

Brooks, Risa A. 2008. *Shaping Strategy: The Civil-Military Politics of Strategic Assessment*. Princeton, NJ: Princeton University Press.

Bury, Patrick. 2021. "Conceptualising the Quiet Revolution: The Post-Fordist Revolution in Western Military Logistics." *European Security* 30(1):112–136.

Cappella Zielinski, Rosella. 2016. *How States Pay for Wars*. Ithaca, NY: Cornell University Press.

Carnap, Rudolf. 1950. "Empiricism, Semantics, and Ontology." *Revue Internationale de Philosophie* 4(11):20–40.

Carroll, Robert J. 2026. "Policy Spaces as Organizational Schemes: A Semantic Theory of Electoral Competition." Working paper.

Chandler, David G. 1966. *The Campaigns of Napoleon*. New York: Macmillan.

Citino, Robert M. 2005. *The German Way of War: From the Thirty Years' War to the Third Reich*. Lawrence, KS: University Press of Kansas.

Coe, Andrew J. and Jane Vaynman. 2020. "Why Arms Control Is So Rare." *American Political Science Review* 114(2):342–355.

Coecke, Bob, Tobias Fritz and Robert W. Spekkens. 2016. "A Mathematical Theory of Resources." *Information and Computation* 250:59–86.

Cohen, Eliot A. 2002. *Supreme Command: Soldiers, Statesmen, and Leadership in Wartime*. New York: Free Press.

Congressional Budget Office. 2021. The U.S. Military's Force Structure: A Primer. Technical report Congressional Budget Office Washington, DC: .

Corradini, Andrea, Ugo Montanari, Francesca Rossi, Hartmut Ehrig, Reiko Heckel and Michael Löwe. 1997. Algebraic Approaches to Graph Transformation. In *Handbook of Graph Grammars and Computing by Graph Transformation*, ed. Grzegorz Rozenberg. Vol. 1 Singapore: World Scientific pp. 163–245.

Dahl, Robert A. 1957. "The Concept of Power." *Behavioral Science* 2(3):201–215.

Debreu, Gerard. 1954. Representation of a Preference Ordering by a Numerical Function. In *Decision Processes*, ed. Robert M. Thrall, Clyde H. Coombs and Robert L. Davis. New York: John Wiley & Sons pp. 159–165.

Dogauchi, Masato. 2008. Historical Development of Japanese Private International Law. In *Japanese and European Private International Law in Comparative Perspective*, ed. Jürgen Basedow, Harald Baum and Yuko Nishitani. Number 48 *in* "Materialien Zum Ausländischen Und Internationalen Privatrecht" Tübingen, Germany: Mohr Siebeck pp. 27–60.

Douhet, Giulio. 1921. *The Command of the Air*. Translated (1942) ed. Washington, DC: Air Force History and Museums Program.

Dupuy, Trevor N. 1979. *Numbers, Predictions, and War: Using History to Evaluate Combat Factors and Predict the Outcome of Battles*. Indianapolis: Bobbs-Merrill.

Ehrig, Hartmut, Karsten Ehrig, Ulrike Prange and Gabriele Taentzer. 2006. *Fundamentals of Algebraic Graph Transformation*. Berlin: Springer.

Ehrig, Hartmut, Michael Pfender and Hans Jürgen Schneider. 1973. "Graph-Grammars: An Algebraic Approach." *Proceedings of the 14th Annual Symposium on Switching and Automata Theory* pp. 167–180.

Epstein, Joshua M. 1990. *Conventional Force Reductions: A Dynamic Assessment*. Washington, DC: Brookings Institution.

Erbel, Mark and Christopher Kinsey. 2018. "Think Again – Supplying War: Reappraising Military Logistics and Its Centrality to Strategy and War." *Journal of Strategic Studies* 41(4):519–544.

Fearon, James D. 2018. "Cooperation, Conflict, and the Costs of Anarchy." *International Organization* 72(3):523–559.

Feaver, Peter D. 2003. *Armed Servants: Agency, Oversight, and Civil-Military Relations*. Cambridge, MA: Harvard University Press.

Federle, Jonathan, André Meier, Gernot J. Müller, Willi Mutschler and Moritz Schularick. 2026. "The Price of War." *American Economic Review* 116(3):791–827.

Fong, Brendan and David I. Spivak. 2019. *An Invitation to Applied Category Theory: Seven Sketches in Compositionality*. Cambridge, UK: Cambridge University Press.

Garfinkel, Michelle R. 1990. "Arming as a Strategic Investment in a Cooperative Equilibrium." *American Economic Review* 80(1):50–68.

Giere, Ronald N. 2004. "How Models Are Used to Represent Reality." *Philosophy of Science* 71(5):742–752.

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

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

Goldman, Emily O. and Leslie C. Eliason, eds. 2003. *The Diffusion of Military Technology and Ideas*. Stanford, CA: Stanford University Press.

Gonda, Tomáš and Robert W. Spekkens. 2023. "Monotones in General Resource Theories." *Compositionality* 5(7).

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

Hirshleifer, Jack. 1991. "The Paradox of Power." *Economics & Politics* 3(3):177–200.

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

Hofstadter, Douglas R. 1979. *Gödel, Escher, Bach: An Eternal Golden Braid*. New York: Basic Books.

Huntington, Samuel P. 1957. *The Soldier and the State: The Theory and Politics of Civil-Military Relations*. Renewed (1985) ed. Cambridge: Harvard University Press.

Hylton, Peter and Gary Kemp. 2023. Willard Van Orman Quine. In *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta and Uri Nodelman. Fall 2023 ed. Metaphysics Research Lab, Stanford University.

Jackson, Matthew O. and Massimo Morelli. 2009. "Strategic Militarization, Deterrence and Wars." *Quarterly Journal of Political Science* 4(4):279–313.

Janowitz, Morris. 1960. *The Professional Soldier: A Social and Political Portrait*. Glencoe, IL: Free Press.

Judson, Harry Pratt. 1888. *Cæsar's Army: A Study of the Military Art of the Romans in the Last Days of the Republic.* Number 13 *in* "Roman Life and Times" reprinted (1961) ed. New York: Biblo and Tannen.

Kagan, Donald and Gregory F. Viggiano, eds. 2013. *Men of Bronze: Hoplite Warfare in Ancient Greece*. Princeton, NJ: Princeton University Press.

Keegan, John. 1976. *The Face of Battle*. London: Jonathan Cape.

Kelly, G M. 1982. *Basic Concepts of Enriched Category Theory*. Number 10 *in* "Reprints in Theory and Applications of Categories" reprinted (2005) ed. London: Cambridge University Press.

Kennedy, Paul M. 1989. *The Rise and Fall of the Great Powers: Economic Change and Military Conflict from 1500 to 2000*. New York: Vintage Books.

King, Anthony. 2013. *The Combat Soldier: Infantry Tactics and Cohesion in the Twentieth and Twenty-First Centuries*. Oxford: Oxford University Press.

Knight, Ian. 1995. *The Anatomy of the Zulu Army: From Shaka to Cetshwayo, 1818–1879*. London: Greenhill Books.

Krepinevich, Andrew F. 1997. "Transforming Defense: National Security in the 21st Century." *Report of the National Defense Panel* . Report to the Secretary of Defense.

Krepinevich, Andrew F. and Barry D. Watts. 2015. *The Last Warrior: Andrew Marshall and the Shaping of Modern American Defense Strategy*. New York: Basic Books.

Kreps, David M. 1988. *Notes on the Theory of Choice*. Underground Classics in Economics New York: Routledge.

Laband, John. 1995. *The Rise and Fall of the Zulu Nation*. Paperback (1998) ed. London: Arms and Armour.

Lanchester, Frederick William. 1916. *Aircraft in Warfare: The Dawn of the Fourth Arm*. London: Constable.

Lawvere, F. William. 1973. "Metric Spaces, Generalized Logic, and Closed Categories." *Rendiconti del Seminario Matematico e Fisico di Milano* 43(1):135–166.

Lowe, E. J. 2006. *The Four-Category Ontology: A Metaphysical Foundation for Natural Science*. Oxford, UK: Oxford University Press.

Lukes, Steven. 2005. *Power: A Radical View*. 2nd ed. Basingstoke: Palgrave Macmillan.

Mac Lane, Saunders. 1986. *Mathematics: Form and Function*. New York: Springer-Verlag.

Machiavelli, Niccolò. 1521. *Art of War*. Lynch (2003) ed. Chicago: University of Chicago Press.

Mahan, A. T. (Alfred Thayer). 1890. *The Influence of Sea Power Upon History, 1660-1783*. Project gutenberg updated (2020) ed. Project Gutenberg.

Martin, Bernd. 1995. *Japan and Germany in the Modern World*. Providence, RI: Berghahn Books.

McCloskey, Deirdre N. 1985. *The Rhetoric of Economics*. Rhetoric of the Human Sciences second (1998) ed. Madison, WI: University of Wisconsin Press.

McNeill, William H. 1982. *The Pursuit of Power: Technology, Armed Force, and Society since A.D. 1000*. Chicago: University of Chicago Press.

Mearsheimer, John J. 2001. *The Tragedy of Great Power Politics*. First norton (2003) ed. New York: Norton.

Meirowitz, Adam and Anne E. Sartori. 2008. "Strategic Uncertainty as a Cause of War." *Quarterly Journal of Political Science* 3(4):327–352.

Millett, Allan R. and Williamson Murray, eds. 1988. *Military Effectiveness*. Boston: Allen & Unwin. 3 vols.

Morris, Donald R. 1965. *The Washing of the Spears: A History of the Rise of the Zulu Nation under Shaka and Its Fall in the Zulu War of 1879*. First paperback (1969) ed. New York: Simon & Schuster.

Murray, Williamson. 2011. *Military Adaptation in War: With Fear of Change*. Cambridge: Cambridge University Press.

Muth, John F. 1961. "Rational Expectations and the Theory of Price Movements." *Econometrica* 29(3):315–335.

Nunn, Frederick M. 1970. "Emil Körner and the Prussianization of the Chilean Army: Origins, Process, and Consequences, 1885–1920." *Hispanic American Historical Review* 50(2):300–322.

Nunn, Frederick M. 1983. *Yesterday's Soldiers: European Military Professionalism in South America, 1890–1940*. Lincoln, NE: University of Nebraska Press.

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

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

Peleg, Bezalel. 1970. "Utility Functions for Partially Ordered Topological Spaces." *Econometrica* 38(1):93–96.

Pieper, Josef. 1948. *Leisure, the Basis of Culture*. Malsbary (1998) ed. South Bend, IN: St. Augustine's Press.

Pollack, Kenneth M. 2019. *Armies of Sand: The Past, Present, and Future of Arab Military Effectiveness*. New York: Oxford University Press.

Posen, Barry R. 1984. *The Sources of Military Doctrine: France, Britain, and Germany between the World Wars*. Cornell Studies in Security Affairs first cornell paperbacks (1986) ed. Ithaca, NY: Cornell University Press.

Powell, Robert. 1993. "Guns, Butter, and Anarchy." *American Political Science Review* 87(1):115–132.

Powell, Robert. 2017. "Research Bets and Behavioral IR." *International Organization* 71(S1):S265–S277.

Quine, Willard V. 1948. "On What There Is." *The Review of Metaphysics* 2(5):21–38.

Reiter, Dan and Allan C. Stam. 2002. *Democracies at War*. Princeton, NJ: Princeton University Press.

Resende-Santos, João. 2007. *Neorealism, States, and the Modern Mass Army*. Cambridge: Cambridge University Press.

Richter, Marcel K. 1966. "Revealed Preference Theory." *Econometrica* 34(3):635–645.

Rosen, Stephen Peter. 1991. *Winning the Next War: Innovation and the Modern Military*. Ithaca, NY: Cornell University Press.

Sater, William F. and Holger H. Herwig. 1999. *The Grand Illusion: The Prussianization of the Chilean Army*. Studies in War, Society, and the Military Lincoln, NE: University of Nebraska Press.

Savage, Leonard J. 1954. *The Foundations of Statistics*. New York: John Wiley & Sons.

Schelling, Thomas C. 1966. *Arms and Influence*. 2008 ed. New Haven, CT: Yale University Press.

Shils, Edward A. and Morris Janowitz. 1948. "Cohesion and Disintegration in the Wehrmacht in World War II." *Public Opinion Quarterly* 12(2):280–315.

Singer, J. David, Stuart Bremer and John Stuckey. 1972. Capability Distribution, Uncertainty, and Major Power War, 1820–1965. In *Peace, War, and Numbers*, ed. Bruce Russett. Beverly Hills, CA: Sage.

Skaperdas, Stergios. 1992. "Cooperation, Conflict, and Power in the Absence of Property Rights." *American Economic Review* 82(4):720–739.

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

Snyder, Glenn H. 1997. *Alliance Politics*. Ithaca, NY: Cornell University Press.

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

Strange, Susan. 1988. *States and Markets*. Reprinted (1993) ed. London: Pinter.

Talmadge, Caitlin. 2015. *The Dictator's Army: Battlefield Effectiveness in Authoritarian Regimes*. Cornell Studies in Security Affairs Ithaca, NY: Cornell University Press.

Thucydides. 432 B.C.E. *The History of the Peloponnesian War*. Project gutenberg (2003) ed. Salt Lake City, UT: Project Gutenberg.

Tilly, Charles. 1985. War Making and State Making as Organized Crime. In *Bringing the State Back In*, ed. Peter B. Evans, Dietrich Rueschemeyer and Theda Skocpol. Cambridge: Cambridge University Press pp. 169–191.

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

van Creveld, Martin. 1985. *Command in War*. Cambridge, MA: Harvard University Press.

Van Creveld, Martin. 2004. *Supplying War: Logistics from Wallenstein to Patton*. Second ed. Cambridge, UK: Cambridge University Press.

Viggiano, Gregory F. and Hans Van Wees. 2013. The Arms, Armor, and Iconography of Early Greek Hoplite Warfare. In *Men of Bronze: Hoplite Warfare in Ancient Greece*, ed. Gregory F. Viggiano and Donald Kagan. Princeton University Press pp. 57–73.

von Clausewitz, Carl. 1832. *On War*. Indexed (1976) howard/paret ed. Princeton, NJ: Princeton University Press.

Waltz, Kenneth N. 1979. *Theory of International Politics*. Boston, MA: McGraw-Hill.

Weber, Max. 1978. *Economy and Society: An Outline of Interpretive Sociology*. Berkeley: University of California Press. Edited by Guenther Roth and Claus Wittich. Originally published 1922.

Weitsman, Patricia A. 2014. *Waging War: Alliances, Coalitions, and Institutions of Interstate Violence*. Stanford, CA: Stanford University Press.