

A Theory of Force^{*}

Robert J. Carroll[†]

March 31, 2025

Abstract

I'm so sorry, reader. So very, very sorry.

Section	Words	Pages
The opener	2,738	6
The apology	2,276	5
The atomic theory	3,592	6
The resource theory	19,872	55
The organization theory	13,813	41
The concept theory	13,369	33
The closer	1,179	2.5
The appendix	3,850	19
\sum	60,689	167.5

^{*}Version 1.0 (November 2024).

[†]Assistant Professor, Department of Political Science, University of Illinois at Urbana-Champaign. Email: rjc@illinois.edu

Contents

0	What Force is Not (or, An Apology to Power)	7
1	Force is Atomic	12
2	Force is a Resource	19
2.1	<i>Mise en Place</i>	20
2.2	The Conversion of Force	27
2.3	The Costs of Conversion	39
2.4	The First Rationality	52
2.5	The Space Between Configurations	58
2.6	Convergence	70
3	Force is Organized	81
3.1	The Structure of Force	83
3.2	Structured Forces	88
3.3	Organization is Organized	91
3.4	Restructuring Forces	101
3.5	Lifting the Metric	114
3.6	The First Rationality, Part Deux	119
4	Force is a Concept	125
4.1	Emulation	127
4.2	The Second Rationality	142
4.3	Strategic Calibration	154
5	What is Force?	162
A	Proofs	165

List of Figures

1	A hoplite molecule.	16
2	Two hoplites in a force configuration.	23
3	Sample categorical diagram of force configurations and conversion morphisms.	32
4	Fashioning a phalanx from two hoplites.	33
5	Training a soldier into a hoplite.	34
6	Two paths to training a soldier	50
7	The Slingshot Effect	64
8	A hoplite molecule, revisited.	71
9	Same presentation of a force structure.	84
10	A Roman legion.	84
11	A binary support relationship.	95
12	Solving the emulation problem.	137

List of Tables

1	Lawvere's (1973) and the Boolean structures.	46
2	A set of processes and their costs.	55
3	Allowable selections for the choice schedule.	55
4	Force structure for the United States Army Armored Brigade Combat Team.	92
5	Brigade-level units in the United States Armed Forces.	93

The whole machinery of our intelligence, our general ideas and laws, fixed and external objects, principles, persons, and gods, are so many symbolic, algebraic expressions. They stand for experience; experience which we are incapable of retaining and surveying in its multitudinous immediacy. We should flounder hopelessly, like the animals, did we not keep ourselves afloat and direct our course by these intellectual devices. Theory helps us to bear our ignorance of fact.

George Santayana, *The Sense of Beauty* (1896)

Principles and opinions can seldom reduce the path of reason to a simple line. As in all practical matters, a certain latitude always remains. Beauty cannot be defined by abscissas and ordinates; neither are circles and ellipses created by their algebraic formulas. The man of action must at times trust in the sensitive instinct of judgment, derived from his native intelligence and developed through reflection, which almost unconsciously hits the right course.

Carl von Clausewitz, *On War* (1832)

Man has always had recourse to violence; sometimes this recourse was a mere crime, and does not interest us here. But at other times violence was the means resorted to by him who had previously exhausted all others in defence of the rights of justice which he thought he possessed. It may be regrettable that human nature tends on occasion to this form of violence, but it is undeniable that it implies the greatest tribute to reason and justice. For this form of violence is none other than reason exasperated. Force was, in fact, the ultima ratio. Rather stupidly it has been the custom to take ironically this expression, which clearly indicates the previous submission of force to methods of reason.

José Ortega y Gasset, *The Revolt of the Masses* (1930)

What is force?

It is common for authors to take great rhetorical pains to persuade the reader that their research question is important: important enough to justify the effort of addressing it, or—at the very least—important enough to justify skimming the first few bits of the answer. No such pains will be taken here. This is not because your humble author thinks his own research question unimportant—far from it—but rather because the question is so obviously important that it requires no further justification. This is evident from the fact that the word “force” occupies a central place in the vocabulary of international relations. States possess “armed forces” and engage in “uses of force” or the occasional “show 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 it is that the word “force” points to, it is clearly a central concept in the study of international relations, salt in the waters we swim. There is no better reason to ask what it 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 of force itself, but also because of the form of the question. Questions of the form “What is X ?” are difficult because they are *ontological* questions. They ask not for a definition of X in terms of other things—or more to the point, a meaning to assign the word “ X ” using other words—but for the *nature* of X itself, its essential features, its defining characteristics, its place in the web of knowledge. To be sure, definitions—and speaking more semantically, coding rules—offer clarity and useful intuitions. But they cannot grasp the full essence of “What is X ?” questions, at least not for any X es worth deeming “important.”

If the form of the question is at least partially to blame for its difficulty, then we would do well to consider the form of the answer. This sets our task: to provide an answer to the question “What is force?” that is both clear and useful, that is both precise and comprehensive, that is both rigorous and accessible. We take as inspiration the famous slogan attributed to philosopher Willard Van Orman Quine: “To be is to be the value of some variable (or the values of some variables).”¹ In other words, to say that something exists is to say that it is one among the values that some variable can take. We care less about one force or another—say, the United States’s Armed Forces in 2024—and more about the general concept of force, the variable that takes these values, the class of things generated by the predicate “— is a force.” To answer the question is to accommodate each of a variable’s values: values past, present and future; values actual, possible, and impossible; values real, ideal, and imaginary. Moreover, it is to accommodate the nature of the variable itself: its structure, its properties, its relations to other variables, its place in the grand scheme of things. Definitions alone are woefully inadequate for this task, as they cannot trace out the contours of the class of things to which a concept refers. They provide neither the values of the variable nor its nature, treating only the symptoms of our ignorance.

And yet, if you and I are to work on this question as a we, we must have some common ground, even if only on an intuitive level. The theory to follow comes equipped with a mathematical model, and such models are useful only after we imbue them with meaning. Indeed, the model is similar to models from other disciplines—physics, logic, computer science—so it is not poised to provide its own meaning like more tailored models often do. In the absence of such guidance, we need an orienting principle; in the absence of such a principle, our interpretations of the same model might diverge too wildly for comfort.

¹The quotation as given is but one of many formulations of what’s now called *Quine’s criterion of ontological commitment*, this particular one due to George Boolos (1984). See relevant articles in the *Stanford Encyclopedia of Philosophy* (e.g., Bricker, 2016; Hylton and Kemp, 2023).

As such, we will not use the word “force” loosely, which is difficult because it admits many meanings. The most relevant senses in which the most recent edition of the *Oxford English Dictionary* defines it are:

force, n. [fɔrs] *I.* Strength, power.

3. Power or might (of a ruler, realm, or the like); *esp.* military strength or power.
4. Concrete *a.* A body of armed men, an army. In plural the troops or soldiers composing the fighting strength of a kingdom or of a commander in the field; also in attributive use or in the possessive, *esp.* during the war of 1939–1945.

Digging in, the most relevant senses of “strength” and “might” are:

strength, n. [streɪŋ(k)θ] *II.* A person or thing that has, gives, or shows the quality of being strong.

10. Military.

- a.* Collectively: troops, armed forces; the ships of a navy; personnel, equipment, or resources for waging war or defending against attack.
- b.* A body of soldiers; an armed force. Cf. sense II.15a. Now literary and somewhat rare.
- c.* The number of people on the muster roll of an army, a regiment, a navy, (now) an air force, etc.; the number of ships in a navy or fleet, or of aircraft in an air force. Also: the personnel making up such a number. Cf. sense II.15a.

15.*a.* The number of people, animals, or things forming a set or unit; a number used or acting together; a complement.

might, n. [maɪt] 3. Great strength, imposing power. Now somewhat rhetorical. *b.* As an attribute of a person or other living creature, or of nation or other group or collection of people: physical or mental strength or powers, commanding influence, military resources, extent of power.

These definitions are striking for their tangibility: ships and aircraft, soldiers and animals, war and fighting, and so on. My impression is that the most ordinary usage of the word “force” relevant for our purposes is indeed material in nature. We are discussing the resources relevant for the production of the means of war, and so proceed as materialists.

We wish to theorize the thing to which these selected definitions point, and with the exception of Section 0 and the conclusion, the remainder of this manuscript is devoted to that task.² This involves both the introduction of theoretical concepts and mathematical machinery. The structure of the manuscript reflects the structure of the theory and its machine, so in previewing the former we preview the latter. We will proceed in four parts.

- § In Section 1, we observe the essential fact that force is *atomic*: comprised of individual units that cannot be further divided. Such units are the basic building blocks of force, the smallest things that can be called forces. These can be as simple as a naked human body or as complex as a modern aircraft carrier. This suggests that atoms arrive in various sorts called *elements*, and the elements of force are indeed the model’s first primitive. From these elements, we may build up more complicated structures called *molecules*, which represent the combination of elements into a single unit. A soldier is one thing, a soldier equipped with and trained in the use of a rifle another, and a soldier trained in the use of a rifle and a tank a third thing still. These molecules are the model’s smallest functional units.
- § In Section 2, we observe that force is a *resource*: a thing that can be used to produce other things. Molecules may be gathered into a single entity called a *configuration*, which collects molecules without forcing them to interact. Just as hydrogen and oxygen can be collected without forming water, so too can a soldier and a tank be collected without forming an integrated fighting force. This process might be concatenated many times, allowing us to consider three hundred hoplites without necessarily fashioning them into a phalanx. This raises questions about how one would go about fashioning a phalanx out of three hundred hoplites, and we answer this by providing a theory of *force conversion*. To make hydrogen and oxygen into water, some process must be applied or some reagent introduced; to turn a recruit and a pile of equipment into a soldier, some training must be conducted or some doctrine introduced. By disentangling the configuration and conversion operations on force, we can begin to understand the relationships between forces. We take a political-economic approach to this, thinking of various conversion processes as costly in the context of cost-calibration scheme. Remarkably, this simple logic provides equips our model with a rich structure—in particular, a metric structure that allows us to reason about inter-force relationships by thinking in spatial terms.

²Section 0 acknowledges the limitations of the concept of force with respect to the (perhaps) more fundamental concept of power. The conclusion concludes—!!!—the manuscript.

- § In Section 3, we observe that force is *organized*: it is not a mere collection of resources, but a collection of resources that are arranged in a particular way. In modern times it is difficult to conceive of a fighting force that is not organized in one way or another, and it's easy to become dazzled with the bureaucratic complexity of modern militaries. Few things represent this complexity better than an *organizational chart*, which is a diagram that shows how the resources of a force are arranged. Thus, this section includes a theory of org charts. We begin by introducing the concept of a *force structure*, which is a way of organizing the resources of a force. Then comes the question of how to assign each component of a force structure with its own resources, and we answer this by providing a theory of *force allocation*. Remarkably, the relationship between a structure and its resources is analogous to the relationship between a force molecules and its elements, so the same mathematical machinery that allows us to reason about the latter allows us to reason about the former. Finally, we introduce a theory of *restructuring*, which allows us to turn one structured force into another just the same way that we turn a recruit and a panoply into a hoplite. This, too, allows us to transport intuitions from the configurational level to the *organizational* level, so that the class of all things one could call a structured force is endowed with a metric structure. This completes the ontological aspect of the program, as we arrive at the variable that takes the values of all things that could be called a force.
- § In Section 4, we observe that force is a *concept*: it is not just a thing in the world, but also a way of thinking, a tool of the mind designed for navigating the world. It is a word we use when issuing comparisons about the things we call forces: one might be “more,” “less,” or “equally” forceful as another. If the pursuit of power includes the pursuit of force, then purpose-driven force-pursuers require some way of comparing forces. Otherwise, they have no way of knowing whether their path is better or worse than another. This raises questions about how decision-makers navigate the vast space of possible forces, and we answer this by providing two theories. The first, a warm-up, is a theory of force *emulation*, wherein an emulator begins from their *status quo* and moves towards a target force by restructuring their own force. The second, the main event, is a theory of force *comparison*, wherein a comparor issues comparisons between pairs of forces *a la* a preference relation. The structure developed in the earlier sections makes it easy to identify the conditions under which such comparisons reduce the vast space of possible forces to a single, numerically-valued dimension, which we call a *force scale*.

This roadmap suggests an arduous journey, and frankly, this is what it is. To make things more palatable, I should like to call your attention to some recurring themes to look out for. Though these themes are not necessarily part of the argument proper, they certainly play an important part in the narrative.

- ◊ **The discrete nature of force.** By the manuscript's end, the reader will likely be tired of hearing that a given class of objects can be stored in a countable set. This is a consequence of the model's atomic nature, itself a consequence of a materialist orientation. The calculus so common to applied researchers of optimal decisions is replaced by a discrete mathematics more amenable to the study of the relationships between forces. This is not to say that the model is not useful for decision-making—just useful in a different way.
- ◊ **The fractal nature of force.** As hinted at in the preview, the model is structured in such a way that the same mathematical machinery can be applied at different levels of analysis. Atoms are part of molecules, which are part of configurations, which are part of structures, and so on. This allows us to reason about the relationships between forces at different levels of abstraction, and to transport intuitions from one level to another.
- ◊ **Motion.** Though the model is not dynamic, its underlying logic is. We conceive of the distance between two forces as the cost it would take to get from the first to the second. Put differently, we reason about relationships between forces in terms of motion, making the space of possible forces a vast landscape. This metaphor suggests that it is the relationships between forces that are most important, not the forces themselves.
- ◊ **Generativity.** Molecules are more than collections of atoms, configurations are more than collections of molecules, and structures are more than collections of configurations; the whole is more than the sum of its parts. The tools we use to reason about combinations must be as minimalist as possible, else we might draw unwarranted conclusions reflecting not just our ignorance but also our lack of analytic control.
- ◊ **Rationality and subjectivity.** The model identifies some key parts of the force-maker's decision-making process, and these moments are marked by a tension between rationality and subjectivity. We seek not to identify a single rational decision-making process, but to understand what would have to be true about the force-maker's reasoning for her to be modeled by the model. We work hard to impose as little structure as possible on this reasoning, and this introduces subtle, but important, wrinkles.

Hopefully, these themes will provide guidance as you navigate the manuscript.

0 What Force is Not (or, An Apology to Power)

But before we proceed, I must acknowledge the omission of the word “power” from the above definitions, not to mention its omission from the remainder of the manuscript after this brief apology. The word is among the most terrifying in the whole of the liberal arts, much less I.R. Of course, this is because it is so sacred to so many theories spanning so many disciplines across so many centuries. What is more, we are all aware of the wide variety of circumstances in which we detect power, from meeting a significant other’s parents to threatening nuclear war. For all its charms, the *O.E.D.* is an insultingly-feeble reference for a word that has its own encyclopedia ([Dowding, 2011](#)).

We in I.R. might cherish the word more than most, as we can trace its usage back to some of our earliest texts. Thucydides, in his *History of the Peloponnesian War*, concludes that “The growth of the power of Athens, and the alarm which this inspired in Lacedaemon, made war inevitable” ([432 B.C.E.](#), Chapter I). But even here, what is meant by power is not entirely clear.³ Certainly, Thucydides means material power in some sense; for example, Themistocles’ build-up of the Athenian navy—twenty triremes a year over ten years—is a clear example of the growth of Athenian power in the material sense. And to be sure, the expansion of the Delian League had an impact on the material power of Athens and its allies, though this is far less important than its political and economic effects. It has been argued that Athens did not actually gain much material power in the run-up to war—see the later chapters of [Kagan \(1969\)](#) for a detailed discussion—but it is hard to explain away 6,000 talents in the treasury.

The problem just described—the increase in power of one state leading to the alarm of another—has been coined *Thucydides’ Trap* by Graham Allison ([2017](#)). All of his examples include shifting power among states, but the nature of that power is nearly always something that goes well beyond mere material power: empire building, alliances, trade, and so on. Though material shifts accompany these changes, they are not the only changes; ordinary users often mean more than material power when employing the word.

³The relevant passage in the original Greek is:

[...] τοὺς Ἀθηναίους ἡγοῦμαι μεγάλους γιγνομένους [...].

The word ἡγοῦμαι is a form of the verb ἡγέομαι; the most relevant sense in which Lidell-Scott-Jones defines it is “to lead, command in war.” μεγάλους is an adjective meaning “great” or “large,” and γιγνομένους is a form of the verb γίγνομαι, which means “to come into being.” The phrase translates more literally as “the Athenians becoming great in their ability to lead.” Strikingly, the more standard words for power—δύναμις, ισχύς, or ἕρωμαι, which point to power in many ways, including military power—are not used here. The word “power” is in the very first English translation of the *History*, due to Thomas Hobbes in 1628, and appears in all others I have seen.

One exceptional user of the word “power” is Kenneth N. Waltz, whose *Theory of International Politics* (1979) turns on the concept. Waltz alludes to Secretary of State Henry Kissinger’s address at the third *Pacem in Terris* convocation (1973). There Kissinger acknowledges multiple forms of power:

The most striking feature of the contemporary period, the feature that gives complexity as well as hope, is the radical transformation in the nature of power. Throughout history power has generally been homogeneous. Military, economic, and political potential were closely related. To be powerful, a nation had to be strong in all categories. Today the vocabulary of strength is more complex.

But Waltz accepts the Kissingerian view only in part: he accepts that there are many forms, but he insists that they are not so neatly separable:

States, because they are in a self-help system, have to use their combined capabilities in order to serve their interests. The economic, military, and other capabilities of nations cannot be sectorized and separately weighed. States are not placed in the top rank because they excel in one way or another. Their rank depends on how they score on *all* of the following items: size of population and territory, resource endowment, economic capability, military strength, political stability, and competence (p. 131, emphasis original).

Later on in the *Theory*, Waltz links the concepts of power and force more explicitly. Clearly, there are some lines Waltz refuses to cross:

[A] confusion about power is found in its odd definition. We are misled by the pragmatically formed and technologically influenced American definition of power—a definition that equates power with control. Power is then measured by the ability to get people to do what one wants them to do when otherwise they would not do it (cf. Dahl, 1957). That definition may serve for some purposes, but it ill fits the requirements of politics. To define “power” as a “cause” confuses process with outcome. To identify power with control is to assert that only power is needed in order to get one’s way. That is obviously false, else what would there be for political and military strategists to do? To use power is to apply one’s capabilities in an attempt to change someone else’s behavior in certain ways (p. 191).

Thus, power has a material basis: the capabilities of states, where presumably these cover the capabilities he listed earlier. Save for political stability and competence, these are material in nature.

Fellow Realist John J. Mearsheimer's *Tragedy of Great Power Politics* (2001) is less subtle in its materialist approach. Mearsheimer asserts: "Power is based on the particular material capabilities that a state possesses. The balance of power, therefore, is a function of tangible assets—such as armored divisions and nuclear weapons—that each great power controls" (p. 55).⁴ Complementing assets already in possession—*i.e.*, military power itself—is latent power, which depends on related, but separate, assets. Among these are wealth and population, which provide the wealth and personnel required to build military forces. To these Mearsheimer attaches technology, yet another necessary ingredient for mustering force. But in the final analysis, "a state's effective power is ultimately a function of its military forces and how they compare with the military forces of rival states" (p. 55). Power *is* force, and force *is* material. Kissinger's notions of distinct, not-necessarily-mutually-reinforcing forms of power take no haven here.

Mearsheimer identifies a gap between latent power and military power: not all states with the capacity to build massive forces—and here he means large, wealthy states—actually do so. Why not? He offers three reasons (pp. 75–82):

1. *Diminishing returns to military power*: "Spending more makes little sense when a state's defense effort is subject to diminishing returns (that is, if its capabilities are already on the 'flat of the curve') or if opponents can easily match the effort and maintain the balance of power."
2. *Heterogeneous efficiency*: "It is also unwise to liken the distribution of economic might with the distribution of military might because states convert their wealth into military power with varying degrees of efficiency."
3. *Multiple sorts of force*: "States can buy different kinds of military power, and how they build their armed forces has consequences for the balance of power. [...] The key issue here is whether a state has a large army with significant power-projection capability. But not all states spend the same percentage of their defense dollars on their army, and not all armies have the same power-projection capabilities."

Mearsheimer the theorist of power becomes Mearsheimer the political economist: states otherwise animated as power-hungry machines are subject to constraints about technology and the relative worth of additional defense spending in the face of diminishing returns—which is to say, the opportunity cost of force.

⁴This is the very thing I will be calling "force" in the sequel, right down to the fact that it is a function of tangible assets. In a sense, then, the model to follow is a model of Mearsheimer's power, which is itself something of a limiting case of other approaches to power. I will retain my humble agnosticism on the matter out of respect for, and agreement with, more nuanced views. However, it is worth examining the Mearsheimer view, as it foreshadows the model to follow.

The most developed theory of power in the hybrid sense Kissinger foresees is due to Susan Strange, who in her *States and Markets* (1988, quotations from pp. 29–30) identifies four structures of power in international political economy:

1. *The Security Structure*: “So long as the possibility of violent conflict threatens personal security, he who offers others protection against that threat is able to exercise power in other non-security matters like the distribution of food or the administration of justice.”
2. *The Production Structure*: “Who decides what shall be produced, by whom, by what means and with what combination of land, labour, capital and technology and how each shall be rewarded is as fundamental a question in political economy as who decides the means of defence against insecurity.”
3. *The Finance Structure*: “Finance—the control of credit—is the facet which has perhaps risen in importance in the last quarter century more rapidly than any other and has come to be of decisive importance in international economic relations and in the competition of corporate enterprises. [...] Its power to determine outcomes—in security, in production and in research—is enormous.”
4. *The Knowledge Structure*: “Knowledge is power and whoever is able to develop or acquire or to deny the access of others to a kind of knowledge respected and sought by others; and whoever can control the channels by which it is communicated to those given access to it, will exercise a very special kind of structural power.”

Strange is clear that these structures are not separate; they are interdependent and mutually reinforcing. Indeed, she takes Realists Past (e.g., Waltz) and Realists Future (e.g., Mearsheimer) to task for hyperfixation on the security structure:

The realist school of thought in international relations has held that in the last resort military power and the ability to use coercive force to compel the compliance of others must always prevail. *In the last resort*, this is undeniably true. But in the real world, not every relationship is put under such pressure. Not every decision is pushed to such extremes. There are many times and places where decisions are taken in which coercive force, though it plays some part in the choices made, does not play the whole, and is not the only significant source of power (pp. 31–32).

Again, Strange grants the material basis of power, but she insists that power is not merely material. Further, it is not hard to see how productive, financial, and epistemic power can influence security power.

Formal models in the I.R. literature, particularly those that carry on the Realist tradition, seem to take a very materialist approach to power. To take two prominent examples, James D. Fearon (1995) and Robert Powell (2006) both write down dynamic models of interstate bargaining where power is measured simply by the ability to win wars; should two states go to war, one of them expects to win with probability p , the other with probability $1 - p$. To become stronger is to increase one's probability of winning, and to become weaker is to decrease it. To be sure, this probability is influenced by things that go beyond material power: morale, tactics, the terms of the battlefield, first- or second-strike capabilities, and so on. But tellingly, Powell embraces the idea that power is a tool of destruction: “[...] the use of power is inefficient in that it destroys some of the flow” of the dynamic benefits states experience without war (p. 181). One could try to tell stories about how these benefits are harmed by non-material factors—perhaps His Majesty's tea tastes more bitter with his kingdom at war—but these seem overwrought. The most natural interpretation is that the benefits are material in nature, and that the destruction of the flow is a destruction of material benefits. It is worth noting that in both of these papers, and many others of their œuvre, the word “power” is used much more often than the word “force”; in Powell (2006), for example, “power” is used over one hundred times, “force” roughly ten.⁵

But since most ordinary users of the word “power” mean more than material power, the model given here should not be interpreted in its light. This is an important scope limitation, but it is necessary to do justice to the literature and to the word itself. That said, most ordinary users of the word “power” grant it some material basis—with varying degrees of centrality—so we remain interested in a concept of extraordinary importance. Personally, my conception of power is closest to Strange's, and the endogenous variable studied here falls short of the structures she identifies. Instead, the model to follow is a model of Strange's security structure of power, where the state organizes its resources to produce force. In keeping with Strange's view of the ongoing interplay of the structures, the model links this process to aspects related to the production, finance, and knowledge structures of power. Here ends the apology on power. Let us begin.

⁵The relative-material-capabilities approach to power is common in the empirical literature on conflict. A recent survey (Carroll and Kenkel, 2019) identified 94 articles in top political science journals over a ten year span that included predictors (or controls) based on ratios of the Composite Index of National Capability (CINC) score (Singer, Bremer and Stuckey, 1972). Many of these articles use variables of the form

$$\Pr(\text{State 1 defeats State 2 in time } t) = \frac{\text{CINC}_{1t}}{\text{CINC}_{1t} + \text{CINC}_{2t}} \in [0, 1],$$

where CINC_{it} is the CINC score of State i at time t .

1 Force is Atomic

Closing one's eyes and imagining a state's force, one does not see a blob; instead, one sees a collection of tanks, ships, planes, and soldiers, each a little piece of force in its own right. Force is atomic in the sense that it is made up of many little pieces of force, which (for lack of a better term) we call *atoms*. For all its nauseating science envy, the metaphor purchases us a great deal of intuition.

1. *Elementariness*. Atoms 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. A good theory of force should respect these differences.
2. *Familiality*. These elements may be organized into families. Helium and neon are both noble gases and thus share certain properties; likewise, the Dreadnought and Iowa battleship classes are both battleships and thus share certain properties. Being constructed by different states in different eras, they could quite plausibly be different force elements; nevertheless, their battleshippiness is a commonality. A good theory of force should respect these commonalities.
3. *Isotopy*. Within a force element there may be subtle differences. Carbon has fourteen known isotopes, each with a different number of neutrons and thus different physical properties; however, any of them behave as carbon does in most chemical reactions. Similarly, a tank might have a different gun, a different engine, or a different crew, but it is still a tank, and it still behaves as a tank does in most military operations. A good theory of force should respect these similarities.
4. *Configurability*. Atoms can be arranged in different ways. A hydrogen atom might arrive in a single atom, or in a deuterium molecule with another hydrogen atom, or in a water molecule with one hydrogen and one oxygen atom; a tank might arrive alone, or with air support, or as part of a larger armored division. The properties of the atom are the same, but the properties of the molecule are different. A good theory of force should respect these differences.

It is hoped that the theory developed here will indeed pay the appropriate respects to these features of force and that the reader will find it useful in understanding the structure of force. There are sure to be moments where the metaphor bends too far for comfort—or breaks entirely—but the reader is encouraged to keep in mind that both the theory and the metaphor are tools, and that the goal is to use them to understand the structure of force better.

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 will typically write this 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 .

Though so deep the primordial ooze as to defy intuitions, we should nevertheless take a moment to assess what commitments we have made. What does it mean to say that force arrives in little packets and that these packets arrive in a countable number of sorts? Well, 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—indeed, something is a force if and only if it is borne by this element—we can say nothing about the element itself, and thus may only name it FORCE. To name it anything else would be to make a commitment about its properties, and we know only that it bears 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 , then, 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 to say that it is a choice, and that it is a choice that should be made consciously. It simply may be the case that one would want to know about LAND FORCE, SEA FORCE, and AIR FORCE, or about ARMORED VEHICLEs, DESTROYERs, and FIGHTERs, or about M1 ABRAMS, USS ARIZONA, and F-22 RAPTORs. Conversely, if $|L| = \aleph_0$, then there are infinitely many sorts of force elements, which leaves open the possibility that a new sort of tool might be invented that would be useful in war. The choice of L is a choice about the granularity of the model, and it is a choice that should be made consciously. In our interpretation, L is however granular it needs to be as to meaningfully encode the thinking of the force-maker. As we have allowed L to be (countably) infinite if need be, we can always add more elements if we need to, *metatheoretically*, to think about the force in a more granular way. Countability restricts attention to those concepts that can be listed, the way bureaucrats classify, list, and count things. Thus, 18MM MORTAR does not present any special problems even though a gun's caliber could be any real number; when force-makers speak of such things, they speak of them as if they were countable, and so we will treat them as such. This requires some special care, but nothing that cannot be managed.

Though L is a primitive, we retain wide latitude in how it is specified, used, and interpreted. The point here is less to identify actual elements than to acknowledge that the state's force is made up of many different sorts of things. Here the natural familiarity 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. The former is finer than the latter, as it appreciates differences between infantry and armor, between destroyers and carriers, and between fighters and bombers. The latter is coarser than the former, as it lumps these differences together.

Classification schemes facilitate the (good faith) conversation you and I are having right now. Thus, I might refer to ARMORED VEHICLES, or to TANKs, or to the M1 ABRAMS, and you will know what I mean. We can always move from very granular to very coarse classification by way of some aggregation function; given an atom of lithium and an atom of sodium, I can always say I have two atoms of alkali metal. Moving from the coarse to the granular is more difficult, as this aggregation is inherently lossy. If I told you that I had two atoms of alkali metal, you would not know if I had two atoms of lithium, two atoms of sodium, or one of each. But this analysis will proceed as if we could always stop and take more granular inventory of the force elements if we so desired, which means that the model is generally agnostic to whether we're working with L itself or with some classification scheme. We will generally refer to the elements as L itself, but the reader should keep in mind that we could always be working with a classification scheme. This theme will recur throughout the manuscript.

States often strive to keep L under control. In the early days of the Roman Republic, the state might have had little control over the elements of its force, as soldiers were expected to provide their own equipment. By the time of the Roman Empire, the state had taken control of the elements of its force, providing soldiers with standardized equipment and training. As new technologies develop, the degree of standardization waxes and wanes. For example, after the development of gunpowder, muskets became a common sort of force; but, early muskets demonstrated wide variety in terms of caliber, barrel length, and firing mechanism. Maurice of Nassau—Prince of Orange and admirer of Cæsar—famously standardized the musket and the drill, making Dutch forces both more effective and cheaper to build and maintain (McNeill, 1982, Chapter 4).

Chemists talk about more than just raw elements; they also talk about compounds. So too must the theorist of force. Prior to bonds, atoms are just atoms; when bonded, they become something more. Just as three moles of water is not the same as two moles of from and one mole of oxygen, so too is a hoplite distinct from a soldier and shield. Let us define the basic unit of force.

1.3 Definition (Force Molecule)

A force molecule is a connected⁶ labeled graph $(n, E, \ell : \underline{n} \rightarrow L)$ where:

1. $n \in \mathbb{N}$ is the size of the molecule;
2. $E \subseteq \underline{n}^2$ is the edge set of the molecule;
3. $\ell : \underline{n} \rightarrow L$ labels the atoms of the molecule with their elemental type.

The reader should think of a force molecule as a sort of Lego model of a piece of force. The atom set $\underline{n} = \{1, \dots, n\}$ is the set of Lego bricks, the bond set is the set of connections between them, and the labeling is the set of stickers on the bricks. Trivial as it might seem, such a molecule conveys a great deal of information: how many atoms it contains (the atom set), how they are connected (the bond set), and what sorts of atoms they are (the labeling). The edge set is a set of ordered pairs of atoms, representing the bonds between them: soldier to shield, captain to ship, pilot to plane.

Now, strictly speaking, two force molecules might be strictly different even if we think they should be treated as the same. For example, consider two simple molecules defined on 3: one with a soldier holding a shield and a spear, the other with a soldier holding a spear and a shield. These molecules are isomorphic, but not identical: after all, in the first molecule the second atom is labeled SHIELD and the third SPEAR, while in the second molecule the second atom is labeled SPEAR and the third SHIELD. We purchase clarity with pedantry in the following definition, which provides us with our first chance to announce our ignorance of those details encoded via isomorphism.

1.4 Remark (Isomorphism of Molecules)

Two force molecules $M_1 = (n, E_1, \ell_1)$ and $M_2 = (m, E_2, \ell_2)$ are isomorphic just in case there is a bijection $\varphi : \underline{n} \rightarrow \underline{m}$ 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 $M_1 \cong M_2$.

Ismorphic molecules share all structural properties; in this work, we treat them as the same, ignoring tedious labeling details.

⁶By a connected graph, we mean that for any two atoms $a_o, a_\bullet \in \underline{n}$, there is a sequence of vertices a_1, \dots, a_k such that $a_1 = a_o$, $a_k = a_\bullet$, and $(a_i, a_{i+1}) \in E$ for all $i \in \{1, \dots, k-1\}$.

To give an example of a force molecule, consider the hoplite molecule in Figure 1. The hoplite was the standard infantry soldier of ancient Greece. Its

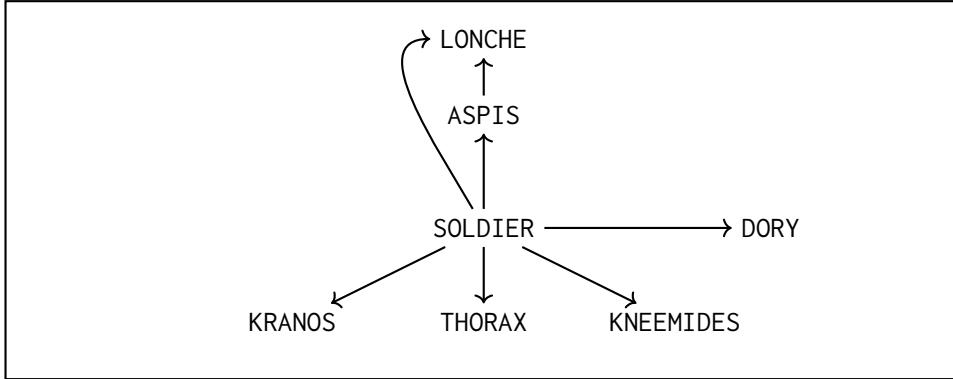


Figure 1: A hoplite molecule, dory/lonche variant.

most characteristic element was a heavy wooden shield, the ἀσπίς (*aspis*), which was often adorned with an identifying device like the λόγχη (*lonche*), the λάμπας (*lampas*), or the ἀστρον (*astron*). The panoply also included a helmet (*κράνος*, *kranos*), a breastplate (*θώραξ*, *thorax*), greaves (*κνημῖδες*, *kneemides*), and a sword (*ξίφος*, *xiphos*) or spear (*δόρυ*, *dory*). Absent a panoply, a hoplite, quite simply, was not a hoplite; in particular, a hoplite without a shield was a ῥιψάσπις (*ripsaspis*), a shield-caster, a term of derision;⁷ the departing soldier was given the command ἢ τὰν ἢ ἐπὶ τὰς (*ē tān ē epi tās*), “with your shield or on it.”

But though a hoplite needed his panoply, a panoply did not a hoplite make. To suggest otherwise would be to ignore the training, the discipline, the battle technique, and the *esprit de corps* that made the hoplite such a formidable force.⁸ Indeed, our machine is designed to avoid the pitfalls of such a reductionist view: the whole may be the same, or more, or less than the sum of its parts, where

⁷ See Donald Kagan and Gregory F. Viggiano’s edited volume *Men of Bronze* (2013), especially the chapter by Viggiano and Hans van Wees (2013), for fascinating details.

⁸ In Section 2, we develop a theory for two important operations on force molecules:

1. the act of gathering various military resources into a single unbonded collection, which we will call a *force configuration*; and
2. the act of turning one force configuration into another, which we will call *force conversion*.

The former involves a binary operation on force molecules (Primitive 2.1) while the latter involves morphisms defined on the free monoid of force molecules generated by that operation (Primitive 2.9). These two primitives form the foundation of our resource theory of force (Section 2). We’ve much to do, and the reader’s kind patience is appreciated.

identity and comparisons are calibrated by a machine of our own making.

The theory takes a set of force molecules as a primitive.

1.5 Primitive (Force Molecules)

There is a set

$$\mathbb{M}_L \subseteq \bigcup_{n \in \mathbb{N}} \{E \in \mathbb{P}_2(\underline{n}) \mid (\underline{n}, E) \text{ is connected}\} \times L^n$$

containing the force molecules. Each force element is a force molecule of size one:

$$(1, \emptyset, \{(1, \ell)\}) \in \mathbb{M}_L \text{ for all } \ell \in L.$$

The reader should think of \mathbb{M}_L as a vast collection of kinds of force, each with its own properties and behaviors. The force molecules are the basic building blocks of the state's force, and the theory will be built up from them. As such, we care about the structure of the set of force molecules; the following useful result gets us started down the path.

1.6 Lemma (Force Molecules are Countable)

\mathbb{M}_L is countable.

Proof. It suffices to show that the right-hand side of the inclusion in Primitive 1.5 is countable. We will do so by showing that it is a countable union of countable sets; choose any $n \in \mathbb{N}$, and observe:

1. The set of labeled connected graphs on \underline{n} is finite. There are $2^{\binom{\underline{n}}{2}}$ possible graphs on \underline{n} ; this sets an upper bound on the number of connected graphs. Now consider a function with domain the connected graphs of \underline{n} and codomain L . Since L is countably infinite, the set of functions from the connected graphs of \underline{n} to L is countably infinite as well.
2. For similar reasons, the set of functions from the finite set \underline{n} to the countable set L is countably infinite.
3. Since the previous two sets are countable, so too is their Cartesian product.

Since this holds for all $n \in \mathbb{N}$, the right-hand side is a countable union of countable sets, and is thus countable; since \mathbb{M}_L is a subset of this set, it is countable as well. We are done. ■

For all its richness, the set of force molecules is countable, and thus isomorphic to some subset of the natural numbers. This is a useful result, as it allows us to enumerate all force molecules, and thus to reason about them systematically without losing track of any of them.

2 Force is a Resource

In Section 1, we developed an atomic theory of force. For all the highfalutin discursiveness, the substantive assertion is quite simple: force is made up of elements, and these elements can be combined into molecules. Formally, this justifies a graph-theoretic approach to force, where the elements are the vertices and the relationships between elements are the edges. The set of all molecules, \mathbb{M}_L , capably represents a wide variety of force configurations.

We justified the atomic approach by arguing that we speak of force as if it were atomic and that we construct molecules from these elements. This is a good start, but it is only a start. To justify deeper interpretations of \mathbb{M}_L and things built from it, we need to dig deeper into the nature of force. To that end, we make our second primary assertion about force: it is a *resource*. The most relevant senses in which the *O.E.D.* defines the word “resource” are as follows:

resource, n. [’ri,sɔrs]

2. A means of supplying a deficiency or need; something that is a source of help, information, strength, etc.
3. *In plural.*
 - a. Stocks or reserves of money, materials, people, or some other asset, which can be drawn on when necessary. In modern use frequently the second element in compounds.
 - b. The collective means possessed by a country or region for its own support, enrichment, or defence.⁹

In all of these senses, a resource is a means to some end: a remedier of deficiencies, a source of help, a stockpile of assets. The relevant concept is deeply pragmatic.

In introducing their “Mathematical Theory of Resources,” Bob Coecke, Tobias Fritz, and Robert W. Spekkens (2016) delineate the *dynamist* and *pragmatic* traditions of theory-building. The former describes systems in the absence of human intervention, where the focus is on the system’s intrinsic properties. The latter describes systems in the presence of human intervention, where the focus is on the system’s utility to humans. Remarkably, “The more a scientific discipline is concerned with aspects close to human life and society, the more relevant this aspect is” (p. 59). It is in the pragmatic tradition that Coecke, Fritz, and Spekkens situate their theory of resources, and it is in this tradition that we situate our theory of force as a resource, too.

⁹Remarkably, the *O.E.D.* supplements this definition a link to the entry for “natural,” as in “natural resources.” Here we think of natural resources as only part of the story, and, what is more, it seems that “natural” is capable of modifying “resources” in other senses of the latter.

2.1 *Mise en Place*

Just as a chemist must gather the substances they wish to combine in various quantities, so too must the force-maker gather the force molecules. The first task is simply to *mise en place*, to gather the necessary ingredients.

2.1 Primitive (Force Configuration)

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

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

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

2. *Associativity: for all $M_1, M_2, M_3 \in \mathbb{M}_L$, we have*¹⁰

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

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

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

The \uplus operator is a way of combining force molecules into a single configuration without performing any conversion. It allows the force-maker to gather the necessary ingredients for the construction of force, namely by using the word “and.” If M_1 is a force molecule representing a soldier and M_2 is a force molecule representing a shield, then $M_1 \uplus M_2$ is a force molecule representing a soldier *and* a shield. This need not be a soldier skilled in the use of a shield, nor a shield designed for a soldier; it is simply a soldier and a shield in the same place at the same time. Likewise, if M_1 and M_2 are two soldiers, there are no implications about their relationship; $M_1 \uplus M_2$ is simply two soldiers in the same place at the same time. The three properties of the \uplus operator are relatively mild and mimic the properties of addition in number theory, which comports with the idea that \uplus represents the word “and” in the present context.

¹⁰ Associativity allows us to write $M_1 \uplus M_2 \uplus M_3$, and indeed to write

$$\biguplus_{i=1}^n M_i := M_1 \uplus \dots \uplus M_n.$$

Letting $\mathcal{M} = M_1, \dots, M_n$ be a sequence of force molecules, we write $\biguplus \mathcal{M}$ for $M_1 \uplus \dots \uplus M_n$. Arbitrary combinations may be written $\biguplus \mathcal{M} \in \mathbb{M}_L^*$, where \mathcal{M} is a sequence of force molecules $\mathcal{M} : \mathbb{N} \rightarrow \mathbb{M}_L$. At risk of obnoxiousness—“risk!”—we explicitly write $\biguplus \mathcal{M}$ when referring to configurations of force molecules; without the \biguplus , we are referring to molecules themselves.

Just as \mathbb{M}_L is a subset of the set of all connected graphs labeled by L , so too is the set of all force configurations a subset of the *free monoid* generated by \mathbb{M}_L under the \uplus operator.

2.2 Primitive (Force Configurations)

The set of all force configurations is

$$\mathbb{M}_L^* \subseteq \widetilde{\mathbb{M}}_L^* := \left\{ \biguplus_{i=1}^n M_i \mid n \in \mathbb{N}, M_1, \dots, M_n \in \mathbb{M}_L \right\}.$$

Naturally, $M \in \mathbb{M}_L^*$ for all $M \in \mathbb{M}_L$.

The concept of a free monoid is quite common in the study of algebraic structures; simply put, it is the set of all combinations of some set of objects under some operation. In number theory, the natural numbers are the free monoid generated by the set $\{1\}$ under the operation of addition; any natural number is a combination of 1 added to itself some number of times. In string theory, the set of all words in some alphabet is the free monoid generated by that alphabet under the operation of concatenation; any word is a combination of letters from the alphabet. And in the present context, the set of all force configurations is the free monoid generated by the set of force molecules under the operation of configuration; any configuration is a combination of force molecules.

Clearly, $\mathbb{M}_L \subseteq \mathbb{M}_L^*$; we have expanded our set of objects of study. However, we have not expanded too much:

2.3 Lemma (Countable Configurations)

\mathbb{M}_L^* is countably infinite.

Proof. It will suffice to show that $\widetilde{\mathbb{M}}_L^*$ is countably infinite, as \mathbb{M}_L^* is a subset of it. We can write $\widetilde{\mathbb{M}}_L^*$ as

$$\bigcup_{n \in \mathbb{N}} \left\{ \biguplus_{i=1}^n M_i \mid M_1, \dots, M_n \in \mathbb{M}_L \right\}.$$

Each of the sets in the union is countable, and the union of countably many countable sets is countable. Thus, \mathbb{M}_L^* is countable. Its infinitude is clear from the fact that we can always add another molecule to a configuration. We conclude that \mathbb{M}_L^* is countably infinite. ■

So, even though we are now talking about combinations of molecules, the structural intuitions we have developed in the study of molecules themselves still apply. We have not lost our way in the dark forest of combinatorial explosion.

The atomic nature of force makes for a more concrete approach to the study of force configurations. There is a natural way to combine connected labeled graphs into a single labeled graph in a way that satisfies these properties, which we will refer to as the *graph union*.

2.4 Definition (Graph Unions)

For force molecules $M_1 = (n_1, E_1, \ell_1 : \underline{n_1} \rightarrow L)$ and $M_2 = (n_2, E_2, \ell_2 : \underline{n_2} \rightarrow L)$, the graph union $M_1 \uplus M_2$ is the force configuration

$$M_1 \uplus M_2 = (n_1 + n_2, E_1 \sqcup E_2, \ell : \underline{n_1 + n_2} \rightarrow L),$$

where the edge set

$$E_1 \sqcup E_2 = \{(i, j) \mid i, j \in \underline{n_1}\} \cup \{(i + n_1, j + n_1) \mid i, j \in \underline{n_2}\}$$

is (essentially) the disjoint union of the edge sets E_1 and E_2 , and where

$$\ell(i) = \begin{cases} \ell_1(i) & \text{if } i \leq n_1 \\ \ell_2(i - n_1) & \text{if } i > n_1. \end{cases}$$

The graph union is a way of combining two connected labeled graphs into a single connected labeled graph. It is viable for configurations:

2.5 Lemma (Graph Unions Configure Molecules)

Graph union satisfies the properties of Primitive 2.1.¹¹

[[Proof](#).]

Observe, however, that the resulting graph (being disconnected) is not itself a molecule, but rather a configuration of molecules—this is just what we wanted from the operation \uplus , which we now know works well for the simplest of configurations, namely the one-molecule configurations. We extend the graph union to configurations in the following definition.

2.6 Definition (Graph Unions on Configurations)

Given two configurations $\biguplus \mathcal{M}_1$ and $\biguplus \mathcal{M}_2$, the graph union $\biguplus \mathcal{M}_1 \uplus \biguplus \mathcal{M}_2$ is

$$\biguplus \mathcal{M}_1 \uplus \biguplus \mathcal{M}_2 = \biguplus_{M_1 \in \mathcal{M}_1} M_1 \uplus \biguplus_{M_2 \in \mathcal{M}_2} M_2.$$

This definition is a natural extension of the graph union to the gathering of configurations; we simply concatenate the concatenated graphs. The properties established in Lemma 2.5 make this an intuitive and useful operation.

¹¹The proof of Lemma 2.5 is straightforward, but it is also a bit tedious and not very illustrative; as such, it is our first proof relegated to Appendix A. The reader ought to be able to prove Lemma 2.5 with a little effort, and indeed it is a good exercise.

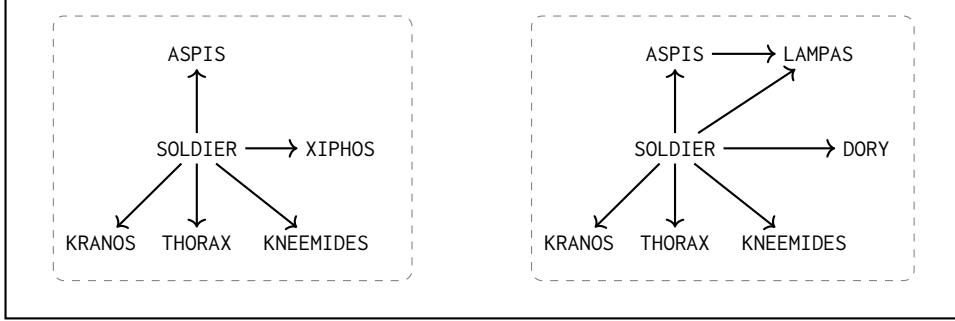


Figure 2: A configuration of two hoplites: one with a sword and unadorned shield, the other with a spear and a decorated shield. The gray dashed rectangles delineate the distinct hoplites, but this is not part of the configuration proper.

For example, consider two hoplites M_1 and M_2 , where the first hoplite carries a sword (a XIPHOS) and an unadorned shield (an ASPIS), and the second hoplite carries a spear (a DORY) and a decorated shield (an ASPIS with a LAMPAS):

$$M_1 = \left(6, \left\{ \begin{array}{l} (1,2), (1,3), \\ (1,4), (1,5), \\ (1,6) \end{array} \right\}, \left\{ \begin{array}{l} (1, \text{SOLDIER}), (2, \text{XIPHOS}), \\ (3, \text{ASPIS}), (4, \text{KRANOS}), \\ (5, \text{THORAX}), (6, \text{KNEEMIDES}) \end{array} \right\} \right),$$

$$M_2 = \left(7, \left\{ \begin{array}{l} (1,2), (1,3), \\ (1,4), (1,5), \\ (1,6), (1,7), \\ (3,4) \end{array} \right\}, \left\{ \begin{array}{l} (1, \text{SOLDIER}), (2, \text{DORY}), \\ (3, \text{ASPIS}), (4, \text{LAMPAS}), \\ (5, \text{KRANOS}), (6, \text{THORAX}), \\ (7, \text{KNEEMIDES}) \end{array} \right\} \right).$$

The graph union of these two hoplites— $M_1 \uplus M_2$ —is a configuration of force molecules representing two hoplites, one carrying a sword and an unadorned shield, and the other carrying a spear and a decorated shield:

$$\left(13, \left\{ \begin{array}{l} (1,2), (1,3), \\ (1,4), (1,5), \\ (1,6), (7,8), \\ (7,9), (7,10), \\ (7,11), (7,12), \\ (7,13), (9,10) \end{array} \right\}, \left\{ \begin{array}{l} (1, \text{SOLDIER}), (2, \text{XIPHOS}), \\ (3, \text{ASPIS}), (4, \text{KRANOS}), \\ (5, \text{THORAX}), (6, \text{KNEEMIDES}), \\ (7, \text{SOLDIER}), (8, \text{DORY}), \\ (9, \text{ASPIS}), (10, \text{LAMPAS}), \\ (11, \text{KRANOS}), (12, \text{THORAX}), \\ (13, \text{KNEEMIDES}) \end{array} \right\} \right).$$

Of course, it is easier to visualize this configuration in a drawing, as in Figure 2. We see that we have two distinct hoplites, each with their own panoply, but the graph union has combined them into a single configuration.

The molecular identities of the hoplites are not, strictly speaking, preserved in the graph union; if presented data in the form of a configuration, we don't necessarily know the molecular identities. That said, it will be useful for us to recover the molecular identities from a configuration. Accordingly, we should consider something akin to an inverse for the \uplus operator. This is *deconfiguration*.

2.7 Primitive (Deconfiguration)

For each force configuration $\biguplus_{i=1}^n M_i \in \mathbb{M}_L^\star$, there is a set of deconfigurations

$$\left\{ \pi_i : \bigcup_{i=1}^n M_i \rightarrow M_i \right\}_{i=1}^n.$$

Less pedantically, one might write

$$\pi_i \left(\biguplus_{i=1}^n M_i \right) = M_i.$$

When useful, we will write

$$\begin{aligned} \pi \left(\biguplus_{i=1}^n M_i \right) &= \left(\pi_1 \left(\biguplus_{i=1}^n M_i \right), \dots, \pi_n \left(\biguplus_{i=1}^n M_i \right) \right), \\ &= (M_1, \dots, M_n). \end{aligned}$$

In the language of category theory, the \uplus operator is a *product* and the π operator is its set of *projections*.¹² Simply put, a product is a way of combining objects, and a projection is a way of decomposing the product into its components. In this sense, the gathering operator \uplus is similar to the Cartesian product of sets, and the deconfiguration operator π is similar to the projection maps from the Cartesian product. But this is not because the action involved is the same; it is because the structure of the operators is the same with respect to the objects they act upon. This is why it works best to focus on the abstract properties of the operators, rather than their concrete actions: the abstract properties are what allow us to reason about the operators in a general way, and the actions make computations possible (and examples illustrative).

¹²For a good start in categorical thinking, accept Brendan Fong and David I. Spivak's *Invitation to Applied Category Theory* (2019). After, move on to Steve Awodey's *Category Theory* (2010) for a more focused treatment. The standard reference is Saunders Mac Lane's *Categories for the Working Mathematician* (1971), but it is a bit more advanced. Emily Riehl's titanic *Category Theory in Context* (2016) fills in the gap between Awodey and Mac Lane.

Primitive 2.7 sets a test for the graph union operator: it must be possible to recover the original molecules from a configuration. We have good news.

2.8 Lemma (Graph Union and Deconfiguration)

For all force configurations $\biguplus \mathcal{M} \in \mathbb{M}_L^*$, there exists a set $\{\pi_i : \biguplus \mathcal{M} \rightarrow M_i\}_{i=1}^n$ as in Primitive 2.7. [Proof.]

This proof—again relegated to Appendix A—is a matter of applying a graph decomposition algorithm to the graph $\biguplus \mathcal{M}$ and then ensuring that the output of that algorithm comports with a given configuration’s component molecules.¹³ This might come across as mere mechanical tedium, but it turns out that Lemma 2.8 is remarkably powerful. After all, it suggests that any configuration of force molecules comes equipped with a ready-made process for recovering the original molecules. Categorically speaking, it suggests that the product-projection relationship discussed above is healthy and well-behaved. Were we to assume that all configurations were possible—*i.e.*, that \mathbb{M}_L^* was the entire free monoid of \mathbb{M}_L under \uplus —then we would be able to begin a category from the molecules alone and assert the existence of all products. This is a very desirable property, as it entails several others; we have not yet introduced the machinery required to fully appreciate this, but we will do so in due course.

Summary of Section 2.1. We began Section 2.1 with a set of force molecules \mathbb{M}_L , itself a subset of the set of all finite connected graphs with labels drawn from the force elements L . We then introduced the \uplus operator, which allows us to combine force molecules into force configurations. In the abstract, we require nothing from the \uplus operator other than that it be associative, commutative, and unital, which helps it mimic operations like addition, union, and concatenation. Indeed, the properties of addition and union imbue the graph-theoretic \uplus operator with the desired properties when it is defined as a simple graph union. We then extended the \uplus operator to configurations of force molecules, which we showed to be a countably infinite set \mathbb{M}_L^* , a subset of the free monoid generated by \mathbb{M}_L under \uplus . This construction contains all allowable configurations of force molecules, and the graph union operator that generates it is well-behaved in the sense that it is deconfigurable. We will now turn our attention to actions

¹³In particular, the proof outlines a *depth-first search* algorithm for decomposing a graph into its maximal connected subgraphs. Other popular approaches include the *breadth-first search* algorithm, the *Kosaraju-Sharir* algorithm, and the *Tarjan* algorithm. These algorithms are well-studied and well-understood, and they are all capable of decomposing a graph into its maximal connected subgraphs. See the proof for references.

within the set of force configurations, which will require further theoretical development. Turn the page.

2.2 The Conversion of Force

Consider the phalanx, which (to repeat) is a particular formation comprised of hoplites. A mere configuration of hoplites is not a phalanx, as the latter stipulates a particular arrangement of the former. The whole, so to speak, must become more than the sum of its parts. In the previous section, we spent time on what the sum of the parts might be; this quite literally took the form of a summation-like operator, \sqcup . But, it has come time for us to think more about relationships between the parts, and how they might be transformed into something new. The following definition is a first step in this direction.

2.9 Primitive (Conversion Morphisms)

For any pair of force configurations $\sqcup \mathcal{M}_1, \sqcup \mathcal{M}_2 \in \mathbb{M}_L^*$, there is a set of conversion morphisms $\text{Hom}(\sqcup \mathcal{M}_1, \sqcup \mathcal{M}_2)$. Either of the two notations

$$\sqcup \mathcal{M}_1 \xrightarrow{f} \sqcup \mathcal{M}_2 \iff f \in \text{Hom}(\sqcup \mathcal{M}_1, \sqcup \mathcal{M}_2)$$

mean process f converts configuration $\sqcup \mathcal{M}_1$ into configuration $\sqcup \mathcal{M}_2$.

By construction:

1. for all $\sqcup \mathcal{M} \in \mathbb{M}_L^*$, we have $\text{id}_{\sqcup \mathcal{M}} \in \text{Hom}(\sqcup \mathcal{M}, \sqcup \mathcal{M})$;
2. for all $\sqcup \mathcal{M}_1, \sqcup \mathcal{M}_2 \in \mathbb{M}_L^*$, we have $\mathfrak{x} \in \text{Hom}(\sqcup \mathcal{M}_1, \sqcup \mathcal{M}_2)$; and
3. if $(\sqcup \underline{\mathcal{M}}_1, \sqcup \overline{\mathcal{M}}_1) \neq (\sqcup \underline{\mathcal{M}}_2, \sqcup \overline{\mathcal{M}}_2)$, then

$$\text{Hom}(\sqcup \underline{\mathcal{M}}_1, \sqcup \overline{\mathcal{M}}_1) \cap \text{Hom}(\sqcup \underline{\mathcal{M}}_2, \sqcup \overline{\mathcal{M}}_2) = \{\mathfrak{x}\}.$$

The set of all conversion morphisms is written

$$\text{Hom}(\mathbb{M}_L^*) := \bigcup_{(\sqcup \mathcal{M}_1, \sqcup \mathcal{M}_2) \in \mathbb{M}_L^* \times \mathbb{M}_L^*} \text{Hom}(\sqcup \mathcal{M}_1, \sqcup \mathcal{M}_2).$$

We think of the conversion morphisms as the processes that convert one set of resources into another; metaphorically, these are the names of the various chemical reactions that turn component pieces into something new. Though a mouthful, Primitive 2.9 bakes little structure into the apparatus: we require only that the ways of converting one collection of molecules into another be set-containable, that each configuration is equipped with a special process labeled $\text{id}_{\sqcup \mathcal{M}}$, that each pair of configurations contains a special process labeled \mathfrak{x} , and that no two different pairs of configurations share a process other than this \mathfrak{x} process. These special processes warrant some discussion.

Maintenance and the Identity Morphism. Each configuration of force molecules $\biguplus \mathcal{M} \in \mathbb{M}_L^*$ is equipped with a special process $\text{id}_{\biguplus \mathcal{M}} \in \text{Hom}(\biguplus \mathcal{M}, \biguplus \mathcal{M})$. In mathematical parlance, this is called the *identity morphism* for the configuration $\biguplus \mathcal{M}$. Substantively, this process does nothing, merely “converting” the configuration into itself:

$$\biguplus \mathcal{M} \xrightarrow{\text{id}_{\biguplus \mathcal{M}}} \biguplus \mathcal{M}.$$

Stretching our interpretation a bit, we might think of this process as the *maintenance* of a configuration; after all, many force configurations take effort to maintain. A soldier must be fed, clothed, and trained; a ship must be crewed, provisioned, and maintained; a fort must be governed, defended, and supplied. And, it seems that there may be many ways to maintain a configuration, not just our special process $\text{id}_{\biguplus \mathcal{M}}$. This is a feature, not a bug; the primitive does not require that there be only one way to maintain a configuration, only that there be at least one way. We will see that $\text{id}_{\biguplus \mathcal{M}}$ is a special maintenance process attached to the interpretation of maintenance *at no cost*. This is a key feature of the identity morphism and the source of its appeal for given applications where the nuances of maintenance are not of interest.

Impossibility and the Null Morphism. Each hom set $\text{Hom}(\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_2)$ includes a special process \mathbf{x} that converts $\biguplus \mathcal{M}_1$ into $\biguplus \mathcal{M}_2$. Rather than a standard construction like the identity morphism, this process is special to our theory of force. It is not essential to the theory but rather is an easy shortcut for the force-maker to know the word “impossible” in a literal sense. If the force-maker is under the impression that there is no way to convert one configuration into another, we use the \mathbf{x} process to indicate this impossibility. In case $\text{Hom}(\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_2) = \{\mathbf{x}\}$, there is literally no way to convert $\biguplus \mathcal{M}_1$ into $\biguplus \mathcal{M}_2$. If instead we had $\text{Hom}(\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_2) = \{\phi, \mathbf{x}\}$, then ϕ is a process that converts $\biguplus \mathcal{M}_1$ into $\biguplus \mathcal{M}_2$ and \mathbf{x} represents ignorance, or rejection, of ϕ .

Uniqueness and the Null Intersection. The final property of the conversion morphisms is that no two different pairs of configurations share a process other than the \mathbf{x} process. This is more of a labeling convention than a substantive assertion; it is a way of ensuring that the force-maker knows which process is being referred to when a process is mentioned. Suppose, for example, that one could f from $\biguplus \mathcal{M}_1$ to $\biguplus \mathcal{M}_2$ and also f from $\biguplus \mathcal{M}_3$ to $\biguplus \mathcal{M}_4$. This could well be the same process, suggesting deep similarities between the odd- and even-numbered configurations. But it could also be two different processes that happen to have the same name, introducing unnecessary confusion. Here we err in the direction that the force-maker should be able to tell which process is being referred to when a process is mentioned.

A few examples might help to clarify the concepts introduced in this section.

1. Consider the process of training a hoplite. One key aspect of the philosophy here is that a hoplite is different than just a soldier and the appropriate panoply; a hoplite is a soldier trained in a particular way, so that the bonds between soldier and shield are strong. Recall from Figure 1 that the input elements for a hoplite molecule include SOLDIER, ASPIS, DORY, and so on. Gather these into the sequence \mathcal{H} , which is the collected raw materials for a hoplite. However, gathering the materials is not enough; they must be converted into a hoplite. We therefore introduce the training process f that converts the raw materials into a hoplite; this yields

$$\biguplus \mathcal{H} \xrightarrow{f} \text{HOPLITE}.$$

This is the process of training a hoplite. There may, of course, be several such methods of training a hoplite—say, f_1 , f_2 , and f_3 —each of which is a different way of converting the raw materials into a hoplite. Thus, $f_1, f_2, f_3 \in \text{Hom}(\biguplus \mathcal{H}, \text{HOPLITE})$.

2. Elaborating on this idea, we might think of dedicated training molecules that serve as catalysts for the training process. These molecules might be combined with the raw materials to produce a hoplite more quickly or more effectively—or in any case, via some process other than the standard training process. To pin the idea down, suppose \mathcal{H} is the sequence of atoms required for a hoplite, and \mathcal{T} is the sequence of atoms required for a dedicated training molecule—for lack of a better term, call this a *trainer*. We might train the trainer via some process

$$\biguplus \mathcal{T} \xrightarrow{g} \text{TRAINER},$$

and then combine the trainer with the raw materials to produce a hoplite via some process

$$\biguplus \mathcal{H} \uplus \text{TRAINER} \xrightarrow{h} \text{HOPLITE} \uplus \text{TRAINER}.$$

Process h , where a trainer is present, is different from process f , where a trainer is not present. Here we have h retain the trainer, as the training of a hoplite rarely expends the trainer. Just as heat and pressure can catalyze chemical reactions, so too can trainers catalyze the training of hoplites. Naturally, one can extend this idea to other contextualizing factors for training, such as the physical environment, the quality of food and other resources, and so on. The gathering construction \uplus allows us to consider all of these factors together.

3. One may consider multiple hoplites trained at the same time under the tutelage of a particular trainer. This is a process of training multiple hoplites at once, and it is a different process from training each hoplite individually. We might have a process that looks like

$$\left(\bigcup_{i=1}^n \mathcal{H} \right) \uplus \text{TRAINER} \xrightarrow{f_n} \left(\bigcup_{i=1}^n \text{HOPLITE} \right) \uplus \text{TRAINER}.$$

This is a different process from training each of the n hoplites individually; perhaps *esprit de corps* is developed, or perhaps the trainer is able to focus on the group as a whole. In such situations, we use the simpler notation

$$(n \times \bigcup \mathcal{H}) \uplus \text{TRAINER} \xrightarrow{f_n} (n \times \text{HOPLITE}) \uplus \text{TRAINER}.$$

This is a process of training n hoplites at once, akin to making n moles of water from n moles of hydrogen and $n/2$ moles of oxygen. The \times operator is a shorthand for repeated application of \uplus .

4. This construction allows us to consider preparedness on a larger scale. Consider the *phalanx*, a formation of hoplites that is more than the sum of its parts. In a phalanx formation, hoplites stand shoulder to shoulder, shields overlapping, so that the formation is more than just a collection of hoplites. Indeed, one hoplite's shield protects both himself (from the front) and the neighbor on his left (from the right), so that the bond set up by the shield is not just between hoplite and shield, but also between hoplite and neighbor. The phalanx is a different force molecule from its underlying collection of hoplites, and so the process by which hoplites are converted into a phalanx—training, marching orders, and so on—is of potential interest. In this case, we have the simple process

$$n \times \bigcup \mathcal{H} \xrightarrow{\tau_n} n \times \text{HOPLITE} \xrightarrow{\phi_n} \text{PHALANX},$$

where τ_n trains n hoplites and ϕ_n trains them to fight as a phalanx. Naturally, this process can be repeated to allow for higher-order constructions made out of phalanxes; composition will be discussed in due course.

These examples do not exhaust the possibilities of the machine introduced in this section, but they do give a sense of the kinds of processes that can be considered. In general, we have developed a theory for the gathering of multiple force molecules and the conversion of these molecules into new molecules (or into the same molecules in different configurations). The construction makes clear the resource-like nature of force, and it allows us to consider the ways in which force is gathered and converted.

The final example given above points toward an important aspect of conversion processes: they can be *composed*. Just as one can convert hydrogen and oxygen into water and then convert water into steam, so too can one convert materials into hoplites and then convert hoplites into phalanxes. This is a key aspect of the conversion processes, formalized as follows.

2.10 Primitive (Composition of Conversion Morphisms)

For every process $f \in \text{Hom}(\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_2)$ and $g \in \text{Hom}(\biguplus \mathcal{M}_2, \biguplus \mathcal{M}_3)$, there is a process $g \circ f \in \text{Hom}(\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_3)$. This process is the composition of f and g . The composition satisfies the following properties for all force configurations:

1. **Associativity:** for all $\biguplus \mathcal{M}_1 \xrightarrow{f} \biguplus \mathcal{M}_2 \xrightarrow{g} \biguplus \mathcal{M}_3 \xrightarrow{h} \biguplus \mathcal{M}_4$, we have

$$h \circ (g \circ f) = (h \circ g) \circ f;$$

2. **Identity:** for all $f \in \text{Hom}(\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_2)$, we have

$$f \circ \text{id}_{\biguplus \mathcal{M}_1} = f = \text{id}_{\biguplus \mathcal{M}_2} \circ f; \text{ and}$$

3. **the Dominance of Impossibility:** for all $f \in \text{Hom}(\biguplus \mathbb{M}_L^{\star})$, we have

$$f \circ \mathbf{x} = \mathbf{x} = \mathbf{x} \circ f.$$

The composition of conversion morphisms is a way of chaining together conversion processes. Consider again the example of training hoplites to fight as a phalanx, encoded as

$$n \times \biguplus \mathcal{H} \xrightarrow{\tau_n} n \times \text{HOPLITE} \xrightarrow{\phi_n} \text{PHALANX}.$$

Primitive 2.10 allows us to consider the process of training hoplites to fight as a phalanx as a single process, $\phi_n \circ \tau_n$, which we could encode:

$$n \times \biguplus \mathcal{H} \xrightarrow{\phi_n \circ \tau_n} \text{PHALANX}.$$

There may be (many) other processes that convert n hoplites into a phalanx; the primitive does not require that there be only one such process. However, it does require that there *is* such a process and that this process can be composed from the training process and the phalanx formation process. We retain wide latitude about what sorts of processes we are composing and how they are composed, but we have a clear way of reasoning about multi-step processes. All told, the process of composition and the properties of Associativity and Identity comprise the core of categorical reasoning: what makes a structure what it is is how its processes can be composed and how they interact with the identity processes.

As we are working with a categorical theory, we can represent the force configurations and conversion morphisms in so-called *categorical diagrams*. Indeed, the examples above are written in a way that mimics the categorical diagrams, but these diagrams can be more general than the examples suggest. An example will help: suppose we had three force configurations $\biguplus \mathcal{M}_1$, $\biguplus \mathcal{M}_2$, and $\biguplus \mathcal{M}_3$, and suppose further that we had conversion morphisms $f \in \text{Hom}(\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_2)$, $g \in \text{Hom}(\biguplus \mathcal{M}_2, \biguplus \mathcal{M}_3)$, and $h \in \text{Hom}(\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_3)$. By composition, we also have the morphism $g \circ f \in \text{Hom}(\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_3)$. Each hom set includes \mathbf{x} , the null morphism, and each configuration includes $\text{id}_{\biguplus \mathcal{M}}$, the identity morphism. For sake of variety, suppose also that there is some $k \in \text{Hom}(\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_1)$. We can represent these processes as follows:

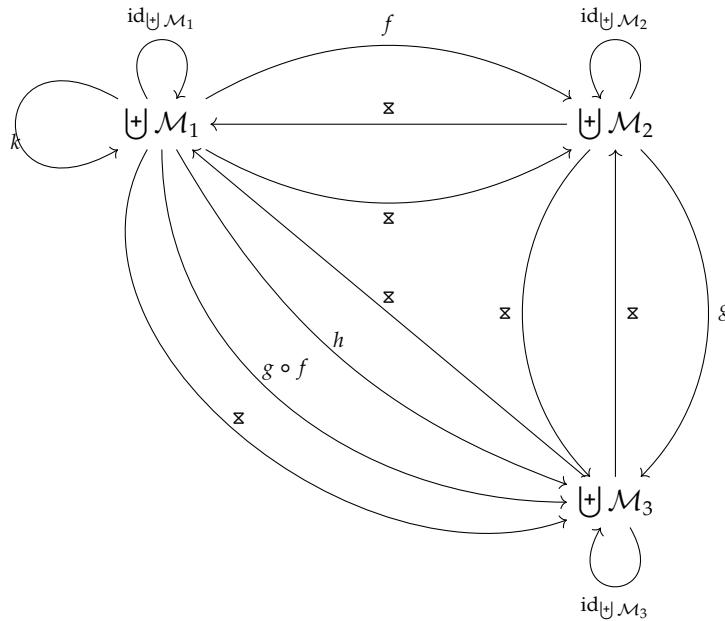


Figure 3: Sample categorical diagram of force configurations and conversion morphisms. For clarity, \mathbf{x} endomorphisms are omitted from the diagram.

This diagram conveys all structural information about the force configurations and the conversion morphisms. The identity morphisms (and k) are self-directed loops staying at the same configuration,¹⁴ and the conversion morphisms are directed arrows moving from one configuration to another. The f and g arrows compel the existence of the $g \circ f$ arrow.

¹⁴Strictly speaking, each configuration ought to have a self-directed loop for the null morphism \mathbf{x} as well, but these are omitted for clarity.

Just as in Section 2.1, we turn our attention from abstract details to the concrete world of the atomic theory of force—*i.e.*, to configurations of force molecules, themselves connected graphs of atoms with labels drawn from the set L . We now ask whether there are simple, graph-theoretic operations for conversion morphisms between force configurations.

To fix intuitions, consider again the two hoplites of Figure 2. The two hoplites

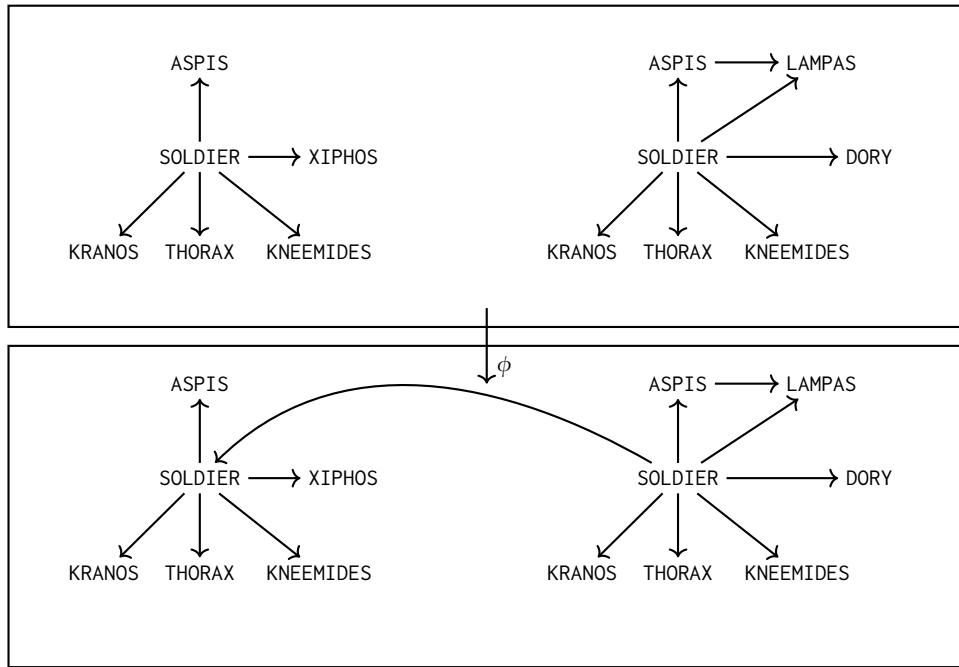


Figure 4: Fashioning a phalanx from two hoplites. The top panel is $M_1 \uplus M_2$, where M_1 is the hoplite on the left and M_2 is the hoplite on the right. The bottom panel is M_3 , the phalanx formed from the two hoplites. We have $\phi \in \text{Hom}(M_1 \uplus M_2, M_3)$.

can be combined into a single force configuration, $M_1 \uplus M_2$, which is the raw material for a phalanx. This is the top panel of Figure 4: two distinct subgraphs in a disconnected union. The phalanx is a different force configuration, and though one could imagine any number of ways of representing the differences between the two hoplites and the phalanx with respect to the graphs, but it's simple enough to say that in our new phalanx, the hoplite on the right protects the hoplite on the left. This is the state of affairs depicted in the bottom panel of Figure 4, where we have added one edge between the two hoplites. This edge—this mere mathematical plaything—represents the bond between the two hoplites, the essence of the phalanx.

The example suggests that reasonably simple graph-theoretic operations can be used to convey the changing of one force configuration into another. Here is another. Just as we added edges to graphs in the previous example, we might also add vertices to the graph, say to demonstrate the training of a soldier in the use of a particular piece of equipment. Indeed, suppose we wanted to represent the training of a soldier into a full-fledged hoplite. Our initial force configuration might be a single vertex labeled **SOLDIER**, and the final force configuration might be a hoplite as we've been discussing. In this case, we need to add vertices

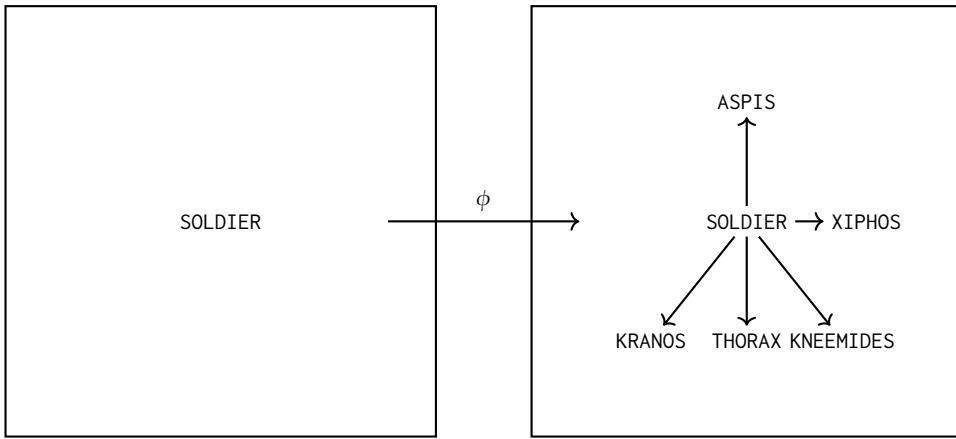


Figure 5: Training a soldier into a hoplite. The left panel is M_1 , the molecule with a single atom of element **SOLDIER**. The right panel is M_2 , the molecule with a full hoplite configuration. We have $\phi \in \text{Hom}(M_1, M_2)$.

for the various pieces in a hoplite's panoply and edges to connect them to the soldier, representing the training process. Thus, we need to add edges as in the previous example, but in tandem we need to add vertices. Again, this is a reasonably simple graph-theoretic operation, and it is enough to represent both the equipping (adding vertices) and the training (adding edges) of a soldier into a hoplite.¹⁵ These simple tools can move mountains, or at least hoplites.

¹⁵It should be noted that adding vertices does not *require* us to add edges; the output need not be connected, as we are linking configurations to configurations rather than molecules to molecules. It just so happens that this simple example is molecule-to-molecule. The reader is encouraged to draw out the version of hoplite training that includes a non-connected **TRAINER** as a catalyst, connected neither to the soldier nor to any of the hoplite's equipment. Literally, you'll just add a vertex labeled **TRAINER** to each panel of Figure 5, making sure not to connect it to anything. You might add a subscript to ϕ to drive home the point that these are two distinct training processes: one with a trainer and one do-it-yourself. Your humble author worries that he's already overtaxed the reader's patience with the diagrams in this section, so he refrains from adding yet another. Plus now the manuscript is interactive.

The previous two examples demonstrated how we could move from one configuration to another by adding edges and vertices; it goes nearly without saying that we could have done the same in reverse by deleting edges and vertices. For example, it was not all that practical for hoplites to remain forever fixed in a phalanx. Just as they could be joined, so too could they be separated.¹⁶ Likewise, a hoplite could be stripped of his equipment as part of a re-equipping process.¹⁷ The point is that the addition and deletion of edges and vertices are simple, graph-theoretic operations that can be used to represent the conversion of one force configuration into another. Let us now encode these four basic operations as *elementary edits* for force configurations.

2.11 Definition (Elementary Edits)

For force configurations

$$\bigcup \mathcal{M}_0 = (n_0, E_0, \ell_0 : \underline{n}_0 \rightarrow L) \in \mathbb{M}_L^*,$$

we define four elementary edits:

1. Vertex Addition: set $n_1 = n_0 + 1$, $E_1 = E_0$, and $\ell_1(i) = \ell_0(i)$ for all $i \in \underline{n}_0$, and $\ell_1(n_1) = \ell_{new}$ for some $\ell_{new} \in L$;
2. Vertex Deletion: set $n_1 = n_0 - 1$, and for the deleted $k \in \underline{n}_0$, set $E_1 = E_0 \setminus \{(k, j) \mid j \in \underline{n}_0\}$ and $\ell_1 = \ell_0|_{\underline{n}_0 \setminus \{k\}}$;
3. Edge Addition: set $E_1 = E_0 \cup \{(i, j)\}$ for some $i, j \in \underline{n}_0$, and leave $n_1 = n_0$ and $\ell_1 = \ell_0$ unchanged; and
4. Edge Deletion: set $E_1 = E_0 \setminus \{(i, j)\}$ for some $i, j \in \underline{n}_0$, and leave $n_1 = n_0$ and $\ell_1 = \ell_0$ unchanged.

¹⁶In this sense, the morphism sending hoplites to a phalanx can represent training, but it just as easily could represent the formation of a phalanx from a group of hoplites. We will soon enough think about how to use such an interpretation to encode the readiness of a force.

¹⁷This was certainly the case with the Zulu Impi warrior. As part of a re-outfitting of his forces, the great leader Shaka kaSenzangakhona forced his warriors to discard their sandals—ISICOCO—and to go barefoot. He switched out their long spears—UMKHONTO, effective for defensive fighting, but useless for the aggressive tactics he had in mind—for shorter stabbing spears called IKLWA. (The latter was named for the sound it made when it was pulled from the body of an enemy; upon removing his blade from the enemy, Shaka would shout “I have eaten!”) The reader is encouraged to take long visits with Donald R. Morris’s *The Washing of the Spears* (1965) and John Laband’s *The Rise and Fall of the Zulu Nation* (1995).

These elementary edits are the building blocks for the conversion of force configurations. They allow us to render configurations more complicated or simpler by adding or deleting vertices or edges.

Obviously—any yet still remarkably!—these edits allow us to convert any configuration into another, so long as we are sufficiently patient. The following is likely obvious, but it is important enough to be stated explicitly.

2.12 Lemma (Elementary Edits Suffice for Conversion)

For any two force configurations $\biguplus \mathcal{M}_1$ and $\biguplus \mathcal{M}_2$, there exists a sequence of elementary edits converting $\biguplus \mathcal{M}_1$ into $\biguplus \mathcal{M}_2$.

Proof. We need not identify an efficient sequence of elementary edits, so let us keep matters simple by breaking the job into two simple parts.

Part 1: converting $\biguplus \mathcal{M}_1$ into $\mathbb{0}_{\mathbb{M}_L}$ by deleting all edges, then all vertices. The details are left to the reader, but the idea is simple: delete edges until $E = \emptyset$, then delete vertices until $n = 0$. This is a sequence of elementary edits that converts $\biguplus \mathcal{M}_1$ into $\mathbb{0}_{\mathbb{M}_L}$.

Part 2: converting $\mathbb{0}_{\mathbb{M}_L}$ into $\biguplus \mathcal{M}_2$ by adding all vertices, then all edges. The details are again left to the reader, but the idea is the same in reverse: add vertices, each time assigning $\ell_{\text{new}} = \ell_2(i)$ for $i \in \underline{n}_2$, then add edges until $E = E_2$. The resulting graph is $\biguplus \mathcal{M}_2$.

Conclusion. We may concatenate these two sequences of elementary edits to obtain a sequence of elementary edits converting $\biguplus \mathcal{M}_1$ into $\biguplus \mathcal{M}_2$. ■

This lemma is a powerful result, as it shows that the elementary edits are sufficient to convert any configuration into any other configuration via a sequence of elementary edits. Indeed, for any two configurations there are many such sequences, and the choice of sequence may depend on the context of the conversion.¹⁸ The reader is encouraged to identify multiple sequences of edits for the conversion morphism of Figure 5. Even more enterprising readers should construct other edits out of the elementary edits—for example, what sequence of elementary edits would be equivalent to simply relabeling one, and only one, vertex in a configuration? (*Hint:* the sequence’s complexity depends on how connected the vertex is to the rest of the configuration.)

¹⁸Strictly speaking, there are infinitely many sequences of elementary edits that convert one molecule into another, since any change we make in one iteration can be undone in the next. Thus, we could toggle back and forth between two molecules for all eternity, and each length-of-toggle would represent a distinct sequence of elementary edits. This is a bit of a silly point, but it is technically correct; we will ignore it for the rest of the manuscript. After all, we will eventually be looking for cost-minimizing sequences of elementary edits, and toggling back and forth between two molecules is not cost-minimizing. We required only that each hom set is a set, not that it is small. And indeed, even with these silly sequences, the set of sequences is still a set.

We can now gather all of the above into a single, coherent package.

2.13 Construction (The Category of Force Configurations)

We construct the category of force configurations $\text{Config}_{\mathbb{M}_L^*}$ as follows:

1. the objects are the force configurations \mathbb{M}_L^* ;
2. the morphisms are the sequences of elementary edits, plus \mathbf{x} ;
3. the identity morphisms are the 0-sequences of edits; and
4. composition of morphisms is concatenation of sequences of elementary edits, along with the special case of \mathbf{x} under the Dominance of Impossibility.

We observe that the concatenation of sequences of elementary edits is itself a sequence of elementary edits, and so the composition of morphisms is a morphism. Further, it clearly satisfies the associative property, and the identity morphisms are clearly identities.

This construction provides a formal framework for the conversion of force configurations. It provides us with the structure we need to further study force configurations and their conversions, particularly from a political-economic perspective when we introduce costs.

We have used the atomic theory of force to imbue force configurations with a rich structure, but it is worth noting that this structure can be forgotten when its details obfuscate the broader picture. It might not be obvious, but in the previous two subsections we have developed a theory of force configurations that is quite general—more general than one that relies on the atomic theory of force and graph-theoretic encodings. The atomic theory of force is a useful tool for understanding the structure of force configurations, but it is not the only tool, and it is hoped that future work will explore other ways of understanding the structure of force configurations within this general framework.

Summary of Section 2.2. In this section, we extended our theory of force configurations to consider the processes by which one configuration is converted into another. We introduced the notion of a conversion morphism, which in the abstract is simply an arrow between two force configurations. In the context of the atomic theory of force, we found that the conversion morphisms could be represented as sequences of elementary edits, which are simple graph-theoretic operations. We showed that these elementary edits are sufficient to convert any force configuration into any other force configuration, and we used this result to construct the category of force configurations. This category provides a formal structure for the conversion of force configurations, and it will be the foundation for our study of the costs of conversion in coming sections.

2.3 The Costs of Conversion

To this point, we have developed a theory of force configurations (themselves composed of force molecules, themselves composed of the elements of force) and the conversions between them. The theory of conversions was introduced in a general, abstract way that (it turns out) was amenable to being rendered concrete in terms of graphs. We continue in this style as we turn attention to the costs of conversion between force configurations, starting from abstract first principles before moving to the concrete world of connected labeled graphs.

In the general setting we have developed, conversion processes are simply arrows sending one force configuration to another. If our task is to assign each of these arrows some datum called a “cost,” we must first understand what the structure of that datum is. Figuring out what sort of datum to attach to a given concept is a famously difficult problem. For example, in classical physics the concept of time is attached to the real numbers \mathbb{R} and the concept of space to the Euclidean space \mathbb{R}^3 . Irrespective of the units we use to navigate these spaces, the underlying structure is the same. So deeply-ingrained is this structure that it is difficult to imagine a different one, which is one reason why so many people find modern physics so difficult to understand. There, time and space are no longer real numbers and Euclidean spaces, but rather parts of a four-dimensional manifold called spacetime equipped with some additional features. The structure of spacetime is different from that of time and space; the machines used to measure spacetime are calibrated differently from those used to measure time and space.¹⁹ And we must ask ourselves here: how do we calibrate the machine measuring the costs of conversion between force configurations?

Your humble author does not want to make promises he cannot keep, so he will not promise to deliver an actual calibration capable of handling soldiers-to-hoplites as capably as uranium-to-warheads. The goal here is not to provide an actual calibration so much as it is to announce what properties such a calibration would have to possess to (1) meaningfully encode the costs of conversion given the ways we talk about such things; and (2) be amenable to the abstract machinery we have developed. What makes for a good measure of cost is a question for the philosophers, but what makes for a good measure of cost *in our theory* is a question for us. Indeed, it is a deep responsibility, given the generality for which we strive, to ensure that the structures we introduce are capable of handling the wide variety of costs one might encounter.

¹⁹The example might be underwhelming given that the structure of spacetime is still quite similar to that of time and space, but the point is that the structure of the datum attached to a concept is not necessarily the same as the structure of the concept itself. One can consider other foundational aspects like charge, spin, or color, none of which admit real-number structure.

How do we begin to understand the structure of the costs of conversion? We first assert the existence of a set of labels used to measure costs, which alerts us to what sort of data we are dealing with.

2.14 Primitive (Cost Labels)

There is a nonempty set Ξ containing the cost labels.

This primitive is the most basic one we can introduce, as it simply asserts the existence of a set of labels used to measure the costs of conversion. It is, to this point, an empty canvas awaiting further details. At present, we have asserted nothing other than the fact that the things we used to measure costs are set-containable. The set might be very simple, like $\{\text{Lo}, \text{Hi}\}$, or some larger finite set $\underline{n} = \{1, 2, \dots, n\}$, or it might be more complex, like the real numbers \mathbb{R} . It might be unidimensional, as in these previous examples, or it might be multidimensional, as in the case of vectors or tensors. It might come equipped with particular properties like order, distance, or topology, or it might be a more abstract structure like a group or a ring. It might be a pointed set, containing at least one element distinguished from the others, or it might be a monoid, containing an operation that combines elements in a particular way. It might combine some or all of these properties into something quite complicated and difficult to understand, or it might distill them into something exquisitely simple. The set of cost labels is a blank slate, and we will fill it in as we go along.

Its job is to calibrate cost measurement.

2.15 Primitive (The Cost Map)

There is a function $\text{cost} : \text{Hom}(\mathbb{M}_L^) \rightarrow \Xi$ assigning a cost to each process.*

This primitive asserts the existence of a function that assigns a cost to each conversion process. But in a deeper sense, it asserts that the cost labels are the things we speak of when we speak of costs. After all, cost and Ξ go hand in hand: what cost is is the function that takes a conversion process and assigns it a cost label from Ξ , and what Ξ is is the set of things targeted by the function cost.²⁰ This is what is meant by “calibrating the machine:” the function cost is the machine, and Ξ is the calibration. In studying Ξ we study the way costs are understood and reasoned about, interpreted and manipulated. As our understanding of cost evolves in time, space, and context, so too must we evolve our understanding of Ξ —and as such, we must assert what properties Ξ should possess to do its job well. We do so in the next few pages.

²⁰This is one reason category theorists are quite doctrinaire about the sources and targets of morphisms: “the same” function defined on different sets is not the same function.

We need some way to compare costs, and the most common parlance used in such comparisons seems to involve words like “more” and “less.” As such, Ξ ought to come equipped with some way to order its elements. We introduce this structure as follows.

2.16 Primitive (Order Structure on Cost Labels)

There is a preorder \geq on the set of cost labels Ξ ; we have

$$\xi_1 \geq \xi_2 \quad \rightsquigarrow \quad \text{“}\xi_1 \text{ is at least as great as } \xi_2\text{.”}$$

Three matters of notation:

1. *we may write $\xi_1 > \xi_2$ just in case $\xi_1 \geq \xi_2$ and not $\xi_2 \geq \xi_1$;*
2. *we may write $\xi_1 \leq \xi_2$ just in case $\xi_2 \geq \xi_1$; and*
3. *we may write $\xi_1 \doteq \xi_2$ just in case $\xi_1 \geq \xi_2$ and $\xi_2 \geq \xi_1$.*

This is a very general assertion, as it does not specify how the costs are ordered, only that they are ordered. We require only that the order satisfies:

1. *Reflexivity:* for all $\xi \in \Xi$, $\xi \geq \xi$, so any cost is at least as great as itself; and
2. *Transitivity:* for all $\xi_1, \xi_2, \xi_3 \in \Xi$, if $\xi_1 \geq \xi_2$ and $\xi_2 \geq \xi_3$, then $\xi_1 \geq \xi_3$, so that the costs are ordered in a consistent way.

We do *not* require completeness, which is to say that we do *not* require that any two cost labels are comparable. This is a deliberate choice, as we do not want to impose any more structure on the cost labels than is strictly necessary. It could well be that some costs are not comparable, and we want our structure to be able to handle that possibility. The atomic structure will be of some use here later.

The preorder \geq reflects the way the force-maker thinks about costs in that it tells us which she deems more or less expensive. There are wide varieties in how she might do so—for example, she might say all costs are the same, or that all distinct costs are incomparable, or that there is some strict ordering of costs. Though the origins of her preferences here lie beyond our scope, we at least store all the possibilities in a set.

2.17 Definition (Set of All Cost Labels)

The set of all preorders on Ξ is denoted $\text{Pre}(\Xi)$.

$\text{Pre}(\Xi)$ is a well-studied construction in set theory, and it possesses several remarkable properties. We will discuss some of them at appropriate moments in the text, but for now we are content to have a set of all preorders on Ξ .

Next, we need to combine costs in some way. In light of sequential processes encoded by composed conversion morphisms, it is only natural to consider the sum of costs, which we introduce as follows.

2.18 Primitive (Sum of Costs)

There is a binary operation $\oplus : \Xi \times \Xi \rightarrow \Xi$ representing the sum of costs. It satisfies:

1. *Monotonicity: for all $\xi_1, \xi_2, \xi^1, \xi^2 \in \Xi$,*

$$\xi_1 \geq \xi^1 \quad \text{and} \quad \xi_2 \geq \xi^2 \implies \xi_1 \oplus \xi_2 \geq \xi^1 \oplus \xi^2;$$

2. *Unitality: there is a void cost $0_\Xi \in \Xi$ such that for all $\xi \in \Xi$,*

$$\xi \oplus 0_\Xi = \xi = 0_\Xi \oplus \xi;$$

3. *Associativity: for all $\xi_1, \xi_2, \xi_3 \in \Xi$,*

$$(\xi_1 \oplus \xi_2) \oplus \xi_3 = \xi_1 \oplus (\xi_2 \oplus \xi_3);$$

4. *Commutativity: for all $\xi_1, \xi_2 \in \Xi$,*

$$\xi_1 \oplus \xi_2 = \xi_2 \oplus \xi_1; \text{ and}$$

5. *Closedness: for all $\xi_1, \xi_2 \in \Xi$, there is an element $\xi_1 \multimap \xi_2 \in \Xi$ satisfying*

$$(\xi_3 \oplus \xi_1) \geq \xi_2 \iff \xi_3 \geq (\xi_1 \multimap \xi_2);$$

call this the hom-element of ξ_1 and ξ_2 .

This primitive asserts that the cost labels can be combined in some way, namely by summing them. As such, the operation \oplus must satisfy a number of properties to ensure that the costs can be combined in a consistent way, and these properties are similar to those we ask of addition of real numbers. These are the first four properties of the operation \oplus . Closedness might seem like an odd property, but it helps us to define the “difference” of two costs. Roughly speaking, $\xi_1 \multimap \xi_2$ is the cost that, when added to ξ_1 , yields ξ_2 . From here, difference might look like:

$$\xi_1 \ominus \xi_2 := \begin{cases} 0_\Xi & \text{if } \xi_2 \geq \xi_1, \\ \xi_2 \multimap \xi_1 & \text{otherwise.} \end{cases}$$

Since \geq is reflexive, we have $\xi \ominus \xi = 0_\Xi$, which confirms our intuitions of subtraction. This will prove useful when we attempt to pin down a metric structure on force configurations, as it will allow us to subtract away unnecessary details to get at the heart of the matter.

We now turn our attention to the existence of special cost labels with respect to the order structure. We have two special needs for our application:

1. we need it to be the case that any subset of Ξ has a greatest lower bound, so that the force-maker can always conceive of some theoretical limiting minimum cost when presented with a set of costs; and
2. we need some special element representing the infinite cost, which we will attach to the impossible process \mathbf{x} .

Let us encode these needs in the following primitive.

2.19 Primitive (Order-Special Cost Labels)

For any set of cost labels $X \subseteq \Xi$, there exists an element $\bigvee X \in \Xi$ satisfying

$$x \geq \bigvee x \text{ for all } x \in X, \text{ and}$$

$$\bigvee x \geq y \quad \text{for all } y \in X \text{ such that } x \geq y \text{ for all } x \in X.$$

Put differently, we say that (Ξ, \geq) has all joins.²¹ We set $\bigvee \Xi \doteq 0_\Xi$.

Moreover, there exists a unique element $\infty \in \Xi$ satisfying $\infty \geq \xi$ for all $\xi \in \Xi$.

This primitive asserts that the cost labels have all joins, which is to say that any subset of cost labels has a greatest lower bound. Had we assumed completeness of \geq , then such an assumption would not be necessary, as the greatest lower bound of any subset would be the infimum of that subset. In the absence of such luxuries, we simply assert that the greatest lower bound exists. We align the order and algebraic structures by sending the void cost 0_Ξ to the bottom of the order structure, so that $\xi \geq 0_\Xi$ for all $\xi \in \Xi$; there may be other costs with this property, but the void cost is order-isomorphic with any such cost. This adds the interpretation that “costs are non-negative,” which is a common property of reasoning in terms of costs. The infinite cost is strictly at the top of the order structure, as $\infty > \xi$ for all $\xi \in \Xi$ other than ∞ itself. Just as we ignored different sorts of impossibility of conversion morphisms—all impossible processes being the same in the abstract—we ignore different sorts of infinite costs. Future work might unpack these different sorts, but for now we are content to have a single

²¹Ordinarily, greatest lower bounds are thought of as meets rather than joins; this is the case when \leq is the primitive, rather than \geq . Here we use \geq to maintain similarity with existing theory on enriched metric spaces, namely that of Lawvere (1973). Having all joins also helps us to properly define the quantale that will ultimately calibrate the costs of conversion.

infinite cost ∞ attached to the impossible process \mathbf{x} . Theoretically, this seems a small price to pay for convenient interpretations and a natural path toward a metric structure on force configurations.

The infinite cost plays a special role in our substantive story, as it represents the cost of the impossible process \mathbf{x} . We observe that it dominates summation, as “regular” infinity does:

2.20 Lemma (Impossibility Absorbs)

For all cost labels $\xi \in \Xi$, we have $\xi \oplus \infty = \infty = \infty \oplus \xi$.

Proof. Choose and fix any $\xi \in \Xi$. We will show that $\xi \oplus \infty \geq \tilde{\xi}$ for all $\tilde{\xi} \in \Xi$; since ∞ is the unique cost label with this property, the argument will suffice. By definition, $\xi \geq 0_\Xi$ and $\infty \geq \tilde{\xi}$; Monotonicity of \oplus delivers $\xi \oplus \infty \geq 0_\Xi \oplus \tilde{\xi} = \tilde{\xi}$, where the final identity obtains from Unitality of \oplus . We infer $\xi \oplus \infty = \infty$; Commutativity of \oplus entails $\xi \oplus \infty = \infty \oplus \xi = \infty$, too, so we are done. ■

Thus, ∞ is the absorbing element of the operation \oplus .

In the spirit of the theory described to this point, we set the cost of the impossible process \mathbf{x} to be infinite and the cost of the identity process to be void, which further links the order and algebraic structures to the abstract theory. Further, we want the addition operation to work smoothly in tandem with composite processes, suggesting that we should pin down how \oplus and \circ interact. We make our special assumptions now.

2.21 Assumption (Special Features)

The map $\text{cost} : \text{Hom}(\mathbb{M}_L^*) \rightarrow \Xi$ satisfies:

1. Costless Identity: for all $\uplus \mathcal{M} \in \mathbb{M}_L^*$, we have $\text{cost}(\text{id}_{\uplus \mathcal{M}}) = 0_{\text{Cost}}$;
2. Impossibility: we have $\text{cost}(\mathbf{x}) = \infty$; and
3. Composability: for all $\uplus \mathcal{M}_1 \xrightarrow{f} \uplus \mathcal{M}_2 \xrightarrow{g} \uplus \mathcal{M}_3$, we have

$$\text{cost}(g \circ f) = \text{cost}(f) \oplus \text{cost}(g).$$

Costless Identity does *not* mean that all maintenance processes are costless—only that the special process $\text{id}_{\uplus \mathcal{M}}$ is costless. Put differently, there may be some force configuration $\uplus \mathcal{M}$ and some maintenance process $f \in \text{Hom}(\uplus \mathcal{M}, \uplus \mathcal{M})$ such that $\text{cost}(f) \neq 0_{\text{Cost}}$.²² Such a process points to costly maintenance, a common feature of the world we inhabit and the tools we use in it. As for Composability, it is a natural assumption to make, as it ensures that the costs of composite processes are the sums of the costs of the component processes. This offers ease of computation and interpretation.

²²Indeed, the impossible process is technically a member of $\text{Hom}(\uplus \mathcal{M}, \uplus \mathcal{M})$, so it is a counterexample to the claim that all maintenance processes are costless.

Thing	General	$\mathbf{Cost}_{\text{Law}}$	\mathbf{Bool}
Labels	Ξ	$\overline{\mathbb{R}}_{\geq 0} = \mathbb{R}_{\geq 0} \cup \{\infty\}$	$\mathbb{B} = \{\text{FALSE}, \text{TRUE}\}$
Order	\geq	$\geq_{\mathbb{R}}$	$\leq_{\mathbb{B}}$ (see text)
Addition	\oplus	$+$	\wedge
Additive Identity	$\mathbb{0}_{\Xi}$	0	TRUE
Hom Object	$x \multimap y$	$\max\{0, y - x\}$	$x \Rightarrow y$ (IF-THEN)
Join	$\bigvee X$	$\inf X$	$\bigwedge X$ (AND)
Infinite Cost	∞	∞	FALSE

Table 1: Lawvere’s (1973) and the Boolean structures.

Though we have taken pains to keep things quite general, two special cases of the cost labels are worth mentioning.²³

1. The first is due to category theorist F. William Lawvere, who used the extended non-negative real numbers $\overline{\mathbb{R}}_{\geq 0} = \mathbb{R}_{\geq 0} \cup \{\infty\}$ to provide a bridge calibration between metric spaces and enriched categories. It allows us to assign a single non-negative real number (or ∞) to each conversion process, mimicking the way we assign a single real number to each point in a metric space. Indeed, Lawvere’s insight was that the structure of the cost labels could be used to calibrate the cost-measurement device, which is the essential idea we seek to generalize here.
2. The second is the Boolean structure, which is a special case of the cost labels where the only costs are FALSE and TRUE; its order is simply $\text{FALSE} \leq_{\mathbb{B}} \text{TRUE}$, plus the usual reflexive statements. Whereas the Lawvere structure allows us to think about costs quite naturally, the Boolean structure is about possibilities; in particular, a process has cost TRUE if it is possible and FALSE if it is impossible. Whereas one combines costs by adding them, one checks possibility paths by taking logical conjunctions—thus, and perhaps unexpectedly, $+$ and \wedge play the same role in these two structures.

Even in examining two simple cases, we see that the cost calibration exercise is crucial in pinning down the nature of the force-maker’s reasoning about costs. There are many such structures, each with its own properties and interpretations, and we have only scratched the surface of what is possible. In asserting Ξ and its properties in the abstract, we assert only that the force-maker abides by *some* set of rules when thinking about costs; save for those, the force-maker is free to think about costs in any way she pleases.

²³This paragraph draws again from the heavenly Fong and Spivak (2019).

Yet again, the atomic theory of force provides a concrete way to understand the abstract structure of cost measurement. Recall that, in the category of force configurations \mathbf{Config}_L , the objects are force configurations (which are a particular kind of graph) and the morphisms are conversion processes (which are a particular kind of graph morphism, namely sequences of elementary edits). We can think of \mathbf{Config}_L as a subcategory of a category of graphs $\mathbf{DirGraph}_L$, where the objects are graphs labeled from L and the morphisms are finite sequences of elementary edits.²⁴ Thus, $\mathbf{DirGraph}_L$ includes all intermediate graphs as objects, whereas \mathbf{Config}_L includes only those graphs that are viable configurations over viable molecules. We introduce $\mathbf{DirGraph}_L$ formally as follows.

2.22 Definition (The Category of Graphs)

The category of graphs $\mathbf{DirGraph}$ has:

1. *Objects the directed graphs with vertices labeled from L ;*
2. *Elementary morphisms the elementary edits, where for each graph G_0 we have*

$$\begin{array}{ll} G_0 \xrightarrow{^0\text{va}_\alpha} G_\alpha, & G_0 \xrightarrow{^0\text{vd}_\beta} G_\beta, \\ G_0 \xrightarrow{^0\text{ea}_\gamma} G_\gamma, & G_0 \xrightarrow{^0\text{ed}_\delta} G_\delta, \end{array}$$

representing the addition and deletion of vertices and edges in G_0 ;

3. *Composition the concatenation of elementary edits; and*
4. *Identity the 0-sequence of elementary morphisms.*

Composition and identity evidently satisfy the usual properties.

Note that, because each pair of elementarily-adjacent graphs is linked via a distinct edit morphism, the morphisms defined by the graphs they link. In other words, not all vertex insertions are the same, nor are all edge deletions, and so on. We will work within the category of graphs $\mathbf{DirGraph}_L$ to understand the structure of cost measurement; we then can apply our understanding to the subcategory of force configurations \mathbf{Config}_L . This is a common strategy

²⁴This is why $\mathbf{DirGraph}_L$ is “a” category of graphs, rather than “the” category of graphs. “The” category of graphs typically takes (undirected) graphs for objects and *graph homomorphisms* for morphisms, which are a different sort of thing. There are stories to tell with this category, but they are not the stories we are telling at present.

in category theory, where one works in a larger category to understand the structure of a smaller one. The larger category provides a broader context in which to understand the smaller one, and the smaller one provides a concrete example of the abstract structures in the larger one.

Now that we have a concrete category in which to work, we can define the cost calibration device. In essence, we do little more than assign each elementary edit morphism a cost label, and then extend this assignment to all morphisms in the category. We define the cost calibration device formally as follows.

2.23 Construction (Cost Functor)

The cost functor is a functor $\text{Cost} : \text{DirGraph}_L \rightarrow \Xi$ that assigns to each elementary edit morphism a cost label. In other words, for all elementary edits

$$G_0 \xrightarrow{f} G_1,$$

we have $\text{cost}(f) \in \Xi$. Composition is additive: if an edit morphism is written as a sequence of elementary edits $f = f_n \circ \dots \circ f_1$, then

$$\text{cost}(f) = \bigoplus_{i=1}^n \text{cost}(f_i).$$

Finally, we have $\text{cost}(\text{id}_G) = \mathbb{0}_\Xi$ for all graphs G .

We have now attached each elementary edit morphism in DirGraph_L to a cost label, and we have extended this assignment to all morphisms in the category by way of \oplus , the means by which we combine costs. Since each elementary edit is defined uniquely by the graphs it links, we now have the possibility of distinct costs for distinct processes.

Suppose, for example, that we had a soldier and intended to train him in the use of a sword and shield. There are several ways to do this; here are two:

1. we could equip him with a sword, train him in its use, then equip him with a shield and train him in its use; or
2. we could equip him with a shield and sword, then train him in the use of the shield, then train him in the use of the sword.

Other routes are possible, but these two are illustrative. The first case involves a sequence of elementary edits: vertex addition (for the sword), edge addition (for the training), vertex addition (for the shield), and edge addition (for the training). The second case involves a different sequence of elementary edits: two rounds of vertex addition (one for the sword, one for the shield) and two rounds of edge addition (one for the sword training, one for the shield training). Though we arrive at the same place in the end, the costs of the two routes may differ: after all, it might be harder to train a soldier in use of a sword when he's holding a shield, or it might be easier to train him in the use of a shield when he's already trained in the use of a sword. Construction 2.23 allows for such distinctions, making for superb flexibility in the force-maker's art.

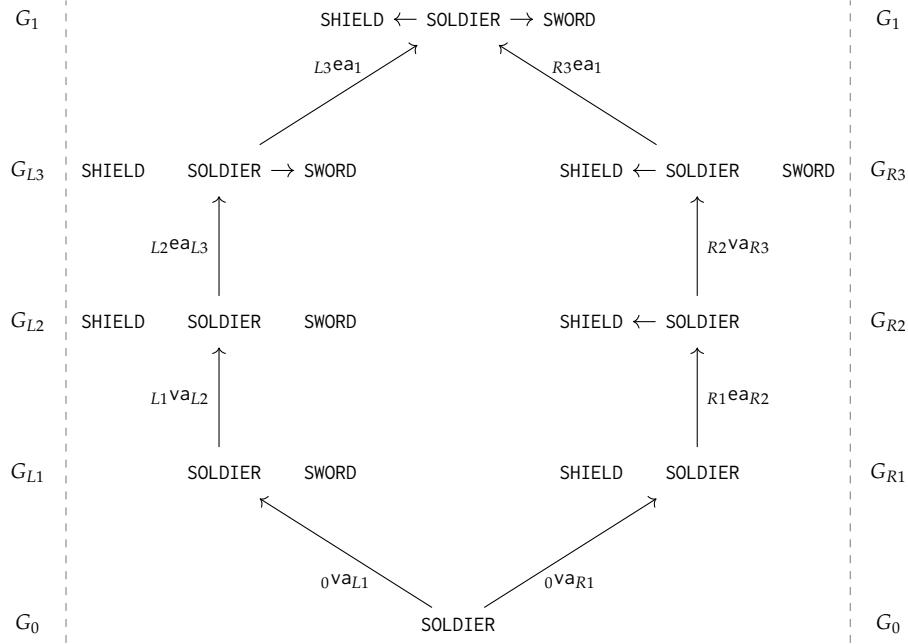


Figure 6: Two paths to training a soldier in the use of a sword and shield. Only the elementary edits are shown, and identity morphisms are omitted for clarity.

Figure 6 depicts these two processes in the category of graphs $\mathbf{DirGraph}_L$. The two processes share common endpoints, but their intermediate steps differ: one involves all equipping then all training, whereas the other involves equipping and training in parallel. The costs of these two processes are not necessarily the same: the left path incurs cost

$$\text{cost} (0va_{L1}) \oplus \text{cost} (L1va_{L2}) \oplus \text{cost} (L2ea_{L3}) \oplus \text{cost} (L3ea_1),$$

whereas the right path incurs cost

$$\text{cost} (0va_{R1}) \oplus \text{cost} (R1ea_{R2}) \oplus \text{cost} (R2va_{R3}) \oplus \text{cost} (R3ea_1).$$

The force-maker can now compare costs to decide which process is more efficient, or she can use the costs to guide her in developing new processes. The properties of the cost labels and the cost calibration device might allow her to manipulate these expressions for further simplification, but for now it suffices to observe potential cost heterogeneity across paths heading to the same destination.²⁵

²⁵This is a good moment to note that the larger category $\mathbf{DirGraph}_L$ is a category of graphs, not a category of force configurations. As such, it is guaranteed to contain all these intermediate graphs, which are not necessarily viable force configurations.

The lattice structure we imposed upon Ξ in Primitive 2.19 facilitates cost-based reasoning by providing a lower bound for any subset of cost labels; in particular, this can help the force-maker navigate the myriad training routes from one configuration to another. For two graphs G_0 and G_1 in $\mathbf{DirGraph}_L$, let $\text{Hom}(G_0, G_1)$ denote the set of all morphisms from G_0 to G_1 . Then the image of cost on $\text{Hom}(G_0, G_1)$ is a subset of Ξ , and as such it has a greatest lower bound—*i.e.*, we can approximate the cheapest way to convert G_0 into G_1 by

$$\bigvee \text{cost}(\text{Hom}(G_0, G_1)).$$

There are not necessarily any morphisms in $\text{Hom}(G_0, G_1)$ that achieve this cost, but the force-maker can use this lower bound to guide her in developing new processes. Were we to equip Ξ with further structure—say, density or a metric—then the force-maker could use these tools to further refine her reasoning about costs, and indeed she might be able to achieve the cost of the cheapest process exactly. We will leave these refinements for future work, content for now to have provided the force-maker with a tool to reason about costs in a structured way.²⁶ We will see that the next development in the theory provides an alternative way to reason about costs, one that is more closely tied to the compositional nature of the force configurations themselves, not to mention a weak notion of rationality on the part of the force-maker.

Summary of Section 2.3. In this subsection, we introduced the cost labels Ξ and discussed what structure it must possess to calibrate the costs of conversion processes. These properties include a weak order structure, a binary operation denoting the combination of costs, and two special cost labels: the void cost $\mathbb{0}_\Xi$ and the infinite cost ∞ . Most importantly, Ξ is the appropriate target for the cost measurement device, which assigns to each conversion process a cost label. In the context of the atomic theory of force, Ξ provides the codomain for elementary edit cost maps in the category of graphs $\mathbf{DirGraph}_L$, which makes cost measurement a rather straightforward affair.

²⁶Frankly, your humble author is skeptical about the tenability of an approach based on choosing the cheapest path between any two force configurations. For starters, it is not obvious that this sort of local cost-minimization is on the mind of real-world force-makers, who often engage in locally-costly processes to achieve globally-optimal results. Moreover, this implies a strong form of rationality on the part of the force-maker, who must always choose the cheapest path across a wide variety of training and production processes. This seems a tall order, and it is not clear that the force-maker is up to the task; this is precisely why we equip her with a weaker form of rationality in the next section. This decision has been quite difficult for your humble author, who has spent many hours pondering the nature of the force-maker’s rationality. The tack taken in the next section is the result of much thought and reflection, and it is hoped that the reader will find it as compelling as your humble author does. But, as always, tastes vary on such matters.

2.4 The First Rationality

We have developed a theory wherein we can gather molecules into configurations and convert these configurations into new configurations via costly processes. These processes are taken to be just as primitive as the configurations themselves, arriving from the sky as a gift from the gods. There are no guarantees that a given force-maker knows how to turn three hundred hopelites into a phalanx, nor now to turn enriched uranium into a nuclear weapon. Put differently, we think of each $\text{Hom}(\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_2)$ as a concrete set of real-world possibilities, rather than a force-maker's current knowledge or practices. This helps stiff-arm questions the reader might have had along the way, such as "how does the force-maker—or the author—know how to convert one configuration into another?" The answer is simple: she doesn't, and neither does your humble author.

But the force-maker is not without her wits. Though she is not a master engineer, she is capable of reasoning about costs and making decisions accordingly. By now we know that she can compare various costs, compute their sums, and consider the residual between two costs. She has a notion of impossibility, which she can use to rule out certain processes, and she has a notion of identity, which she can use to maintain the status quo. She can reason about costs in a structured way, and she can use this reasoning to guide her in developing new processes. And yet, in the light of the humiliating complexity of the problem at hand, it would appear unwise to expect her to behave as a perfectly rational agent (and the same goes for your humble author, who clearly is not one, either).

We will now introduce a new primitive, the *choice schedule*, via which the force-maker selects a process from each set of possibilities.

2.24 Primitive (Choice Schedule)

There is a selection of processes

$$\begin{aligned} \mathbf{CS} : \mathbb{M}_L^* \times \mathbb{M}_L^* &\longrightarrow \text{Hom}(\mathbb{M}_L^*), \\ (\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_2) &\longmapsto \mathbf{CS}(\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_2) \in \text{Hom}(\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_2), \end{aligned}$$

representing the force-maker's choice of a process from each set of possibilities.

The choice schedule is the force-maker's tool for selecting a process from each set of possibilities. We do not specify how the force-maker makes her selections, only that she can make them. Indeed, nothing precludes the force-maker from shirking her duties, setting $\mathbf{CS}(\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_2) = \mathbf{x}$ for all pairs of configurations $\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_2 \in \mathbb{M}_L^*$, declaring all conversions impossible.

Lurking behind the simple definition in Primitive 2.24 lies a deep assumption worth unpacking. We have asserted that she is capable of choosing a process from each set of possibilities, but we have not said how she makes these choices. This is intentional, as it would take more pages than we have to describe how one might choose training processes for hoplites, enrichment processes for uranium, and so on. But precisely because of this lack of structure, it is not immediately obvious just what sort of unifying principle exists to allow us to define the choice schedule in the first place.

Now, one thing all processes have in common is that they are sent to the cost structure \mathfrak{E} via the cost measurement functor cost . It is tempting to use this structure to define the choice schedule, using the unified space of costs. And indeed, this is one tack we could have taken—say,

“the force-maker selects the process with the lowest cost.”

There are two damning problems with this approach:

1. Substantively, this amounts to a very strong form of rationality, requiring the force-maker to always choose the cheapest process. But she might lack the knowledge to do so, or she might have other reasons for choosing a more expensive process—say, because a cheap process is too risky, or because a more expensive process sends contracts to the right defense contractors, *or whatever*. It could be *anything*, and we have no way to know.
2. Formally, this approach would require us knowing that each of the infinite cost minimization problems are well-posed, which is a tall order. One usually makes such guarantees by appealing to some kind of compactness property on the set of options and some kind of continuity property on the cost function, but we have not topologized the set of options and we have not endowed the cost function with any kind of structure. We could certainly do so, and perhaps we will in future work—the graph structure of the force configurations makes this easy to sketch out—but the substantive point raised above is enough of a headwind to make us think twice about this approach. Things are hard enough as it is.

Owing to these problems, we have left the terms of the choice schedule slack, allowing the force-maker to choose as she pleases. But precisely because of the slackness, we find ourselves in a situation where we must choose one option from an infinite number of menus via an unspecified process. This is, of course, a famous philosophical problem blunted by most applied scholars, including this one, via the (in)famous *Axiom of Choice*. That axiom asserts that, given a collection of nonempty sets, there is a way to choose one element from each set, and it is a foundational principle in mathematics.

Highfalutin disclaimers aside, we arrive at our rationality assumption.

2.25 Assumption (Compositional Awareness)

The cost functor cost and the choice schedule \mathbf{CS} satisfy

$$\bigoplus_{i=1}^n (\text{cost} \circ \mathbf{CS})\left(\biguplus \mathcal{M}_{i-1}, \biguplus \mathcal{M}_i\right) \geq (\text{cost} \circ \mathbf{CS})\left(\biguplus \underline{\mathcal{M}}, \biguplus \overline{\mathcal{M}}\right),$$

for all pairs of configurations $\biguplus \underline{\mathcal{M}}, \biguplus \overline{\mathcal{M}} \in \mathbf{M}_L^*$ and all chained selections

$$\biguplus \underline{\mathcal{M}} = \biguplus \mathcal{M}_0 \xrightarrow{\mathbf{CS}(\biguplus \mathcal{M}_0, \biguplus \mathcal{M}_1)} \dots \xrightarrow{\mathbf{CS}(\biguplus \mathcal{M}_{n-1}, \biguplus \mathcal{M}_n)} \biguplus \mathcal{M}_n = \biguplus \overline{\mathcal{M}}.$$

We call this the Compositional Awareness of the force-maker.

In words, two things are happening here.

1. First, the force-maker understands that processes can be put into a sequence. If she knows she can get from $\biguplus \mathcal{M}_0$ to $\biguplus \mathcal{M}_1$ via f and from $\biguplus \mathcal{M}_1$ to $\biguplus \mathcal{M}_2$ via g , then she knows she can get from $\biguplus \mathcal{M}_0$ to $\biguplus \mathcal{M}_2$ via $g \circ f$. If she is aware enough of f and g to choose them from their respective menus, then she is aware enough of $g \circ f$ to choose it from its menu. This seems a perfectly reasonable assumption, as many officers with master's degrees from specialized war colleges can attest; their second program began *after* their first program ended.
2. The real “rationality” bet is the second part of the assumption, which asserts that the force-maker is thinking in terms of costs. The assumption entails that the force-maker would never choose a path from $\biguplus \mathcal{M}_0$ to $\biguplus \mathcal{M}_n$ that is more expensive than the cheapest path of which she is aware. If she is aware enough of f_1, \dots, f_n to choose them from their respective menus, then the cost of $f_n \circ \dots \circ f_1$ sets an upper bound on the cost of the path she chooses. In this sense, the force-maker is rational—consistent in some sense with the costs of conversion. To repeat, none of the respective f_i need be optimal in its locale.

The force-maker’s rationality is thus rather weak, more about structural awareness than about the strict optimization motivating goofy *homo economicus* caricatures. Bounded rationality is still rationality, and these assumptions, while modest, remain capable of question. Yet, despite their limitations, they provide a framework capable of delivering results sufficiently robust as to justify their inclusion. This allows the force-maker to navigate complexity with enough strategic foresight, even if not perfectly optimal, while ensuring that her decisions reflect an understanding of costs and composition over time.

For example, suppose we have three force configurations— $\biguplus \mathcal{M}_1$, $\biguplus \mathcal{M}_2$, and $\biguplus \mathcal{M}_3$ —and the conversion morphisms listed in Table 2.²⁷ Two f s convert

Processes				Composites			
Process	Source	Target	cost	Process	Source	Target	cost
f_1	$\biguplus \mathcal{M}_1$	$\biguplus \mathcal{M}_2$	1	$g_1 \circ f_1$	$\biguplus \mathcal{M}_1$	$\biguplus \mathcal{M}_3$	6
f_2	$\biguplus \mathcal{M}_1$	$\biguplus \mathcal{M}_2$	3	$g_1 \circ f_2$	$\biguplus \mathcal{M}_1$	$\biguplus \mathcal{M}_3$	8
g_1	$\biguplus \mathcal{M}_2$	$\biguplus \mathcal{M}_3$	5	$g_2 \circ f_1$	$\biguplus \mathcal{M}_1$	$\biguplus \mathcal{M}_3$	3
g_2	$\biguplus \mathcal{M}_2$	$\biguplus \mathcal{M}_3$	2	$g_2 \circ f_2$	$\biguplus \mathcal{M}_1$	$\biguplus \mathcal{M}_3$	5
h_1	$\biguplus \mathcal{M}_1$	$\biguplus \mathcal{M}_3$	5				
h_2	$\biguplus \mathcal{M}_1$	$\biguplus \mathcal{M}_3$	6				

Table 2: A set of processes and their costs, structured in $\Xi = \overline{\mathbb{R}}_{\geq 0}$.

$\biguplus \mathcal{M}_1$ into $\biguplus \mathcal{M}_2$, and two g s convert $\biguplus \mathcal{M}_2$ into $\biguplus \mathcal{M}_3$; the resulting four composites convert $\biguplus \mathcal{M}_1$ into $\biguplus \mathcal{M}_3$, along with two direct h s. We do not require that $\mathbf{CS}(\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_2) = f_1$, even though f_1 is the cheapest process linking $\biguplus \mathcal{M}_1$ and $\biguplus \mathcal{M}_2$; likewise, we do not need $\mathbf{CS}(\biguplus \mathcal{M}_2, \biguplus \mathcal{M}_3) = g_2$.

What we *do* require is that the choice schedule is consistent with choices with respect to costs. Table 3 lists the allowable combinations of selections from

$\mathbf{CS}(\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_2)$	$\mathbf{CS}(\biguplus \mathcal{M}_2, \biguplus \mathcal{M}_3)$	\oplus	$\mathbf{CS}(\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_3)$
f_1	g_1	6	$g_1 \circ f_1$ or h_1 or h_2
f_1	g_2	3	$g_2 \circ f_1$
f_2	g_1	8	$g_1 \circ f_2$ or h_1 or h_2
f_2	g_2	5	$g_2 \circ f_2$ or h_1

Table 3: Allowable selections for the choice schedule. The first two columns list the selected direct processes from $\biguplus \mathcal{M}_1$ to $\biguplus \mathcal{M}_2$ and from $\biguplus \mathcal{M}_2$ to $\biguplus \mathcal{M}_3$. The fourth column tells us which processes from $\biguplus \mathcal{M}_1$ to $\biguplus \mathcal{M}_3$ satisfy Compositional Awareness, given the choices in the first two columns.

$\text{Hom}(\biguplus \mathcal{M}_1)$ to $\text{Hom}(\biguplus \mathcal{M}_3)$ as a function of the selections from $\text{Hom}(\biguplus \mathcal{M}_1)$ to $\text{Hom}(\biguplus \mathcal{M}_2)$ and from $\text{Hom}(\biguplus \mathcal{M}_2)$ to $\text{Hom}(\biguplus \mathcal{M}_3)$. The example makes clear that the first two choices constrain the third choice by ensuring that the selected process from $\biguplus \mathcal{M}_1$ to $\biguplus \mathcal{M}_3$ is at least as cheap as the sum of the costs of the selected processes from $\biguplus \mathcal{M}_1$ to $\biguplus \mathcal{M}_2$ and from $\biguplus \mathcal{M}_2$ to $\biguplus \mathcal{M}_3$.

²⁷Here we ignore identity morphisms, which we will discuss more in due course.

The selection of processes allows us to tell stories about the evolution of the force-maker’s practices. Consider, for example, the process of constructing an aircraft carrier. The first aircraft carriers were not custom-built; they were converted from existing ships. For example, the British aircraft carrier H.M.S. *Argus*, deployed in 1918, was constructed from the hull of a partially-built Italian liner called the *Conte Rosso* with the addition of a flight deck procured from a *Blackburne*-class reconnaissance ship.²⁸ This process was expensive, requiring two ships and their subsequent fusion, but it was the only way to get an aircraft carrier at the time. We can think of f as the process of converting raw materials to a liner and a reconnaissance ship and g as the process of converting a liner and a reconnaissance ship to an aircraft carrier. Carriers were originally made via process $g \circ f$, but as the technology of carrier construction advanced, it became possible to construct carriers directly, without the need for the intermediate step. Such hs were cheaper than $g \circ f$, and so the choice schedule shifted over time; the first British custom-built carrier, H.M.S. *Hermes*, was deployed in 1923.

The choice schedule requires the choice of a morphism for each pair of configurations, including the pair of the same configuration. Among the available options lives the identity morphism, which carries zero cost. But, some configurations might require maintenance, and the cost of maintaining a configuration in its current state might not be zero. However, Compositional Awareness imposes a requirement on maintenance: the cost of maintaining a configuration in its current state must be less than the sum of converting it to *anything else* and then converting it back! For any fixed configuration $\biguplus \mathcal{M} \in \mathbb{M}_L^*$, define

$$\leftrightarrow (\biguplus \mathcal{M}') := (\text{cost} \circ \text{CS})(\biguplus \mathcal{M}, \biguplus \mathcal{M}') \oplus (\text{cost} \circ \text{CS})(\biguplus \mathcal{M}', \biguplus \mathcal{M}),$$

which is the cost of converting $\biguplus \mathcal{M}$ to $\biguplus \mathcal{M}'$ and then back to $\biguplus \mathcal{M}$. Compositional Awareness takes the form

$$\leftrightarrow (\biguplus \mathcal{M}') \geq (\text{cost} \circ \text{CS})(\biguplus \mathcal{M}, \biguplus \mathcal{M}) \text{ for all } \biguplus \mathcal{M}' \in \mathbb{M}_L^*.$$

The tenability of this requirement depends on the time horizon we wish to impose on the force-maker’s practices; the shorter the horizon, the more tenable the requirement. Shorter time horizons make for smaller costs, and fewer available conversions. The former suggests that the “always select the identity” strategy is a limiting, instantaneous case of the choice schedule, and we allow ourselves this interpretation as a maintenance-free baseline of theoretical import. This latter point suggests that many $\leftrightarrow (\biguplus \mathcal{M}') = \infty$, as there might be no way to convert a configuration into another and back again.

²⁸Details are drawn from Norman Polmar’s adorable *Aircraft Carriers* (2006), a comprehensive history of the aircraft carrier and veritable treasure trove of information.

Yet another interesting wrinkle in our structure is what happens once the force-maker deems a given conversion impossible. By construction, each hom set includes the impossible process \mathbf{x} , and the Dominance of Impossibility (Primitive 2.10) ensures that any chain one could have built with the impossible process is itself impossible. But strictly speaking, Dominance of Impossibility is not required for this to hold. Since the cost ∞ —the cost of the impossible process \mathbf{x} —is absorbing with respect to \oplus (Lemma 2.20), the cost of any chain that includes the impossible process is infinite. Since $\infty \geq \xi$ for any $\xi \in \Xi$, the cost of any chain that includes the impossible process is at least as great as the cost of any other chain. Thus, Compositional Awareness binds not at all with respect to processes including at least one link the force-maker deems impossible. Thus, we have a consistency not just in the processes the force-maker truly chooses, but also in the conversions she deems impossible.²⁹ Since ∞ may be assigned to any process—not just the impossible process—the force-maker may use it to indicate that a process is too costly, too risky, or otherwise undesirable, similarly disconnecting it from the rest of the space of configurations. Thus, literal impossibility is not required for the force-maker to treat a process as impossible, and the force-maker’s practices are consistent with her beliefs about the costs of conversion. This is what we mean by Compositional Awareness.

Summary of Section 2.4. In this subsection, we introduced the choice schedule **CS**, which encodes the force-maker’s choice of a process from each set of possibilities. We then introduced the Compositional Awareness assumption, which requires that the force-maker’s choices are consistent with the costs of conversion. In a sense, the assumption brings all of our implicit notions of how the force-maker navigates the space of configurations into the light, providing a framework for reasoning about her practices. We also discussed the implications of the choice schedule for the force-maker’s practices, including the maintenance of configurations and the treatment of impossible processes. The force-maker’s practices are consistent with her beliefs about the costs of conversion, and the choice schedule provides a simple space for structured reasoning. In a sense, then, the assumption is painfully obvious either in its acceptability or its rejection. What might not be so obvious, however, is that this mild behavioral postulate is enough to deliver an even deeper structure for the space of force configurations.

²⁹It should be noted that Dominance of Impossibility is useful for other reasons. For example, since $\infty \in \text{Hom}(\biguplus \mathcal{M}, \biguplus \mathcal{M})$ for all $\biguplus \mathcal{M} \in \mathbb{M}_L^*$, composability would require us to consider infinite chains of impossible processes $\mathbf{x} \circ \mathbf{x} \circ \dots$. Dominance of Impossibility ensures that these chains evaluate to \mathbf{x} , which is a reasonable result. We ultimately will be using infinite costs more than we will be using impossible processes, so the Dominance of Impossibility is less important than the absorbing property of ∞ .

2.5 The Space Between Configurations

We have developed a theory wherein we can gather molecules into configurations, convert these configurations into new configurations via costly processes, and select processes that convert one configuration into another. It might not be obvious, but our construction is amenable to distance-like measures between configurations. This introduces an interesting question: what makes distance distance? Our answer is inspired by the pathbreaking work of category theorist F. William Lawvere (1973), who helped distill the essence of distance to its simplest form in a decidedly beautiful way.

What is distance? In the most general sense, distance is a measure of separation between two points: two points that are not very separated have a small distance between them, while two points that are very separated have a large distance between them. This is a simple idea, and our intuitions of how to execute such measurement are quite strong. In the Euclidean plane, for example, we can measure the distance between two points by drawing a straight line between them and assigning a non-negative real number representing the length of that line. This is the most common way to measure distance, but it is not the only way. Once one learns the Pythagorean Theorem, for example, one can measure distance in terms of the lengths of the sides of a right triangle. This is a different way to measure distance, but it is still a way to measure distance all the same. And indeed, in many contexts there exist many ways to measure distance, each with its own strengths and weaknesses. Thus, it is worth thinking about what distance is in traditional terms.

2.26 Definition (Metric Space, Traditional Style)

A traditional metric space is a pair (X, d) , where X is a set and $d : X \times X \rightarrow \mathbb{R}_{\geq 0}$ is a function satisfying the following properties:

1. Traditional Codomain: $d(x, y) \in \mathbb{R}_{\geq 0}$ for all $x, y \in X$;
2. Identity of Discernibles: $d(x, x) = 0$ for all $x \in X$;
3. Discrimination: $x \neq y$ implies $d(x, y) > 0$ for all $x, y \in X$;
4. Symmetry: $d(x, y) = d(y, x)$ for all $x, y \in X$; and
5. the Triangle Inequality: $d(x, y) + d(y, z) \geq d(x, z)$ for all $x, y, z \in X$.

Each property captures an essential feature of distance, and all of them appeal to the intuitions we gain from real-world experience and elementary geometry.³⁰

³⁰ Classical as these properties are, their formal introduction dates back only to Maurice Fréchet's classic "Sur Quelques Points du Calcul Fonctionnel" (1906).

We discuss the properties just enumerated in turn.

1. *Traditional Codomain*: $d(x, y) \in \mathbb{R}_{\geq 0}$ for all $x, y \in X$. This property simply asserts that distance is something measured in non-negative real numbers. It calibrates the ruler just the same way that Ξ calibrates the device that measures costs in the context of force configurations. Notice its inability to capture the idea of infinite distance, which has been key in our theory.
2. *Identity of Discernibles*: $d(x, x) = 0$ for all $x \in X$. This property asserts that the distance between a point and itself is zero. It is a natural property that carries a few different meanings. One way to see it is that the system is static: each location is a fixed point in space, and the distance between a point and itself is zero because the point is already there. Another is that the points are small, and the distance between a point and itself is zero because the point is so small that it has no size.
3. *Discrimination*: $x \neq y$ implies $d(x, y) > 0$ for all $x, y \in X$. This property asserts that the distance between two different points is strictly positive. It captures the idea that distance is a measure of separation, but it is restrictive in the sense that the same point cannot be in two places at once.
4. *Symmetry*: $d(x, y) = d(y, x)$ for all $x, y \in X$. This property asserts that the distance between two points is the same regardless of the order in which they are considered. In the context of a resource theory, the costs of conversion might be different depending on the direction of the conversion: assembly and disassembly are not necessarily the same, much like running up- and downhill are not the same.
5. the *Triangle Inequality*: $d(x, y) + d(y, z) \geq d(x, z)$ for all $x, y, z \in X$. For better or worse, the Triangle Inequality is the most interesting and essential feature of a metric space. It encodes, and significantly expands, the maxim that the shortest distance between two points is a straight line. This is just the sort of homespun wisdom that points toward essential truths, and its logic dates back (at least) to the famous Proposition 20 from Book I of Euclid's *Elements* (300 B.C.):

In any triangle two sides taken together in any manner are greater than the remaining one.

This property is a powerful one, and it is the one that most clearly distinguishes a metric space from a mere set.

These five properties are the essence of introductory notions of distance, reinforced by centuries of mathematical practice and our real-world experiences.

Lawvere (1973) takes issue with three of the properties just discussed:

None of our results in this paper will depend on the additional Frechet axioms:

$$\begin{aligned} &\text{if } X(a, b) = 0 \text{ then } a = b \\ &X(a, b) < \infty \\ &X(a, b) = X(b, a). \end{aligned}$$

The first of these is not very natural from the categorical viewpoint since it corresponds to requiring that isomorphic objects are equal [...]. Allowing ∞ among the quantities is precisely analogous to including the empty set among abstract sets, and it is done for similar reasons of completeness [...]. The non symmetry is the more serious generalization, and moreover occurs in many naturally arising examples, such as $X(a, b) = \infty$ work required to get from a to b in mountainous region X (p. 138).

Lawvere's critique is that the Identity of Discernibles and Discrimination properties are not necessary for a distance-like measure, and that the Symmetry property is not necessary for a distance-like measure to be useful. In fact, he argues that the Symmetry property is not even necessary for a distance-like measure to be natural, as it is violated in many natural examples. Accordingly, we arrive at a more general definition of a metric space, with the construction now being called a *Lawvere metric space*.³¹

2.27 Definition (Metric Space, Lawvere Style)

A Lawvere metric space is a pair (X, d) , where X is a set and $d : X \times X \rightarrow \overline{\mathbb{R}}_{\geq 0}$ is a function satisfying the following properties:

1. Extended Codomain: $d(x, y) \in \overline{\mathbb{R}}_{\geq 0}$ for all $x, y \in X$;
2. Identity of Discernibles: $d(x, x) = 0$ for all $x \in X$; and
3. the Triangle Inequality: $d(x, y) + d(y, z) \geq d(x, z)$ for all $x, y, z \in X$.

What has been added is the possibility of infinite distance; what's been removed are the maxims that points are *only* where they are and that moving backwards is the same as moving forwards; and what remains are the maxims that points are where they are and that the shortest distance between two points is a straight line. This is the essence of real-valued distance in its most general form.

³¹One sometimes hears "extended pseudoquasimetric space," where "extended" points to the inclusion of ∞ , "pseudo-" to the relaxation of Discrimination, and "quasi-" to the relaxation of Symmetry. It simply sounds less elegant than "Lawvere metric space."

It might be obvious to the reader that the idea will now be to generalize Lawvere’s notion of distance to our context. Lawvere’s observations seem sound when applied here, as (1) we do not wish to impose equality upon isomorphic configurations; (2) we have need for infinite distance to stand in for impossible conversions; and (3) the costs of conversion might be different depending on the direction of the conversion. As such, something like a Lawvere metric space seems a natural fit for our context. And yet, it seems unlikely that the force-maker thinks about distance-as-cost in unidimensional, real-valued terms. The processes at hand are too intricate, the span of configurations too vast, and the costs of conversion too varied for the force-maker to think in terms of a simple, real-valued distance. But, per our discussion in Section 2.3, the force-maker might also ask questions where distance is far simpler than what the real numbers offer, such as when she asks whether a given conversion is possible or not. We therefore work with Ξ as our codomain, and we generalize the notion of distance to our context.

2.28 Definition (Metric Space, Weak Style)

A weak metric space calibrated by Ξ is a pair (X, d) , where X is a set and $d : X \times X \rightarrow \Xi$ is a function satisfying the following properties:

1. Weak Codomain: $d(x, y) \in \Xi$ for all $x, y \in X$, where Ξ satisfies the properties discussed in Section 2.3;
2. Identity of Discernibles: $d(x, x) = 0_{\Xi}$ for all $x \in X$; and
3. the Triangle Inequality: $d(x, y) \oplus d(y, z) \geq d(x, z)$ for all $x, y, z \in X$.

We retain the general idea of Lawvere’s metric space, but we replace the real numbers with Ξ ; this also involves replacing the real 0 with 0_{Ξ} , the real + with \oplus , and the real \geq with \geq .

At stake at the moment is whether the force-maker’s practices are consistent with a notion of distance. Nearly any data structure we might think of her using to reason about the space of configurations will have some notion of distance, and it is this notion of distance that we are trying to capture. In formal terms, we are asking whether the force-maker as we have structured and motivated her can envision the space of configurations as a weak metric space calibrated by Ξ . Since Ξ has been left open to interpretation, this question is not as restrictive as it might seem: we merely need her to be calibrated by *some* Ξ that satisfies the properties discussed in Section 2.3. Put differently, we do not wonder whether the force-maker abides by some rule but rather by some class of rules, and we are interested in the structure of the space of configurations that emerges from her practices conditional on the rules she abides by.

Let us define our candidate weak distance function.

2.29 Construction (Candidate for Weak Distance)

We define the function

$$d : \mathbb{M}_L^* \times \mathbb{M}_L^* \longrightarrow \Xi, \\ (\bigsqcup \mathcal{M}_1, \bigsqcup \mathcal{M}_2) \longmapsto \xi(\bigsqcup \mathcal{M}_1, \bigsqcup \mathcal{M}_2) \ominus \xi(\bigsqcup \mathcal{M}_1, \bigsqcup \mathcal{M}_1),$$

where \ominus is the subtraction operation in Ξ and $\xi := \text{cost} \circ \text{CS} : \mathbb{M}_L^* \times \mathbb{M}_L^* \rightarrow \Xi$.

The candidate distance function d is the “difference” between the costs of converting $\bigsqcup \mathcal{M}_1$ into $\bigsqcup \mathcal{M}_2$ and converting $\bigsqcup \mathcal{M}_1$ into itself, where “difference” has been put in scare quotes because it is not a subtraction in the traditional sense.³² It still serves its purpose by measuring the separation between two configurations in terms of the costs of conversion. This is an important quantity: the gap between the costs of turning $\bigsqcup \mathcal{M}_1$ into $\bigsqcup \mathcal{M}_2$ and the costs of maintaining $\bigsqcup \mathcal{M}_1$ in its current state. This subtraction has two important effects:

1. It zeroes out the cost of maintaining a configuration in its current state, as the distance between any configuration and itself will be 0_{Ξ} .
2. The more nuanced effect is relational: the distance between two configurations contains information about the relationship between the cost of converting one configuration into another and the cost of maintaining the first configuration in its current state. If the initial state has high “inertia”—i.e., low maintenance costs—then the distance between it and other configurations will be characterized primarily by the costs of conversion. If, on the other hand, the initial state has low inertia—i.e., high maintenance costs—then the distance between it and other configurations will be smaller, as the costs of maintaining the initial state will be subtracted from the costs of conversion. If the costs of conversion are small compared to the maintenance costs, then the target configuration “attracts” the initial configuration, and the distance between them will be small.

Thus, the candidate distance function allows for a dynamic approach to distance.

³²Recall that we defined

$$\xi_1 \ominus \xi_2 := \begin{cases} 0_{\Xi} & \text{if } \xi_2 \geq \xi_1, \\ \xi_2 \multimap \xi_1 & \text{otherwise,} \end{cases}$$

where \multimap is the hom object in the quantale Ξ —i.e., $\xi_2 \multimap \xi_1$ satisfies

$$(\xi_3 \oplus \xi_2) \geq \xi_1 \iff \xi_3 \geq (\xi_2 \multimap \xi_1) \text{ for all } \xi_3 \in \Xi.$$

These dynamics play a key role in shaping the space of configurations. As such, let us begin from the stylized case where there are no maintenance costs.

2.30 Definition (The Static Cases)

We say a choice schedule \mathbf{CS} is static just in case

$$(\text{cost} \circ \mathbf{CS})(\biguplus \mathcal{M}, \biguplus \mathcal{M}) = \mathbb{0}_{\Xi} \text{ for all } \biguplus \mathcal{M} \in \mathbb{M}_L^{\star}.$$

If $\mathbf{CS}(\biguplus \mathcal{M}, \biguplus \mathcal{M}) = \text{id}_{\biguplus \mathcal{M}}$ for all $\biguplus \mathcal{M} \in \mathbb{M}_L^{\star}$, we say \mathbf{CS} is identically static.

The identically static case is the limiting, instantaneous case of the choice schedule, where the identity morphism is always selected. It means that that force-maker's way of thinking about the space of configurations is to treat each configuration as a fixed point in space. This is the most intuitive way to think about force configurations in the context of what makes for a strong force.

The static case makes our lives quite easy.

2.31 Proposition (The Static Case)

If \mathbf{CS} is static, then $(\mathbb{M}_L^{\star}, d)$ is a weak metric space calibrated by Ξ .

Proof. Weak Codomain is immediate from the definition of d , and Identity of Discernibles is immediate because \mathbf{CS} is static. We therefore focus only on the Triangle Inequality. Choose any three configurations $\biguplus \mathcal{M}_1$, $\biguplus \mathcal{M}_2$, and $\biguplus \mathcal{M}_3$ in \mathbb{M}_L^{\star} ; to ease notation, write

$$\xi_{ij} := \xi(\biguplus \mathcal{M}_i, \biguplus \mathcal{M}_j) \text{ for all } i, j \in \{1, 2, 3\}.$$

The Triangle Inequality may be written

$$(\xi_{12} \ominus \xi_{11}) \oplus (\xi_{23} \ominus \xi_{22}) \geq \xi_{13} \ominus \xi_{11}.$$

Since \mathbf{CS} is static, $\xi_{ii} = \mathbb{0}_{\Xi}$ for all $i \in \{1, 2, 3\}$, so we may simplify the above to

$$\xi_{12} \oplus \xi_{23} \geq \xi_{13}.$$

This is Compositional Awareness, and so the Triangle Inequality holds. \blacksquare

Remarkably, the static case demonstrates that Compositional Awareness is *precisely* the correct rationality postulate to impose on the force-maker's practices: it is fully equivalent to the Triangle Inequality, the most interesting and essential feature of a metric space. Thus, minimally-consistent reasoning in terms of static cost gives rise to a notion of distance *and vice versa*. This is a powerful result, suggesting that the force-maker's practices are consistent with a notion of distance, and that the space she envisions is structured in a way that is amenable to reasoning. Such unexpected joys are the fruits of our labor.

The more speculative case is when the force-maker's practices are not static. Recall that we can write the Triangle Inequality in compact notation as

$$(\xi_{12} \ominus \xi_{11}) \oplus (\xi_{23} \ominus \xi_{22}) \geq \xi_{13} \ominus \xi_{11}.$$

This always holds in case the right-hand side is 0_{Ξ} —i.e., when $\xi_{11} \geq \xi_{13}$. So to make things difficult, suppose that the right-hand side is not 0_{Ξ} , yielding

$$\begin{aligned} & (\xi_{12} \ominus \xi_{11}) \oplus (\xi_{23} \ominus \xi_{22}) \geq \xi_{11} \multimap \xi_{13}, \\ \iff & (\xi_{12} \ominus \xi_{11}) \oplus (\xi_{23} \ominus \xi_{22}) \oplus \xi_{11} \geq \xi_{13}. \end{aligned}$$

We have a triangle inequality that is not quite a triangle inequality: the two “legs” of the triangle are residuals of the costs of conversion, and the “hypotenuse” is the cost of conversion. One of the legs is “de-residualized” by adding its maintenance costs back—again, recall that this is not a true difference but rather a residual in Ξ —but the other leg is not. Put differently, ξ_{11} is put back into the mix, but ξ_{22} is not. This sets an upper bound on how large ξ_{22} can be, as it could well be that we subtract away so much from ξ_{23} that the Triangle Inequality is violated, even when Compositional Awareness holds. In other words, if $\uplus \mathcal{M}_2$ is too costly to maintain that it naturally wants to decay into $\uplus \mathcal{M}_3$, then the Triangle Inequality might be violated. We might call this a *slingshot effect*, baked into the DNA of the candidate distance function.

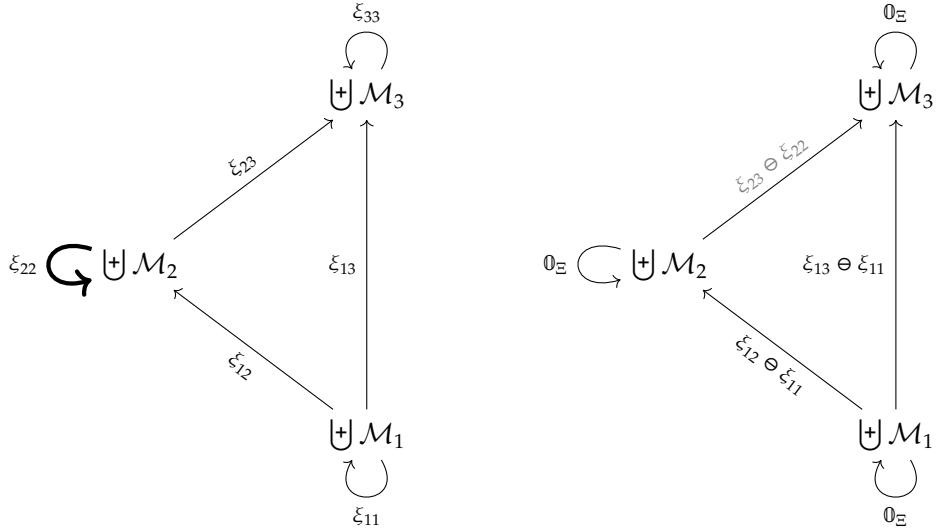


Figure 7: The slingshot effect: the Triangle Inequality is violated when the costs of maintaining $\uplus \mathcal{M}_2$ are so high that it naturally wants to decay into $\uplus \mathcal{M}_3$.

To the degree that the force-maker is thinking dynamically, the slingshot effect is a feature, not a bug. Nevertheless, it is a feature best left for another day, and we turn our attention to determining how large it is allowed to be without violating the Triangle Inequality. Notice that this requires gradually strengthening Compositional Awareness to begin to account for maintenance costs and associated slingshot effects, suggesting that the force-maker must think a little bit harder about the costs of conversion to allow for distance-oriented dynamic reasoning.

2.32 Definition (Dynamic Awareness)

We say a choice schedule \mathbf{CS} satisfies Dynamic Awareness just in case

$$(\xi_{12} \ominus \xi_{11}) \oplus (\xi_{23} \ominus \xi_{22}) \oplus \xi_{11} \geq \xi_{13}$$

for all $\biguplus \mathcal{M}_1$, $\biguplus \mathcal{M}_2$, and $\biguplus \mathcal{M}_3$ in \mathbb{M}_L^ .*

Dynamic Awareness is the rationality postulate that ensures that the slingshot effect is not too large. It is stronger than Compositional Awareness, as it requires not just that the force maker understands composition of processes and sums of costs, but also residuals of costs.

Dynamic Awareness carries a second interpretation, one less focused on rationality and more on military discipline. The force-maker is not just aware of the costs of conversion, but also of the costs of maintenance. The latter are not just a nuisance, but a discipline: they are the costs of keeping the force-maker's house in order. In a well-disciplined force, the configurations stay where they are—presumably doing as they are told—so there is no urge for them to decay into other configurations, and thus no urge for the slingshot effect to be large. Consider, as an easy example, George Washington's ability to keep his troops in line at Valley Forge during the winter of 1777–1778. His initial force of 12,000 troops was reduced to 10,000 (or so) by the end of the winter, but the force was kept in line and ready to fight thanks to training from (*inter alia*) Friedrich Wilhelm von Steuben. The configuration surely decayed, but this decay was bounded by the incredible efforts of the officer corps and the rugged resilience of the (poorly equipped) troops. This is the discipline that Dynamic Awareness points to: the ability to keep the force in line, to manage maintenance costs to the point that they do not induce unnecessary decay, and to keep the slingshot effect in check. And again, this deep substantive interpretation is fully equivalent to the Triangle Inequality in the dynamic setting, linking the force-maker's practices to the structure of the space of configurations. Though Washington likely did not reason in terms of a weak metric space, his practices were fully equivalent to those of a force-maker who did.

We can offer a concrete interpretation of the costs of force conversion in the context of the atomic theory of force. Recall to this point that we have:

1. defined a set of force molecules \mathbb{M}_L (Primitive 1.5) as connected graphs labeled by the elements of force L ;
2. defined a notion of graph union (Definition 2.4) and used it to construct a set of all force configurations \mathbb{M}_L^* (Primitive 2.2);
3. defined four elementary edit operations (Definition 2.11) and demonstrated their sufficiency for converting configurations (Lemma 2.12), giving rise to the category of force configurations **Config** (Construction 2.13) with morphisms the sequences of edits; and
4. attached to each elementary edit some cost in Ξ (Section 2.3).

It turns out that in the special case where this final step is done with non-negative, finite real numbers, we arrive at a well known construction in graph theory and computer science: the *edit distance*.³³

2.33 Definition (Edit Distance)

The edit distance between two force configurations $\biguplus \mathcal{M}_1$ and $\biguplus \mathcal{M}_2$ is the minimum cost of converting $\biguplus \mathcal{M}_1$ into $\biguplus \mathcal{M}_2$ using the elementary edits when $\text{cost} : \mathbb{M}_L^* \times \mathbb{M}_L^* \rightarrow \mathbb{R}_{\geq 0}$. We denote the edit distance by $\text{ed}(\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_2)$.

This definition of the edit distance contains two important features, namely that it addresses both choice (via the minimality requirement) and cost (via the cost of conversion). Taken together, these mean that choice implicitly chooses minimum-cost sequences of elementary edits, which is a well-posed notion with finite sequences of edits measured in real numbers. More formally, let $\biguplus \underline{\mathcal{M}}$ and $\biguplus \overline{\mathcal{M}}$ be two force configurations, and define $\text{Paths}(\biguplus \underline{\mathcal{M}}, \biguplus \overline{\mathcal{M}})$ as

$$\left\{ \biguplus \underline{\mathcal{M}} = \biguplus \mathcal{M}_0 \xrightarrow{\phi_1} \biguplus \mathcal{M}_1 \xrightarrow{\phi_2} \dots \xrightarrow{\phi_n} \biguplus \mathcal{M}_n = \biguplus \overline{\mathcal{M}} \mid n \in \mathbb{N} \right\},$$

Then the edit distance is given by

$$\text{ed}(\biguplus \underline{\mathcal{M}}, \biguplus \overline{\mathcal{M}}) = \min_{(\phi_1, \dots, \phi_n) \in \text{Paths}(\biguplus \underline{\mathcal{M}}, \biguplus \overline{\mathcal{M}})} \sum_{i=1}^n \text{cost}(\phi_i),$$

where cost measures the cost of each elementary edit in real numbers. This straightforward definition distills the problem down nicely.

³³The edit distance is a well-appreciated construction in the context of string theory, where it is used to measure the distance between two strings. See Kaspar Riesen's *Structural Pattern Recognition with Graph Edit Distance* (2015) for a comprehensive introduction.

The edit distance is indeed a distance function, as we now show.

2.34 Lemma (Edit Distance as a Metric)

The edit distance is a weak metric calibrated by $\mathbb{R}_{\geq 0}$.

Proof. We must show that the edit distance satisfies Identity of Discernibles and the Triangle Inequality, Weak Codomain being immediate. Identity of Discernibles is immediate from the definition of the edit distance, as the minimum cost of converting a configuration into itself is 0, via the identity morphism of no edits. We therefore focus on the Triangle Inequality. Choose any three force configurations $\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_2, \biguplus \mathcal{M}_3$ in \mathbb{M}_L^* . The Triangle Inequality is

$$\text{ed}(\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_2) + \text{ed}(\biguplus \mathcal{M}_2, \biguplus \mathcal{M}_3) \geq \text{ed}(\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_3),$$

or more explicitly as

$$\sum_{i=1}^{n_{12}} \text{cost}(\phi_i^{12}) + \sum_{j=1}^{n_{23}} \text{cost}(\phi_j^{23}) \geq \sum_{k=1}^{n_{13}} \text{cost}(\phi_k^{13}),$$

where ϕ_i^{12} , ϕ_j^{23} , and ϕ_k^{13} are the elementary edits in the minimum-cost sequences of edits converting $\biguplus \mathcal{M}_1$ into $\biguplus \mathcal{M}_2$, $\biguplus \mathcal{M}_2$ into $\biguplus \mathcal{M}_3$, and $\biguplus \mathcal{M}_1$ into $\biguplus \mathcal{M}_3$, respectively. Now suppose to the contrary that the Triangle Inequality is violated, so that

$$\sum_{i=1}^{n_{12}} \text{cost}(\phi_i^{12}) + \sum_{j=1}^{n_{23}} \text{cost}(\phi_j^{23}) < \sum_{k=1}^{n_{13}} \text{cost}(\phi_k^{13}),$$

where we know this from the completeness of \geq on $\mathbb{R}_{\geq 0}$. But then the right-hand side would fail to be the minimum cost of converting $\biguplus \mathcal{M}_1$ into $\biguplus \mathcal{M}_3$, since the sequence of elementary edits

$$\biguplus \mathcal{M}_1 \xrightarrow{\phi_1^{12}} \dots \xrightarrow{\phi_{n_{12}}^{12}} \biguplus \mathcal{M}_2 \xrightarrow{\phi_1^{23}} \dots \xrightarrow{\phi_{n_{23}}^{23}} \biguplus \mathcal{M}_3$$

would be a sequence of edits converting $\biguplus \mathcal{M}_1$ into $\biguplus \mathcal{M}_3$ at a cost strictly less than the minimum cost. This is a contradiction, and so the Triangle Inequality holds. We conclude that the edit distance is a weak metric calibrated by $\mathbb{R}_{\geq 0}$. ■

Naturally, one could come up with conditions under which the edit distance is a traditional metric, but we leave that for another day.³⁴ What matters here is that

³⁴See a recent paper by Francesc Serratosa (2021) for a set of weak conditions under which the edit distance is a metric; no within-molecule triangle inequality is required.

a concrete construction abides by the rules of distance in our atomic theory of force, providing intuitions about how the abstract construction works. We are content to leave the edit distance in weak form to avoid unnecessary clutter.

One remarkable thing about our set-up is that the metric space (\mathbb{M}_L^*, d) evolves in response to changes in the force-maker's practices or her understanding of costs. Which configurations are close and which are far apart depend on cost and **CS**, and so the space of configurations is dynamic. Suppose, for example, that cost and **CS** were functions of some parameter $\lambda \in \Lambda$. Then we could define a family of metric spaces $(\mathbb{M}_L^*, d_\lambda)$, where d_λ is the distance function associated with cost_λ and CS_λ . This family of metric spaces would allow us to track how the space of configurations changes as the force-maker's practices change. Even though the configurations themselves stay in one place, the distances between them shift as the force-maker's practices evolve; alternatively, the distances could stay in place and the configurations themselves could shift. In either case, we learn the deep lesson that points are fully determined by their relationships to other points, and that the space of configurations is a dynamic entity that evolves in response to the force-maker's practices. Just which deep parameter λ is responsible for this evolution is a question for another day, but it could be a function of technological progress, political change, or even the force-maker's own understanding of the world.

Summary of Section 2.5. We first thought through what an appropriate notion of distance would look like, eventually setting on the conception of a weak metric space calibrated by a quantale. We then considered the static case, where the force-maker treats each configuration as a fixed point in space, and showed that the Triangle Inequality is equivalent to Compositional Awareness, thus linking cost-based reasoning to distance-based reasoning in something like a principle of least action.³⁵ Remarkably, the dynamic case required us to take Compositional Awareness as a starting point and gradually strengthen it to account for maintenance costs, leading to the notion of Dynamic Awareness. The analysis honed in on the steadiness of a force configuration in the context of its likelihood to decay into another, suggesting that discipline is necessary for the force-maker to reason about costs in distance form. We will now use the distance function to construct a topology on the set of all force configurations, allowing us to reason about the structure of the space of configurations in terms of neighborhoods and open sets. Then, and only then, will we be prepared to issue a verdict on force as a concept in the context of political economy.

³⁵The principle of least action is a fundamental principle in physics, stating that the path taken by a system between two points is the one that minimizes the action, a quantity that is the integral of the Lagrangian over time. The principle of least action is equivalent to the Euler-Lagrange equations, which are the equations of motion in classical mechanics. In this context, a weak form of cost minimization serves as a principle of least action, suggesting that the force-maker's practices are consistent with a notion of distance.

2.6 Convergence

We have introduced distance in the service of a larger goal: namely, to construct a political-economic topology for the set of all force configurations \mathbb{M}_L^* . A topology is a way of understanding the structure of a set, particularly with respect to concepts like continuity and convergence. Formally, a topology on \mathbb{M}_L^* is a collection \mathcal{T}_Ξ of subsets of \mathbb{M}_L^* satisfying three properties.

1. *The Empty Set and the Whole Set*: $\emptyset, \mathbb{M}_L^* \in \mathcal{T}_\Xi$;
2. *Finite Intersections*: if $U_1, \dots, U_n \in \mathcal{T}_\Xi$, then $\bigcap_{i=1}^n U_i \in \mathcal{T}_\Xi$; and
3. *Arbitrary Unions*: if $\{U_i\}_{i \in I} \subseteq \mathcal{T}_\Xi$, then $\bigcup_{i \in I} U_i \in \mathcal{T}_\Xi$.

Any collection of subsets satisfying these properties is a topology, but we are especially interested in topologies that reflect the underlying political-economic problem at hand.

We begin with the idea of a *neighborhood* of a configuration.

2.35 Definition (Neighborhoods)

Let $\biguplus \mathcal{M} \in \mathbb{M}_L^*$ and $\xi \in \Xi$. A neighborhood of $\biguplus \mathcal{M}$ is a set

$$\mathcal{N}_\xi(\biguplus \mathcal{M}) := \{\biguplus \mathcal{N} \in \mathbb{M}_L^* \mid \xi > d(\biguplus \mathcal{M}, \biguplus \mathcal{N})\}.$$

A neighborhood of a configuration is the set of all configurations that can be converted into the configuration at a cost strictly less than ξ . This is a natural way to think about the space of force configurations: two configurations are close if they can be converted into one another at a low cost. The neighborhood of a configuration is the set of all configurations that are close to it in this sense.

Neighborhoods are the building blocks of a topology, and we now use them to define the open sets of a topology.

2.36 Definition (The Candidate Conversion Cost Topology)

A set $U \subseteq \mathbb{M}_L^*$ is open just in case for all configurations $\biguplus \mathcal{M} \in U$, there exists a cost $\xi \in \Xi$ such that $\mathcal{N}_\xi(\biguplus \mathcal{M}) \subseteq U$. The collection of open sets is denoted $\mathcal{T}_{\mathbb{M}_L^*}$.

In words, a set of configurations is open if, for every configuration in the set, there is a cost such that all configurations close to it at that cost are also in the set. This generalizes our notion of open sets in Euclidean space: a set is open if, for every point in the set, there is a radius such that all points within that radius are also in the set. Here ξ is our (rather abstract) radius and $\mathcal{N}_\xi(\biguplus \mathcal{M})$ is our (rather abstract) ball. This is the essence of a topology: it is a way of understanding the structure of a set in terms of neighborhoods and open sets.

Yet again, an example will help to clarify the concept. Consider again the hoplite as depicted in Figure 8. What are the neighborhoods of this hoplite

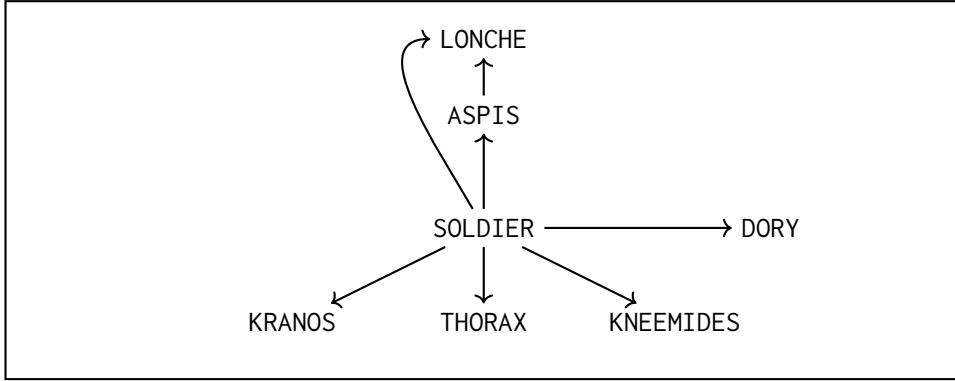


Figure 8: A hoplite molecule, dory/lonche variant.

configuration? Well, in this case, we might think that the decorative LONCHE is a bit of a luxury, and presumably it is not too difficult to convert the hoplite into a hoplite without a LONCHE—this would require the removal of one vertex and two edges. This, if given a small cost radius ξ , the neighborhood of our original hoplite molecule would be itself, the hoplite molecule without the LONCHE, and any other molecule similarly “close” to the original with respect to the costs of conversion. This is why our notion of distance was such a useful concept: it allows us to reason about the structure of the space of configurations in terms of costs of conversion as if they were distances in a metric space.³⁶ A larger neighborhood would include these configurations and others further away—say, ones where the weapon switches from a DORY to a XIPHOS or the ASPIS is replaced by a PELTA. But in general, the idea here is that each configuration is surrounded by a class of abstract balls measured in costs, and within each ball are all the configurations that are close to the original in terms of costs of conversion. This is the essence of a neighborhood, and it generalizes our intuitions from elementary analysis to the space of force configurations.³⁷

³⁶ It ought to be noted that the metric structure is not necessary for the construction of a topology, which indeed is meant to generalize the notion of distance. And yet, few people in the social sciences voluntarily work with non-metrizable topologies, as they are difficult to reason about. The metric structure is a useful crutch, as it provides a systematic way for us to apply something we can conceive of—the costs of conversion—to something trickier.

³⁷ Observe that we have $d(\biguplus \mathcal{M}, \biguplus \mathcal{M}) = 0_{\Xi}$ by construction, so neighborhoods are never empty. As we have allowed distinct configurations to have zero distance between them, it could well be that $\mathcal{N}_{0_{\Xi}}(\biguplus \mathcal{M}) \neq \{\biguplus \mathcal{M}\}$, suggesting that multiple points are in the same place. Thus,

We confirm that the collection of open sets $\mathcal{T}_{\mathbb{M}_L^*}$ forms a topology on \mathbb{M}_L^* .

2.37 Proposition (The Conversion Cost Topology)

$\mathcal{T}_{\mathbb{M}_L^*}$ forms a topology on \mathbb{M}_L^* ; we call it the conversion cost topology.

Proof. We handle each requirement in turn.

1. *The Empty Set and the Whole Set.* The empty set \emptyset is open vacuously, as there are no configurations in it and thus no neighborhood tests to fail. The whole set \mathbb{M}_L^* is open, as for any configuration $\biguplus \mathcal{M} \in \mathbb{M}_L^*$ and any cost $\xi \in \Xi$, we have $\mathcal{N}_\xi(\biguplus \mathcal{M}) \subseteq \mathbb{M}_L^*$.
2. *Finite Intersections.* Let $U_1, \dots, U_n \in \mathcal{T}_{\mathbb{M}_L^*}$; we must show that $\bigcap_{i=1}^n U_i \in \mathcal{T}_{\mathbb{M}_L^*}$. We may assume that $\bigcap_{i=1}^n U_i \neq \emptyset$, as otherwise the result is immediate (by the first part of the proof). So, choose and fix an arbitrary $\biguplus \mathcal{M} \in \bigcap_{i=1}^n U_i$. We must identify a cost $\xi \in \Xi$ such that $\mathcal{N}_\xi(\biguplus \mathcal{M}) \subseteq \bigcap_{i=1}^n U_i$. Since $\biguplus \mathcal{M} \in U_i$ for all $i \in \{1, \dots, n\}$, there exist costs $\xi_i \in \Xi$ such that $\mathcal{N}_{\xi_i}(\biguplus \mathcal{M}) \subseteq U_i$. Choose any ξ satisfying $\xi_i \geq \xi$ for all $i \in \{1, \dots, n\}$; at minimum, this includes $0_{\text{Cost}} \in \Xi$ by construction. Then we have $\mathcal{N}_\xi(\biguplus \mathcal{M}) \subseteq \mathcal{N}_{\xi_i}(\biguplus \mathcal{M}) \subseteq U_i$ for all $i \in \{1, \dots, n\}$, and so $\mathcal{N}_\xi(\biguplus \mathcal{M}) \subseteq \bigcap_{i=1}^n U_i$.
3. *Arbitrary Unions.* Let $\{U_i\}_{i \in I} \subseteq \mathcal{T}_{\mathbb{M}_L^*}$; we must show that $\bigcup_{i \in I} U_i \in \mathcal{T}_{\mathbb{M}_L^*}$. Choose and fix an arbitrary $\biguplus \mathcal{M} \in \bigcup_{i \in I} U_i$. Then there exists some $i \in I$ such that $\biguplus \mathcal{M} \in U_i$, and thus a cost $\xi_i \in \Xi$ such that $\mathcal{N}_{\xi_i}(\biguplus \mathcal{M}) \subseteq U_i$. Then we have $\mathcal{N}_{\xi_i}(\biguplus \mathcal{M}) \subseteq \bigcup_{i \in I} U_i$, and so $\bigcup_{i \in I} U_i \in \mathcal{T}_{\mathbb{M}_L^*}$.

As all three requirements are satisfied, $\mathcal{T}_{\mathbb{M}_L^*}$ forms a topology on \mathbb{M}_L^* . ■

The proof of Proposition 2.37 is a straightforward exercise,³⁸ but it is worth noting that the topology $\mathcal{T}_{\mathbb{M}_L^*}$ is generated by the map $\text{cost} \circ \text{CS}$, which represents the force-maker's conversion practices and associated costs. It also relies on the structure Ξ , both in terms of its extremal elements and its ordering. These elements all encode various aspects of the problem at hand and the force-maker's reasoning about it. Moreover, that the proof is straightforward bodes

the space might not satisfy the Hausdorff separation axiom, but this is a feature, not a bug.

³⁸One gains intuitions from the standard textbooks in general topology, such as James R. Munkres's *Topology* (2000) or John L. Kelley's *General Topology* (1975). As with so many things, one can also find these definitions in Charalambos D. Aliprantis and Kim C. Border's *Infinite Dimensional Analysis: A Hitchhiker's Guide* (2006).

well should future work consider, say, continuous evolution of practices, costs, and calibrations. Such work lies beyond the present scope, but the present construction ought to provide reasonable foundations for it.

Topologies may have many properties, but we are especially interested in *second-countability* and *Hausdorffness*. In words, a topology is second-countable if it is simple enough to be characterized by a countable set of sets.

2.38 Definition (Second-Countability)

A topology \mathcal{T} is second-countable just in case there exists

$$\mathcal{B} = \{B_i\}_{i \in \mathbb{N}},$$

called a basis, such that every open set in \mathcal{T} can be written as a union of elements of \mathcal{B} :

$$U = \bigcup_{i \in I} B_i \quad \text{for some } I \subseteq \mathbb{N}.$$

Second-countability is a useful property, as it ensures that the topology is not too large. It is a generalization of the idea of a countable basis for a topology, which is a familiar concept from elementary analysis. Owing to the countability of \mathbb{M}_L^* , we have the following.

2.39 Corollary (Second-Countability of $\mathcal{T}_{\mathbb{M}_L^*}$)

The convergence cost topology $\mathcal{T}_{\mathbb{M}_L^}$ is second-countable.*

Clearly, we can set $\mathcal{B} = \mathbb{M}_L^*$; any open set can be written as a union of elements of from this basis, as any open set is comprised merely of the configurations in \mathbb{M}_L^* itself. Though simple, Corollary 2.39 is one of the most important results in this theory, as it will allow us to discern a sufficient condition for force representations. We are nearly there!

We next define the Hausdorff property.

2.40 Definition (Hausdorffness)

A topology \mathcal{T} is Hausdorff just in case for all distinct configurations $\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_2 \in \mathbb{M}_L^$, there exist neighborhoods $\mathcal{N}_{\xi_1}(\biguplus \mathcal{M}_1)$ and $\mathcal{N}_{\xi_2}(\biguplus \mathcal{M}_2)$ such that*

$$\mathcal{N}_{\xi_1}(\biguplus \mathcal{M}_1) \cap \mathcal{N}_{\xi_2}(\biguplus \mathcal{M}_2) = \emptyset.$$

In words, a topology is Hausdorff if every pair of distinct configurations can be separated by neighborhoods. This is a generalization of the idea of distinct points in a metric space being separated by open balls, and it is a fundamental property of topological spaces. In our setting, however, Hausdorffness is not necessarily guaranteed, for reasons to be discussed presently.

Why might the topology $\mathcal{T}_{\mathbb{M}_L^*}$ fail to be Hausdorff? The most natural reason is that there exist two distinct configurations with distance 0_{Ξ} between them.³⁹ When such is the case, any neighborhood including one is sure to include the other, and so the two configurations cannot be separated by neighborhoods. This is a natural consequence of our construction, as we have allowed for the possibility of distinct configurations with zero distance between them, be it in the static or the dynamic case. It therefore stands to reason to ask whether we can recover Hausdorffness by (1) imposing a condition on the cost function cost; or (2) constructing a quotient space that identifies configurations with zero distance between them. Let us consider the first option first; we begin by introducing a *regular distance function*.

2.41 Definition (Regular Distance Function)

A distance function $d : \mathbb{M}_L^ \times \mathbb{M}_L^* \rightarrow \Xi$ is regular just in case for all distinct configurations $\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_2 \in \mathbb{M}_L^*$, we have $d(\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_2) > 0_{\Xi}$.*

As an example, the graph edit distance is regular so long as all elementary edits has a strictly-positive cost. Thus, in the context of the atomic theory of force, we do indeed recover the Hausdorff property so long as elementary edits are costly. A similar property obtains in the dynamic case where all inter-configuration distances are strictly larger than the associated slingshot maintenance cost. Regardless of the reasons, we have the following.

2.42 Lemma (Regular Spaces are Hausdorff)

If the distance function d is regular, then the topology $\mathcal{T}_{\mathbb{M}_L^}$ is Hausdorff.*

Proof. This is nearly immediate; choose any two configurations $\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_2 \in \mathbb{M}_L^*$. Since d is regular, we have $d(\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_2) > 0_{\Xi}$. Set $\xi_1 = \xi_2 = 0_{\Xi}$; then from regularity, we have

$$\mathcal{N}_{0_{\Xi}}(\biguplus \mathcal{M}_1) \cap \mathcal{N}_{0_{\Xi}}(\biguplus \mathcal{M}_2) = \{\biguplus \mathcal{M}_1\} \cap \{\biguplus \mathcal{M}_2\} = \emptyset,$$

implying that the topology $\mathcal{T}_{\mathbb{M}_L^*}$ is Hausdorff. ■

Thus, if the distance function behaves according to our spatial intuitions, the associated topology it generates is indeed Hausdorff, which simplifies the analysis considerably by providing a clear separation between configurations.

³⁹In just a moment, we will use distance- 0_{Ξ} configurations to discuss military preparedness, a hallmark of the organized nature of force. But for now, we are concerned with the topological properties of the space of configurations.

The second option for recovering Hausdorffness is to construct a quotient space that identifies configurations with zero distance between them. This is a more radical solution, as it involves changing the space of configurations itself, or at minimum simplifying it in a structured way. The construction requires a simple assumption: that the distance function d is symmetric for zeroes—*i.e.*,

$$d\left(\bigcup \mathcal{M}_1, \bigcup \mathcal{M}_2\right) = 0_{\Xi} \quad \text{implies} \quad d\left(\bigcup \mathcal{M}_2, \bigcup \mathcal{M}_1\right) = 0_{\Xi}.$$

Thus, we remove the possibility of left- and right-zeroes, as we are only interested in configurations that are close to one another in terms of the distance function.

1. *Equivalence Relation.* Define an equivalence relation \sim on \mathbb{M}_L^* by writing $\bigcup \mathcal{M}_1 \sim \bigcup \mathcal{M}_2$ just in case

$$d\left(\bigcup \mathcal{M}_1, \bigcup \mathcal{M}_2\right) = d\left(\bigcup \mathcal{M}_2, \bigcup \mathcal{M}_1\right) = 0_{\Xi}.$$

This is an equivalence relation, as we can check.

- (a) *Reflexivity.* Clearly, $\bigcup \mathcal{M} \sim \bigcup \mathcal{M}$, as $d\left(\bigcup \mathcal{M}, \bigcup \mathcal{M}\right) = 0_{\Xi}$.
- (b) *Symmetry.* d is symmetric at the zeroes, so this is immediate.
- (c) *Transitivity.* Let $\bigcup \mathcal{M}_1, \bigcup \mathcal{M}_2, \bigcup \mathcal{M}_3 \in \mathbb{M}_L^*$ be such that

$$\bigcup \mathcal{M}_1 \sim \bigcup \mathcal{M}_2 \quad \text{and} \quad \bigcup \mathcal{M}_2 \sim \bigcup \mathcal{M}_3.$$

Then we have $d\left(\bigcup \mathcal{M}_1, \bigcup \mathcal{M}_2\right) = d\left(\bigcup \mathcal{M}_2, \bigcup \mathcal{M}_3\right) = 0_{\Xi}$. From the Triangle Inequality, we have

$$0_{\Xi} \oplus 0_{\Xi} \geq d\left(\bigcup \mathcal{M}_1, \bigcup \mathcal{M}_3\right),$$

implying that $d\left(\bigcup \mathcal{M}_1, \bigcup \mathcal{M}_3\right) = 0_{\Xi}$. Since \sim is symmetric, the same logic can show that $d\left(\bigcup \mathcal{M}_3, \bigcup \mathcal{M}_1\right) = 0_{\Xi}$, and so $\bigcup \mathcal{M}_1 \sim \bigcup \mathcal{M}_3$, as required.

2. *Equivalence Classes.* The equivalence classes of \sim are the sets

$$[\bigcup \mathcal{M}] = \left\{ \bigcup \mathcal{N} \in \mathbb{M}_L^* \mid \bigcup \mathcal{M} \sim \bigcup \mathcal{N} \right\}.$$

These equivalence classes are the sets of configurations that are close to one another in terms of the distance function d . In case the distance function is regular, the equivalence classes are the singleton sets, as the distance between any two distinct configurations is strictly positive. In case the distance function is not regular, at least one equivalence class will contain more than one configuration. The equivalence classes are the building blocks of the quotient space:

$$\mathbb{M}_L^*/\sim = \left\{ [\bigcup \mathcal{M}] \mid \bigcup \mathcal{M} \in \mathbb{M}_L^* \right\}.$$

3. *Obtaining Distances on the Quotient Space.* Now that we have our equivalence classes, we can define a distance function on the quotient space \mathbb{M}_L^*/\sim by picking the lowest distance between any two configurations in the equivalence classes—that is, by setting $d_\sim([\bigcup \mathcal{M}_1], [\bigcup \mathcal{M}_2])$ to be the infimum of the set of distances between $\bigcup \mathcal{M}_1$ and $\bigcup \mathcal{M}_2$:

$$\bigvee \{d([\bigcup \mathcal{N}_1], [\bigcup \mathcal{N}_2]) \mid [\bigcup \mathcal{N}_1] \in [\bigcup \mathcal{M}_1], [\bigcup \mathcal{N}_2] \in [\bigcup \mathcal{M}_2]\}.$$

This approximation of the distance function is well-defined, as Ξ is a quantale and thus has infima. Notice that this definition does not necessarily deliver symmetry of distances, but then again, we were only working with a weak metric space to begin with. As a matter of course, we observe that

$$d_\sim([\bigcup \mathcal{M}], [\bigcup \mathcal{M}]) = 0_\Xi,$$

and indeed symmetry at the zeroes implies that this is the only zero. Let us now check the Triangle Inequality: choose any three equivalence classes $[\bigcup \mathcal{M}_1], [\bigcup \mathcal{M}_2], [\bigcup \mathcal{M}_3] \in \mathbb{M}_L^*/\sim$. Let us first consider the sum on the left-hand side of the Triangle Inequality:

$$d_\sim([\bigcup \mathcal{M}_1], [\bigcup \mathcal{M}_2]) \oplus d_\sim([\bigcup \mathcal{M}_2], [\bigcup \mathcal{M}_3]).$$

By definition, this is the sum of the infima of the sets of distances between configurations in the equivalence classes. It involves four configurations: one from $[\bigcup \mathcal{M}_1]$, two from $[\bigcup \mathcal{M}_2]$, and one from $[\bigcup \mathcal{M}_3]$. The first infimum describes a path from $\bigcup \mathcal{M}_1$ to $\bigcup \mathcal{M}_2$, and the second infimum describes a path from $\bigcup \mathcal{M}_2$ to $\bigcup \mathcal{M}_3$. In particular, it need not be the case that $\bigcup \mathcal{M}_2 = \bigcup \mathcal{M}_2$, since the target of the first path might not be the same as the source of the second. Thus, the overall path described on the left-hand side is

$$\bigcup \mathcal{M}_1 \longrightarrow \bigcup \mathcal{M}_2 \longrightarrow \bigcup \mathcal{M}_2 \longrightarrow \bigcup \mathcal{M}_3,$$

which is still a path from $\bigcup \mathcal{M}_1$ to $\bigcup \mathcal{M}_3$. As such, it is part of the set over which the infimum is taken on the right-hand side of the Triangle Inequality, and so the Triangle Inequality holds.

We may conclude:

2.43 Proposition (Quotient Spaces)

$(\mathbb{M}_L^*/\sim, d_\sim)$ is a weak metric space calibrated by Ξ ; its metric topology is Hausdorff.

Thus, the construction works as promised, so long as the distance function is symmetric at the zeroes. This is a mild condition, so we are happy to leave it at that without further digging into the structure of the quotient space.

In the context of the atomic theory of force, there are other topologies that might be of interest, not just the political-economic one derived from the costs of conversion. We have been telling a particular story about the nature of force, but other contexts might require different stories. The graph edit distance topology used to this point induces an enriched category in **Graph**, and we have used that to define a weak metric and subsequently a topology. This approach is especially useful for thinking about how a force-maker might reason about the space of configurations. Other approaches might be more useful for other purposes; consider, for example, a subgraph distance approach. This is similar to the graph edit distance approach, but it measures the distance between two graphs by the number of edits required to turn one into a *subgraph* of the other, rather than the number of edits required to turn one into the other. Consider, for example, a situation where $\biguplus \mathcal{M}_1$ is a large army of uniformly-equipped hoplites and $\biguplus \mathcal{M}_2$ is a single hoplite equipped just the same way as those in $\biguplus \mathcal{M}_1$. The graph edit distance between these two configurations would be quite large, as the two graphs are quite different. But, we would have $d_{\text{subgraph}}(\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_2) = 0$, as the single hoplite is a subgraph of the large army. This is a different way of thinking about the space of configurations, one more naturally concerned with the structure of the configurations—and in particular, their natural structural hierarchies—than with the costs of conversion between them. This is a different story about the nature of force, and it might be more useful in different contexts.⁴⁰

Summary of Section 2.6. We have constructed a topology on the space of force configurations, \mathbb{M}_L^* , by using the costs of conversion between configurations as a proxy for distances. The open sets in this topology are the neighborhoods of configurations defined in political-economic terms via the weak metric space structure. Since the set of all force configurations is countable, the topology naturally inherits an important topological property called second-countability. Less obvious is another property called Hausdorffness, which is not guaranteed by the construction. We discussed two ways to recover Hausdorffness: by imposing a mild condition on the distance function or by constructing a quotient space that identifies configurations with zero distance between them, again under mild restrictions on the distance function. The first of these delivers Hausdorffness directly, whereas the latter does so after issuing the more radical

⁴⁰Naturally, there are topologies on graph spaces other than the graph edit distance topology and the subgraph distance topology. In uncertain settings on large graphs, one might consider more probabilistic approaches like the one defined by Benjamini-Schramm convergence, which is a topology on the space of graphs that is induced by a random walk on the graph.

change of simplifying the space of configurations itself.

Summary of Section 2. This section, the core of the manuscript, has been a journey through a resource-theoretic approach to force. We entered into the section armed only with the set of all force molecules, \mathbb{M}_L , a subset of the set of all finite connected graphs labeled by the elements of force L . Included in this set are the isolated elements, such as a soldier or a weapon, as well as the more complex molecular structures, such as a hoplite or a phalanx. Then:

1. In Section 2.1 we introduced a gathering operation \uplus that allows us to combine force molecules into force configurations, resulting in the set of all force configurations, \mathbb{M}_L^* , itself a subset of the free monoid on \mathbb{M}_L generated by \uplus .
2. In Section 2.2 we introduced hom sets linking force configurations, where these encode the processes by which one configuration can be converted into another. Each $\text{Hom}(\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_2)$ contains all ways in which $\biguplus \mathcal{M}_1$ can be converted into $\biguplus \mathcal{M}_2$, including an impossible process \mathbf{x} .
3. In Section 2.3 we introduced a cost calibration structure Ξ , a quantale that allows us to assign costs to the conversion processes. Just as time and space are (typically) calibrated by real numbers, so too are the costs of conversion calibrated by elements of Ξ .
4. In Section 2.4 we introduced a choice structure \mathbf{CS} that allows us to reason about the force-maker's decision-making process. It provided an opportunity to impose a mild behavioral postulate called *Compositional Awareness*, which essentially encodes the idea that the force-maker understands composition and the combination of costs as well as we do.
5. In Section 2.5 we used all of these structures to introduce a generalized weak metric structure on the space of force configurations, \mathbb{M}_L^* . We saw that, in the static case where maintenance processes are trivialized, Compositional Awareness is precisely the same as the Triangle Inequality. In the dynamic case, Compositional Awareness must be strengthened to Dynamic Awareness, a more stringent condition.
6. Finally, in Section 2.6 we used the aforementioned metric structure to construct a topology on the space of force configurations, \mathbb{M}_L^* . The topology has a countable basis and is Hausdorff, provided the distance function is regular or the space is quotiented by the zero distances.

We now turn our attention to what can be done with this structure, and in particular to how we can reason about the organization and representation of force. These are the subjects of the next two sections, respectively.

3 Force is Organized

It is difficult to conceive of the force produced by a state—or even a state-like entity—without considering how it is organized. One can imagine a completely unstructured collection of soldiers, each acting under their own initiative, but we'd like not refer to such a collection as a “force.” Such a thing might be a mob, a riot, or a band of brigands, but it is not a force in the sense we are interested in. The force produced by a state is necessarily organized, and this is a key part of the force itself. We therefore cannot study force without studying its structure.

This is a hoary topic, predating the modern nation-state by over a millennium. As evidence, consider the detail with which Julius Cæsar describes the organization, and reorganization, of Roman forces in his *Commentaries*. The main unit of the Roman infantry was the legion, each of which included 3,500–5,000 men (Judson, 1888). During the Gallic Wars, Cæsar commanded at least twelve such legions, which were assigned numbers corresponding to the order in which they were formed. Each legion was divided into ten cohorts; each cohort was divided into three maniples; and each maniple was divided into two centuries (or *ordo*). The centuries were the basic tactical unit; depending on the historian consulted, they included between 60–120 infantrymen.⁴¹ The Roman cavalry was drawn from the citizenry, and it was organized into separate units called *ala*, each of which was divided into *turmae*. Organization was a less-rigid affair in the cavalry, and many horsemen were organized into ad-hoc units as needed. Artillery units—catapults, ballistae, and the like—were even less organized, and they were often attached to infantry units as needed.

The Roman example, though a bit overplayed, is exceedingly important. Niccolò Machiavelli, for example, considered the Roman legion the epitome of military organization, and he used it as a model for his own military reforms in Florence, right down to the disdain of cavalry and artillery; these reforms are discussed at length in Machiavelli's *Art of War* (1521). The aforementioned Maurice of Nassau admired Julius Cæsar's organization of force, and he sought to emulate it in his own army (McNeill, 1982, pp. 129–130). Perhaps most drastically, the reorganization of the French army into the corps system by Napoleon Bonaparte is credited with much of his success in the field (see, e.g., Chandler, 1966), and these corps were themselves organized along the lines of the Roman legion. Since part of the motivation for this section is that *force-makers seem to care a great deal about organization*, the Roman example is a good place to

⁴¹Judson (1888, p. 1) defines a tactical unit as “a body of troops under a single command, by a combination of several of which a higher unit is formed.” This fractal-like structure is a key part of the Roman military organization, and indeed of military organization more generally.

start given its prominence in the minds of so many real-world force-makers.

3.1 The Structure of Force

Organization being a key part of force, we ought to incorporate it into our theory. We begin by establishing a primitive notion of force structure, which is a partial order on a set of force units representing the chain of command.

3.1 Primitive (Force Structure)

A force structure is a finite partial order (V, \preccurlyeq) , where each $v \in V$ is a force unit and \preccurlyeq is a subordination relation on V , where $v_1 \preccurlyeq v_2$ means that v_1 is subordinate to v_2 . The set of all force structures is denoted

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

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

The chain of command encoded by the partial order \preccurlyeq is a key part of the force structure. Being a partial order, \preccurlyeq is reflexive, transitive, and antisymmetric. Reflexivity will not factor into the theory much, as it is a trivial property: every force unit is (in a sense) subordinate to itself, suggesting a minimal degree of autonomy implicit in being called a unit at all. More important, of course, is transitivity, which ensures a coherent chain of command: if a company is subordinate to a battalion and a battalion subordinate to a brigade, then the company is subordinate to the brigade, too. Finally, antisymmetry ensures that the chain of command is not circular: no two distinct units can be subordinate to one another, ensuring that the chain of command is a tree. The substantive properties encoded by transitivity and antisymmetry seem reasonable in the context of organized force, where much of the point of being organized at all is to avoid violations of these properties.

It should be noted here that we can weaken the antisymmetry requirement to allow for more general organizational structures without any technical difficulty: nearly all of the results in this section will hold for preordered sets—*i.e.*, sets equipped with a relation that is merely reflexive and transitive. The antisymmetry requirement represents substantive motivations more than technical ones, as it encodes the important real-world property that no two distinct units are mutually subordinate. When we add non-command relationships among units to the force structure, we will keep things quite general, allowing for a wide variety of organizational structures for support, training, logistics, and so on. But it seems to your humble author that modern chains of command do not tolerate circularity, and so we will keep the antisymmetry requirement in place for subordination. Of course, the reader is free to consider more general organizational structures if they wish.

We have enough apparatus to envision the structuring of force in a way quite close to how force-makers do—after all, the proclivity of excellent military minds to reorganize their forces suggests that the organization of force is a key part of their reasoning. Figure 9 shows what one of our Roman cohorts looks like in a force structure: superior to three maniples, each of which is superior to two centuries. Figure 10 shows the complete force structure of a Roman legion, with a single cohort at the top, followed by three maniples, each of which is followed by two centuries; the fractal nature just described is quite evident in this presentation. So, for example, we have that $\text{Century 1.1} \preceq \text{Maniple 1} \preceq \text{Cohort}$ and that $\text{Cohort 01} \preceq \text{Legion}$, among various other subordination relations.

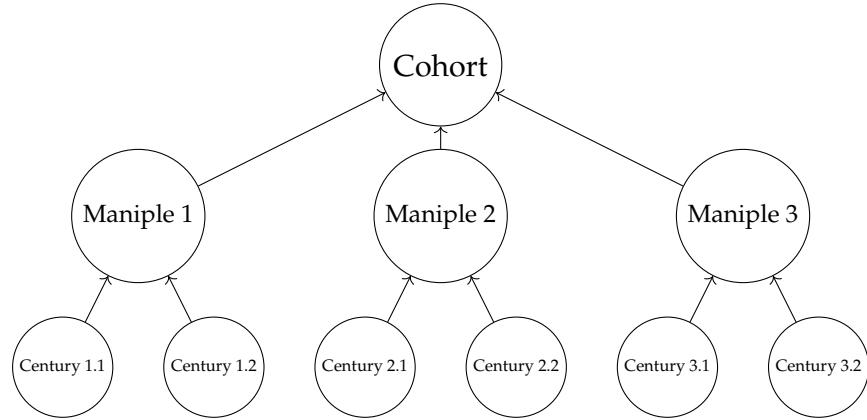


Figure 9: Sample presentation of a force structure; the arrows represent subordination.

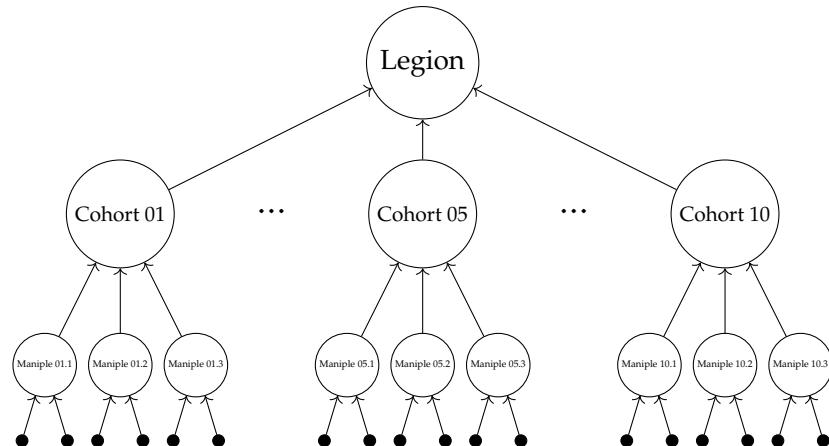


Figure 10: The complete force structure of a Roman legion; the black dots represent centuries.

The Roman legion was a model of organization, copied by many—though not all—of Rome’s successors. Indeed, nostalgia for the old legion system is a recurring theme in Machiavelli’s *Art of War* (1521), who laments the *condottieri*—literally, the “contractors”—who had supplanted the Roman legions and provided medieval Italian city-states with their military muscle. Medieval and Renaissance Italian city-states could count on loosely-organized local militias drawn from the populace, but they were rarely called upon to do more than fortify a wall or chase down a recalcitrant noble. These local militias, the vestiges of feudal levies, were almost exclusively infantry; cavalry—more expensive to maintain, but by this point the dominant arm of the battlefield—was provided by the *condottieri*, who were essentially private military contractors. An Italian city-state did not hire at the company level, but rather did business with individual captains, who in turn hired companies of soldiers. The *condottieri* were not subordinate to the city-state, but rather to the captain who hired them. City-states, eager for decisive victories but lacking in effective control, often watched with frustration while their hired captains engaged in mealy-mouthed battles, each unwilling to commit to a decisive engagement. The lack of political unity—or more to the point, the political body’s ineffective control over force—is the origin of Machiavelli’s frustration.⁴²

Condottieri companies varied widely in their organization, but they were still organized. Michael Mallett (1974) outlines the organization of several such companies; large companies were divided into squadrons, each led by a *squadriere* with a contract with the *condottiere*; leading fighters in a squadron had their own contracts with the *squadriere*. The *condottiere* had his own squadron, known as the *casa*, and this was typically the largest and best-equipped. For example, mid-1400s *condottiere* Tiberto Brandolini had a company of 400 lances and 300 infantrymen; these were divided into his *casa* and seven cavalry squadrons and eight infantry companies. His *casa* included men-at-arms, light cavalry, his master of the horse, his marshal, his chaplain, two cooks, six chancellors, trumpeters, a billeting officer, and a munitions officer (p. 108). So, a given company was well-organized, but it did not fall into the same integrated command structure.

Over time, cities became more responsible for their own force; Italian armies included large, city-provided infantry forces supplemented by hired cavalry. These forces were folded into the city-state’s own force structure. So-called *lanze spezzate*—“broken lances”—were cavalrymen who had broken away from their *condottiere* and joined the city-state’s army. The famous reforms of Charles VIII of France helped complete the transition from *condottieri* to city-state armies.

⁴² Michael Howard (1976, Chapter 2) claims Machiavelli makes too strong of a point against the *condottieri*, but he also describes many of the same problems.

Our machinery allows for disjointed force structures like those of the *condottieri* system. We have not required that the force structure be connected; that is, we have not required that there be a single unit that commands all others. This is a feature, not a bug; it allows us to model the *condottieri* system as a force structure. We can differentiate Cæsar's legions from the *condottieri* companies by requiring that the force structure have a single unit that commands all others, a property we will call *unity of command*. We begin by formalizing a command chain in a force structure.

3.2 Definition (Command Chain)

A command chain from force unit v_0 to force unit v_n is a sequence $(v_0, \dots, v_n) \in V^{n+1}$ of force units where

$$v_n \preccurlyeq v_{n-1} \preccurlyeq \dots \preccurlyeq v_1 \preccurlyeq v_0.$$

In case there exists a command chain from v_0 to v_n , we say that v_0 commands v_n .

For example, perhaps v_0 is a brigade, v_1 a battalion, and so on down to v_n , a fireteam. In this case, the brigade commands the fireteam.⁴³

Machiavelli's lamentations had something to do with the lack of a command chain; in particular, he wished that the force embodied what we will call *unity of command*. Let us formalize this.

3.3 Definition (Unity of Command)

A force structure (V, \preccurlyeq) has unity of command just in case there exists a force unit $v_\star \in V$ such that there is a command chain from v_\star to every other unit in V .

Figures 9 and 10 both feature unity of command, with the cohort and the legion, respectively, serving as the root of the tree. One can spoil the property by removing a single arrow from either diagram. Note that v_\star might represent an individual like a general, a queen, or a president; it just as easily can represent a committee, a consulate, or a Senate. Unity of command is a feature of the force structure that ensures that all force units are ultimately subordinate to a single unit. In graph-theoretic terms, this means that the Hasse diagram for the force structure is a *rooted tree*, which is a useful concept in graph theory.

⁴³Here questions of reflexivity arise. If the subordination relation is reflexive, then the brigade commands itself, providing a command chain of length zero. If the subordination relation is not reflexive, then the brigade does not command itself. We will consider chains of length zero to be valid command chains, but this is rather trivial.

To this point, we have considered both unified command structures like the Roman legion and disjointed command structures like the *condottieri*. It is interesting to wonder whether these can be considered as special cases of a more general class of force structures. One way to do so—among a great many that spring to mind without much effort—is to attach to each subordination relation a *weight* that represents the degree of subordination, the quality of the command, or some other relevant feature. The resulting structure is a *weighted force structure*, which we encode with the following construction.

3.4 Construction (Weighted Force Structures)

A weighted force structure is a category enriched in the symmetric monoidal preorder

$$\mathbf{Command} := ([0, 1], \leq, \times, 1).$$

Explicitly:

1. the objects of the category are V , the set of force units;
2. for each pair of force units $v_1, v_2 \in V$, the hom-object $\text{Hom}(v_1, v_2) \in [0, 1]$ is the strength of command from v_1 to v_2 ;
3. for each force unit $v \in V$, we have $\text{Hom}(v, v) = 1$;
4. $\text{Hom}(v_1, v_2) > 0$ implies $\text{Hom}(v_2, v_1) = 0$; and
5. for enrichment, we require

$$\text{Hom}(v_1, v_2) \times \text{Hom}(v_2, v_3) \leq \text{Hom}(v_1, v_3)$$

for all $v_1, v_2, v_3 \in V$.

The strength of command is a measure of the degree of subordination between two force units. The strength of command is a number between 0 and 1, where 0 means that the two units are not subordinate to one another and 1 means that one is completely subordinate to the other. The enrichment requirement demands that the strength of command decays as we add more layers to the organizational structure: a command directly from v_1 to v_3 is stronger than flows from v_1 to v_3 via some intermediary v_2 . The difference between $\text{Hom}(v_1, v_3)$ and $\text{Hom}(v_1, v_2) \times \text{Hom}(v_2, v_3)$ is a measure of the *friction* in the command chain, and it is a key part of the structure of the force. In this structure, a force structure with unity of command might mean that there is a single force unit v_* such that $\text{Hom}(v_*, v) = 1$ for all $v \in V$ and $\text{Hom}(v, v_*) = 0$ for all $v \neq v_*$. Conversely, a disjointed command structure might mean that there exist multiple units v_1, \dots, v_n such that $\text{Hom}(v_i, v_j) = 0$ for all $i \neq j$. The construction accommodates the full variety of intermediate cases where subordination ties are stronger or weaker.

3.2 Structured Forces

We have considered the structure of force in terms of a partial order on a set of force units, but we have not yet considered the content of those units. A force structure that includes one soldier per unit is quite different from one that includes two per unit, and each of these is different from one that includes one soldier and one tank per unit. In other words, the mere fact that two force structures share similar skeletons does not mean that they are the same. To that end, we now load the force units with force content.

3.5 Primitive (Structured Forces)

There is a map

$$\text{assign} : V \longrightarrow \mathbb{M}_L^*,$$

which assigns a configuration to each force unit.

A structured force is a triplet (V, \leq, assign) , where (V, \leq) is a force structure and assign is such an assignment map. The set of all structured forces is denoted

$$\mathbb{F}^*(\mathbb{M}_L^*) \subseteq \{(V, \leq, \text{assign}) : V \rightarrow \mathbb{M}_L^* \mid (V, \leq) \in \mathcal{F}\}.$$

Remarkably, a structured force is a kind of decorated graph, just like a force molecule or configuration: instead of having atoms for vertices and functional relationships for edges, we have force units for vertices and subordination relations for edges. The same logic and notion of structure applies to all three, so the intuitions honed in previous sections continue to apply.

In a sense, $\mathbb{F}^*(\mathbb{M}_L^*)$ represents the final summit of the mountain: it is nothing short of the set of all ways to put together an organized force out of pieces of force content comprised of the elements of force. One cannot help but be overwhelmed by $\mathbb{F}^*(\mathbb{M}_L^*)$: humiliating in its complexity, startling in its comprehensiveness, beautiful in its variety, terrifying in its potential. And yet, the structure we've designed for \mathbb{M}_L^* will help us understand the structure of $\mathbb{F}^*(\mathbb{M}_L^*)$. For example, structured forces can be enumerated.

3.6 Proposition (The Set of All Structured Forces)

$\mathbb{F}^*(\mathbb{M}_L^*)$ is countable.

[*Proof.*]

This is an immediate consequence of the countability of \mathcal{F} and the countability of \mathbb{M}_L^* (Lemma 2.3). As a consequence of this fact, any topology on $\mathbb{F}^*(\mathbb{M}_L^*)$ will be second-countable, which will prove most useful at the climax of our journey. But, a few final steps remain before we can fully savor the view from the summit.

Figures 9 and 10 provide the (V, \preccurlyeq) part of the force structure, but they do not provide the assignment map assign . We now decorate the force structure with the content of each force unit, taking a *tactical-units* approach. Let us first establish the basic molecule of force content as the *legionary*, which we will denote by $M \in \mathbb{M}_L$ —that is, each M represents one standard legionary, which is some connected graph where a *SOLDIER* is equipped with some appropriate panoply. By a *century* we will mean eighty such legionaries,

$$C = \biguplus_{i=1}^{80} M_i,$$

where each M_i is a legionary. We relabel the *SOLDIER* part of M_1 with the special label *CENTURION*, where this recognizes the special role of the commanding centurion in the century. Thus, for Centuries $j \in \{1, \dots, 6\}$, we have

$$\text{assign}(\text{Century } j) = \biguplus_{i=1}^{80} M_{(j-1)80+i},$$

which assigns a set of eighty unique legionaries to each century. This fills in the smallest part of the force structure, the century, with the appropriate content of eighty legionaries with a special centurion. Next, maniples are composed of three centuries. The most straightforward way to proceed is simply to use \uplus to combine the content of the three centuries, so that Maniple $k \in \{1, 2, 3\}$ is assigned the content of three centuries,

$$\text{assign}(\text{Maniple } k) = \biguplus_{j=1}^2 \text{assign}(\text{Century } 2(k-1) + j).$$

This is a simple way to combine the content of the centuries into the content of the maniples; we will consider another approach in the next subsection. The cohort is composed of three maniples: for each cohort $m \in \{1, \dots, 10\}$, we have

$$\text{assign}(\text{Cohort } m) = \biguplus_{k=1}^3 \text{assign}(\text{Maniple } 3(m-1) + k).$$

Finally, the legion is composed of ten cohorts: for each legion $n \in \{\text{I}, \dots, \text{XIII}\}$, we have

$$\text{assign}(\text{Legion } n) = \biguplus_{m=1}^{10} \text{assign}(\text{Cohort } 10(n-1) + m).$$

This completes the assignment map for the Roman legion when we think of things in terms of tactical units, which simply combines assign and \uplus in a \preccurlyeq -respecting way.

The construction above suggests a general approach for the tactical-units method of decorating a force structure.

3.7 Construction (Tactical-Units Approach to **assigning**)

The tactical-units approach to assigning is a method of decorating a force structure (V, \preccurlyeq) with configurations.

1. Define the bottom units,

$$\text{bottom} := \{v \in V \mid \nexists z \in V \setminus \{v\} \text{ such that } z \preccurlyeq v\},$$

which are the units with no distinct subordinates.

2. Define some primitive $\widetilde{\text{assign}} : \text{bottom} \rightarrow \mathbb{M}_L^*$ that assigns a configuration to each bottom unit.

3. Define the subordination relation $\text{sub} : V \rightarrow \mathbb{P}(\text{bottom})$, which assigns to each unit the set of bottom units subordinate to it:

$$\text{sub}(v) := \{z \in \text{bottom} \mid z \preccurlyeq v\}.$$

4. Define the assignment map $\text{assign} : V \rightarrow \mathbb{M}_L^*$ by the formula

$$\text{assign}(v) := \bigcup_{z \in \text{sub}(v)} \widetilde{\text{assign}}(z).$$

Since \preccurlyeq is a partial order, the operation is well-defined.

The tactical-units approach is a general method for decorating a force structure with content from the bottom up. It should be noted here that one could use some kind of operation other than \uplus to combine the content of the subordinates, and in particular one could relax Commutativity to retain some notion of order in the combination. This order would allow us to put the six centurions of a maniple in a particular order, for example, reflecting seniority and within-officer-class rank. Similarly, one could attach, as a final step in the construction, a target configuration to the output of the assignment map, so that the final output is a function of the bottom elements, but not necessarily directly through \uplus . Thus, the combination of two separate sets of eighty separate legionaries could result in a configuration with the legionaries in some specified order, perhaps reflecting a particular military formation. Surely other similar refinements are possible, but the basic structure of the tactical-units approach is as described above. It helps us appreciate the nested structure of force, the boxes upon boxes upon boxes, the subtle organizational dynamics one might not notice when scanning the forces standing on either side of a battlefield.

3.3 Organization is Organized

Little has changed; at the time of this writing, the United States Army takes as its basic tactical unit the brigade combat team (B.C.T.). Each is organized into a number of battalions, which in turn are organized into companies. The force structure for an armored B.C.T. is given in Table 4.

Company	Personnel	A.V.S	Trucks			
			L	M	H	Misc.
<i>Brigade H.Q.</i>						
H.Q. Company	137	4	29	6	0	0
<i>Battalion Total</i>	137	4	29	6	0	0
<i>Field Artillery</i>						
H.Q. Battery	233	19	33	10	0	0
Field Artillery ($\times 3$)	91	14	7	1	6	0
<i>Battalion Total</i>	506	61	54	13	18	0
<i>Cavalry</i>						
H.Q. Troop	116	17	14	6	0	0
Cavalry Troop ($\times 3$)	94	17	1	1	0	0
Armor Company	63	15	2	1	0	0
<i>Battalion Total</i>	461	83	19	10	0	0
<i>Infantry</i>						
H.Q. Company	177	25	19	5	1	0
Rifle Company ($\times 2$)	137	15	2	1	0	0
Armor Company	64	15	2	1	0	0
<i>Battalion Total</i>	515	70	25	8	1	0
<i>Armor I</i>						
H.Q. Company	176	25	19	5	1	0
Rifle Company	137	15	2	1	0	0
Armor Company ($\times 2$)	64	15	2	1	0	0
<i>Battalion Total</i>	515	70	25	8	1	0
<i>Armor II</i>						
H.Q. Company	176	25	19	5	1	0
Rifle Company	137	15	2	1	0	0
Armor Company ($\times 2$)	64	15	2	1	0	0
<i>Battalion Total</i>	515	70	25	8	1	0

Company	Personnel	A.V.S	Trucks			
			L	M	H	Misc.
<i>Brigade Engineer</i>						
H.Q. 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
<i>Battalion Total</i>	449	33	59	19	8	35
<i>Brigade Support</i>						
H.Q. 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 F.S.	154	5	14	14	21	1
Cavalry F.S.	111	6	7	9	7	1
Infantry F.S.	147	6	12	12	21	1
Armor I F.S.	147	7	12	12	21	1
Armor II F.S.	147	7	12	12	21	1
Brigade Engineer F.S.	141	3	14	14	13	1
<i>Battalion Total</i>	1,272	46	112	113	180	14
<i>Brigade Combat Team Total</i>	4,222	437	348	185	209	49

Table 4: Force structure for the United States Army Armored Brigade Combat Team. “F.S.” denotes a forward support company. “A.V.” denotes the number of armored vehicles. “L,” “M,” and “H” denote the number of light, medium, and heavy trucks, respectively. “Misc.” denotes miscellaneous vehicles. Data courtesy the Congressional Budget Office (2021).

Like legions and corps before them, B.C.T.s reflect the state’s organization of its force. The brigade is broken into battalions, which are broken into companies; not shown are even smaller units, such as platoons and squads. Battalions are assigned their own specialties: infantry, armor, cavalry, field artillery, and so on. They also receive particularized support from engineer, signal, and intelligence companies, as well as from medical, maintenance, and distribution companies. The line from Cæsar to Napoleon runs through modern forces, too.

Unit	#	Personnel (per unit)		Cost (per unit)	
		Direct	Total	Direct	Total
<i>Department of the Army</i>					
Armored B.C.T.s	12	4,040	16,330	690	3,160
Stryker B.C.T.s	7	4,680	16,670	600	3,060
Infantry B.C.T.s	13	4,560	15,910	580	2,920
<i>Department Total</i>	–	140,520	519,480	20,020	97,300
<i>Department of the Navy</i>					
Aircraft Carriers	11	3,360	6,600	620	1,470
Carrier Air Wings	9	1,750	4,880	440	1,140
Destroyers	72	350	710	80	180
Submarines	53	200	400	100	190
Amphibious Ships	33	750	1,480	160	360
Marine Infantry Battalions	24	1,900	6,320	200	990
<i>Department Total</i>	–	158,860	389,360	31,920	95,990
<i>Department of the Air Force</i>					
Fighter Squadrons	41	420	1,260	80	270
Bomber Squadrons	3	1,360	4,790	450	1,200
Cargo Squadrons	16	500	1,510	110	330
Tanker Squadrons	28	560	1,920	140	430
Reaper Squadrons	23	380	1,020	70	220
<i>Department Total</i>	–	53,720	167,410	11,920	37,050
<i>Total</i>	–	353,100	1,076,250	63,860	230,340

Table 5: Brigades in the United States Armed Forces. Costs measured in millions of 2021 dollars.

Like Roman forces, American forces demonstrate a strongly fractal nature: in 2021 the United States Army commanded twelve armored B.C.T.s, each costing \$3.2 billion to maintain—and it is but one of the many units that make up the United States Armed Forces. Table 5 provides yet another ledger for our grim accountant to tally, this time in terms of annual costs in both personnel and dollars. In 2021, the United States Armed Forces coordinated over one million personnel and cost of over \$230 billion to maintain—and this is merely in terms of the units themselves, not to mention the procurement of equipment, the training of personnel, and the myriad other costs associated with the military-industrial complex. Total military spending in 2021 was over \$700 billion, or roughly 3% of the United States’s gross domestic product of \$23 trillion.

The force structure in Table 4 bespeaks two important differences from the Roman legion.

1. *Formal headquarters.* The b.c.t. has a formal headquarters, staffed by a company of personnel. It has one company's worth of personnel, along with a light outfit of armored vehicles and light trucks. But observe also that each of the brigade's constituent battalions has a formal headquarters, again staffed by a company of personnel. The table only goes as granular as a company, but it turns out that modern army companies also have formal headquarters, and indeed their own tactical units—the platoons—have headquarters, too. Now, the Roman legion had a headquarters, as well, and even a lowly centurion might have an assistant. But the overall structure of the Roman legion suggested that the chain of command was as much about conveying information and basic tactical coordination as it was about bureaucracy and administration. The development of headquarters is a crucial aspect of the revolution in military affairs. For example, in his titanic *Command in War*, historian Martin van Creveld describes Napoleon's high-level Imperial Headquarters as itself a highly-structured, well-outfitted unit broken into independent subdivisions equipped with large, diverse staffs 1985, pp. 65–78, but the lower-level units were not so well-structured. This reflected Napoleon's preference for unified command and direct access to information, but it also reflected the limitations of the time: the technology of the day did not allow for the kind of real-time communication that we take for granted today.

Deeper headquarters structures reflect developments in the officer corps. Two stodgy classics in civil-military relations, Samuel P. Huntington's *The Soldier and the State* (1957) and Morris Janowitz's *The Professional Soldier* (1960) trace the development of the officer corps in the United States military. Huntington, taking a much longer view, argues that the officer corps has evolved from a warrior aristocracy to a professional bureaucracy; the goal of this development is to ensure that the military remains subordinate to objective civilian control. Famously, this is handled by taking the pragmatic mind of the soldier and imbuing it with the qualities of professionalism: expertise, responsibility, and corporateness. The end result is a force that is more effective not only in combat, but also in the broader context of civil-military relations. Janowitz paints a much different picture, arguing that modern officers are less like lawyers and more like the police: skilled in organizing and administering violence in the service of the state. In either case, the density of the officer corps is purpose-driven and reflects the needs of the state.

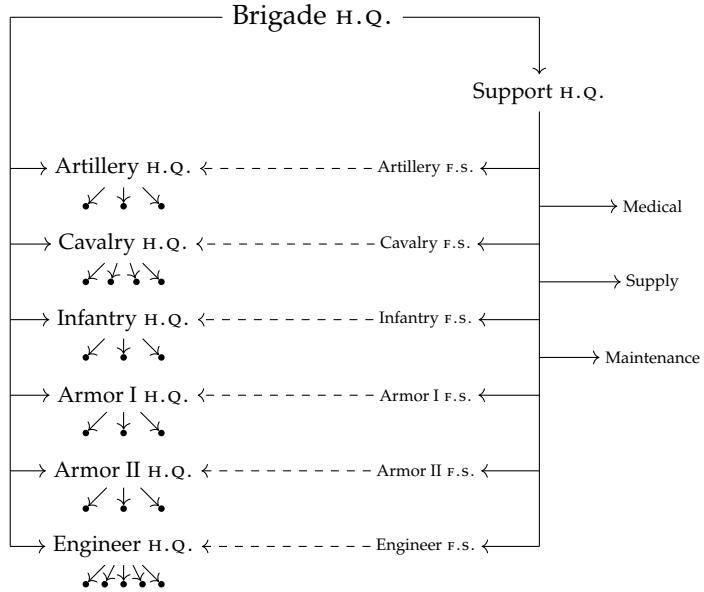


Figure 11: A binary support relationship. Solid lines represent superordination, while dashed lines represent support. Black nodes represent other companies.

2. *Formal non-command relationships.* The B.C.T. has a complete battalion dedicated to support. Some of its companies provide general support to the entire brigade, such as the Medical Company, while others provide specialized field support to particular combat battalions, such as the Cavalry F.S. Company. The development of such formalized support relationships reflects an underappreciated aspect of the revolution in military affairs—namely the revolution in military logistics.⁴⁴ This less-appreciated revolution has been the subject of a spate of recent research (e.g., ?Bury, 2021), much of it inspired again by a canonical treatment due to van Creveld (2004), this time his *Supplying War*. Van Creveld is unimpressed by the role of hyper-organized logistics in the outcome of wars, and indeed he argues that the ability to improvise is a key aspect of logistical success. By linking combat units to support units, the B.C.T. ensures that the right supplies are in the right place at the right time. This suggests that a force structure ought to include relationships that are not strictly about command and subordination but also about support, training, joint operations, the sharing of resources, and so on.

⁴⁴Here we mean the part of logistics pertaining to the delivery of supplies, rather than their procurement. This latter aspect is also important and is taken up in a companion paper.

The first of these differences suggests a different interpretation of décor where higher-level units are concerned.

3.8 Construction (Strategic Approach to assigning)

The strategic approach to assigning is a method of decorating a force structure (V, \preccurlyeq) . Bottom nodes are interpreted as the largest units without formal administration; each is assigned a configuration meant to represent the unit's tactical capabilities. Those units higher in the chain of command are assigned configurations representing administrative staff. The assignment map is simply the function giving bottom units their tactical capabilities and nonbottom nodes their staff configurations.

The strategic approach is a general method for decorating a force structure from the top down. It is a more modern approach, reflecting the development of formal headquarters and support relationships as evidenced in the force structure in Table 4. The reader is advised to use the tactical-units approach to aggregate bottom units to some desired level of granularity, and then to use the strategic approach to populate the various administrative units.

The second of these differences is the more subtle, and it is sufficiently beautiful that it's worth pausing to appreciate. It is not difficult to think about binary non-command relationships in the context of a force structure—the field support companies in Figure 11 are a good example. But some relationships are far more difficult to visualize.

1. Consider how we might model the Medical Company's support relationships, where there is not a dedicated target company (as was the case for field support companies) but rather a set of companies that might need medical support. We could certainly reduce the problem to a set of binary relationships, but this would be a poor model of the actual relationships. For example, we would not immediately have information at hand about how overtaxed the Medical Company is, as we'd need some operation for adding the individual burdens of support.
2. More ambitiously, consider what's going on at a given military base, where the base command might not be directly involved in the chain of command for the units stationed there. What is the base commander's relationship to the units? What is the relationship between the units themselves? What of the base commander's subordinates, who presumably serve a variety of higher-order functions themselves?

It is not hard to consider other examples, but the point is that the relationships in a force structure are not always so simple as the chain of command. And of course, various sorts of relationships can be combined in a single force structure.

The problem posed above suggests the need for a more general approach to decorating a force structure, which requires yet another layer of abstraction.

3.9 Primitive (Non-Command Relations)

There is a countable family

$$\mathbf{Rel} := \{\mathcal{R}^r\}_{r \in \mathbb{N}}$$

of non-command relations, where each \mathcal{R}^r is a directed hypergraph on V —i.e.,

$$\mathcal{R}^r := \{\mathbf{R}_i^r := (D_i^r, C_i^r)\}_{i \in I^r},$$

where $\emptyset \subseteq D_i^r, C_i^r \subseteq V$ are the domain and codomain of the i th hyperedge for relation of sort $r \in \mathbb{N}$, and I^r is the finite index set of hyperedges in \mathcal{R}^r .

The non-command relations are a general method for decorating a force structure with relationships that are not strictly about command and subordination. A few examples might help to illustrate the concept:

1. The Medical Company's job is *one-to-many*, as it must support a variety of units, so that

$$D^{\text{med}} = \{\text{Medical Company}\},$$

$$C^{\text{med}} = \{\text{the various units it supports}\}.$$

One could imagine enriching the binary support relationships for field-support companies in a similar way.

2. The base commander's job is *many-to-many*, as the base commander might coordinate with commanding officers, so that one relationship might be

$$D_1^{\text{base}} = \{\text{Base Commander, Commanding Officer 1}\},$$

$$C_1^{\text{base}} = \{\text{the various subunits in Unit 1}\};$$

and another might be

$$D_2^{\text{base}} = \{\text{Base Commander, Commanding Officer 2}\},$$

$$C_2^{\text{base}} = \{\text{the various subunits in Unit 2}\}.$$

This seems especially useful for coalition forces, where commanders might need to coordinate with a variety of other commanders. It is not hard to imagine a laughably-complicated org chart for a coalition force, where the base commander is coordinating with a variety of other commanders, who are in turn coordinating with their own subordinates. The reader is encouraged to sharpen some crayons—a *lot* of crayons—and draw one.

3. Of course, one could model non-directed relationships, as well, and this can be done at any desired order. A given three-way relationship—say, three units that are coordinating on a joint operation—might be

$$(D^{\text{joint}_3}, C^{\text{joint}_3}) = (\{v_1, v_2, v_3\}, \{v_1, v_2, v_3\}),$$

suggesting that the relationship is fully characterized by an undirected ternary relationship between the three units. In case each \mathbf{R}_i^r features $|I_i^r| = 1$ and $|D_i^r| = |C_i^r|$ for all $r \in \mathbb{N}$, we say **Rel** is *undirected*.

4. Remarkably, the construction as stated provides an inexpensive theoretical way to parameterize not just the relationships in question but the units themselves. After all,

$$(D^{\text{id}_v}, C^{\text{id}_v}) = (\{v\}, \{v\})$$

is a perfectly valid non-command relationship, and it is not hard to imagine that the units themselves might be decorated in this way. For some parameter $\theta_v \in \Theta_v$ featuring $|\Theta_v| \leq N_0$, we might choose a single hyperedge $\mathbf{R}_{\theta_v}^{\text{id}_v}$ for each $v \in V$. This indeed provides a natural way to parameterize the units themselves by including one or more such hyperedges in **Rel**.

These examples should help provide a sense of the power of the non-command relations: we have at our disposal a countable family of relationships, each potentially of a different sort and density, that can be used to decorate a force structure in a variety of ways.

Let us formally define the set of all possible relation families for a given V . The set of hypergraphs on V is

$$\mathcal{H}_V := \{(D_i, C_i)_{i \in I} \mid 0 < |I| \leq N_0 \text{ and } \emptyset \subseteq D_i, C_i \subseteq \underline{n} \text{ for all } i \in I\};$$

since V is finite, so too is \mathcal{H}_V , as there are only finitely many ways of pulling out two subsets of V . The set of all possible non-command relation families is then given by

$$\mathcal{Z}_V := \left\{ \{\mathcal{R}^r\}_{r \in \mathbb{N}} \mid \mathcal{R}^r \in \mathcal{H}_V \text{ for all } r \in \mathbb{N} \right\} \cong \mathbb{N}^{\mathcal{H}_V} = \mathbb{N}^{N_0} \cong \mathbb{N}.$$

In words, we have shown that the set of hypergraphs on any V is finite and thus the set of ways of giving each natural number a hypergraph is countable. Since V was chosen arbitrarily, we may conclude:

3.10 Lemma (Countability of Non-Command Relations)

For all finite V , the set of non-command relations on V is countable.

Wonders never cease.

We can now define what we mean by an org chart.

3.11 Construction (Org Charts)

By an org chart on V we mean a triplet

$$\mathbf{Org}_V := (V, \preccurlyeq, \mathbf{Rel}),$$

where (V, \preccurlyeq) is a force structure and $\mathbf{Rel} \in \mathcal{Z}_V$ is a non-command relation family.

The org chart is a generalization of the force structure, one that includes relationships that are not strictly about command and subordination. It can encode more detailed unit- and relation-level responsibilities, not to mention the wide variety of logistical, support, informational, and other relationships that are necessary for a modern force to function. The org chart is a powerful tool for understanding the structure of a force, and it is a natural way to model the relationships that are not strictly about command and subordination. Indeed, we've even seen that it can point us toward parameterization on deeper levels.

Conceiving of a structured force as a set of units, a subordination relation, an assignment mapping, and a constellation of non-command relationships, we arrive at the following.

3.12 Corollary (The Relational Variant of Structured Force is Countable)

The set of all structured forces with non-command relations,

$$\mathbb{F}_R^*(M_L^*) \cong \mathbb{F}^*(M_L^*) \times \mathcal{Z}_V,$$

is countable.

This is a corollary of the countable nature both of structured forces $\mathbb{F}^*(M_L^*)$ and of non-command relations \mathcal{Z}_V , which we established in Proposition 3.6 and Lemma 3.10, respectively.⁴⁵ The countability of the structured forces is a remarkable fact, given the complexity of the force structure: we now have a

⁴⁵The reader can rest assured that this is the *final* countability result in this manuscript. This is not a manuscript about countability but rather about force, but it happens that countability is a useful property for cheaply obtaining second-countability in the spaces we are considering. It also happens to be beautiful that countability survived the trip:

$$L \leadsto M_L \leadsto M_L^* \leadsto \mathbb{F}^*(M_L^*) \leadsto \mathbb{F}_R^*(M_L^*),$$

or in words, from elements to molecules to configurations to structured forces to structured forces with non-command relationships.

complete theory of military configuration and organization that is easy to work with. This is the force-maker's possibilities space, and it is a beautiful thing. We now turn our attention to how she navigates this space by turning one structured force into another, a process running parallel to the conversion of configurations.

3.4 Restructuring Forces

We have constructed a theory of structured forces where some set of units is organized by both a chain of command and a set of non-command relationships; moreover, each of these units is assigned a configuration from \mathbb{M}_L^* . We now turn our attention to how to turn one of these structured forces into another, first with respect to the org chart and then with respect to the configurations. Though still “elementary” in the sense that they do not depend on operations beyond set theory, these operations are more complex than those for configurations. Even the simplest-possible operations are not trivial:

3.13 Primitive (Elementary Restructuring Operations for Units)

Given a structured force

$$(V_0, E_0, \mathbf{Rel}_0, \text{assign}_0) \in \mathbb{F}_R^*(\mathbb{M}_L^*) ,$$

where V_0 is the initial set of units, E_0 is the edge set associated with the initial chain of command, \mathbf{Rel}_0 is the initial non-command relation family, and assign_0 is the initial assignment map, we define two elementary unit restructuring operations on V :

1. Unit Addition at v : for $v \notin V_0$, set

$$V_1 := V_0 \cup \{v\} ,$$

$$E_1 := E_0 ,$$

$$\mathbf{Rel}_1 := \mathbf{Rel}_0 , \text{ and}$$

$$\text{assign}_1 := \text{assign}_0 \cup \{(v, \bigcup \mathcal{M})\} \text{ for some } \bigcup \mathcal{M} \in \mathbb{M}_L^* ; \text{ and}$$

2. Unit Deletion at v : for $v \in V_0$, set

$$V_1 := V_0 \setminus \{v\} ,$$

$$E_1 := E_0 \setminus (\{(v, w) \mid w \in V_0\} \cup \{(w, v) \mid w \in V_0\}) ,$$

$$\mathbf{Rel}_1 \text{ to be discussed in the sequel, and}$$

$$\text{assign}_1 := \text{assign}_0 \setminus \{(v, \text{assign}_0(v))\} .$$

Unit Addition is similar to Vertex Addition for force molecules (Definition 2.11), so we need not say much about it; the original intuitions ought to carry over. It simply adds a new unit and assigns it some configuration from \mathbb{M}_L^* . The resulting structured force is simply the old one plus some isolated unit with some name and some configuration. The output, though potentially still incomplete regarding command and non-command relationships, is a useful intermediate step in a micro-stylized restructuring process, a sequence of steps so small as to appear meaningless, the operational atoms of the verb “to restructure.”

But Unit Deletion is a different beast. To be sure, it is not very difficult to remove a vertex from a graph, nor even to remove its contributions to the subordination relationships. This is for two reasons:

1. *Subordination is a binary relation.* The statement " $v_1 \leq v_2$ " encodes all of the information we need to know about the subordination relationship between v_1 and v_2 , and (what is more) this relationship encodes no additional meaning one could have gotten for unrelated " $v_3 \leq v_4$ " statements. These statements all provide independent data about the subordination relationships in the force structure.
2. *Subordination is a partial order.* Since the subordination relation is a partial order, removing a vertex and all the edges it touches produces another partial order. In other words, the set of structured forces is closed under the removal of a unit and all its incident edges.

But the non-command relationships are a different story. Recall that a given **Rel** is a countable family of hypergraphs, each a finite set of directed multi-unit relationships. One hypergraph might be labeled " M " and it would contain each of the myriad relationships between the myriad Medical Companies and their myriad charges. Suppose we planned to remove one of those Medical Companies, call it i . Suppose that Medical Company i coordinated with an Intelligence Company to support an Infantry Company, so that one of the hyperedges in the Medical Company's hypergraph was

$$\mathbf{R}_i^M = (\{\text{Medical } i, \text{Intelligence}\}, \{\text{Infantry}\}) \in \mathcal{R}^M \in \mathbf{Rel}.$$

How do we proceed? Three possibilities suggest themselves:

1. *Remove \mathcal{R}^M from **Rel**:* perhaps removing a single unit from a single hypergraph is enough to warrant scrapping all information once the unit is removed. This would suggest a deep and nuanced web of relationships *among the relationships* in the Medical Support substructure.
2. *Remove \mathbf{R}_i^M from \mathcal{R}^M :* perhaps removing a single unit is "small" in the grand picture but important enough that no relationship that had included that unit should be retained. In this motivating example, we'd remove R^M , since the associated Intelligence-Infantry relationship need not feel "adjacent" to the original one.
3. *Remove Medical i from \mathbf{R}_i^M :* or, if this relationship does, in fact, feel adjacent, we might remove only the unit's contributions to the hypergraph, leaving the rest of the simpler-order relationships intact.

Frankly, all three of these possibilities seem to have their place.

1. The first is a sort of nuclear option, where the removal of a single unit is enough to warrant the removal of all relationships that had included that unit, even seemingly-unrelated relationships that are same-in-kind. It could work here were the force-maker occupied with how evenly-distributed the support burden was across the Medical Companies. This option seems most appealing when considering the removal of a high-level command web in a coalition force, where the removal of a single unit might warrant the removal of all relationships that had included that unit, even if they were with other coalition forces. *Everything* is different in the absence of that unit, and the removal is large in a global sense.
2. The second, a predictable middle ground, is where the removal of a single unit is enough to warrant the removal of all relationships that had included that unit, even if other units remained to maintain the relationship. This is certainly the most relevant option for our motivating example, since the resulting Intelligence-Infantry relationship does not feel adjacent to the original Medical-Intelligence-Infantry relationship. The removal is large in a local sense but small in a global sense.
3. The third is a scalpel, where the removal of a single unit is enough to warrant the removal of only the relationships that had included that unit, leaving the rest of the relationships intact. This is the most appealing option when the force-maker is concerned with the overall structure of the force, and the removal is small in a local sense.

As such, we encode all three possibilities in the definition of Unit Deletion.

3.14 Primitive (Classes of Unit Deletion)

We define three classes of Unit Deletion for unit $v \in V$:

1. Global Deletion: *remove all relationships that had included the unit*,

$$\mathbf{Rel}_1 = \{\mathcal{R}_0^r \in \mathbf{Rel}_0 \mid \nexists \mathbf{R}_{i0}^r \in \mathcal{R}_0^r \text{ such that } v \in D_{i0}^r \cup C_{i0}^r\};$$
2. Hyperedge Deletion: *remove all hyperedges that had included the unit*,

$$\mathcal{R}_1^r := \{\mathbf{R}_{i0}^r \in \mathcal{R}_0^r \mid v \notin D_{i0}^r \cup C_{i0}^r\} \text{ for all } r \in \mathbb{N}; \text{ and}$$
3. Surgical Deletion: *remove the unit from all hyperedges that had included it*,

$$\mathcal{R}_1^r := \{(D_{i0}^r \setminus \{v\}, C_{i0}^r \setminus \{v\}) \mid \mathbf{R}_{i0}^r \in \mathcal{R}_0^r\} \text{ for all } r \in \mathbb{N}.$$

Thus, our technical conundrum has proven substantively rich.

Primitive 3.13 provides simple generalizations of vertex addition and deletion to the structured forces. They are elementary in the sense that they do not use any other operations save for those made available from basic set theory. We also need to generalize edge addition and deletion, both for subordination relationships and for non-command relationships. We begin with the former.

3.15 Primitive (Elementary Restructuring Operations for Subordination)

Given a structured force

$$(V_0, E_0, \mathbf{Rel}_0, \mathbf{assign}_0) \in \mathbb{F}_R^*(\mathbb{M}_L^*) ,$$

we define two elementary subordination restructuring operations on E :

1. Edge Addition at (v, w) : for $v, w \in V_0$ with $(v, w), (w, v) \notin E_0$, set⁴⁶

$$V_1 := V_0,$$

$$E_1 := (E_0 \cup \{(v, w)\})^+$$

$\mathbf{Rel}_1 := \mathbf{Rel}_0$, and

$\mathbf{assign}_1 := \mathbf{assign}_0$; and

2. Edge Deletion at (v, w) : for $(v, w) \in E_0$, set

$$V_1 := V_0,$$

$$E_1 := (E_0 \setminus \{(v, w)\})^+$$

$\mathbf{Rel}_1 := \mathbf{Rel}_0$, and

$\mathbf{assign}_1 := \mathbf{assign}_0$.

Edge Addition is a simple generalization of the vertex addition operation, and it is a natural way to add a new subordination relationship to a force structure. The resulting structured force is simply the old one plus a new subordination relationship, where particular care is taken to ensure that the output remains a partial order. Edge Deletion is a simple generalization of the vertex deletion operation, and it is a natural way to remove a subordination relationship from a force structure; again, we take small pains to ensure that the output remains a partial order. Overall, these two generalizations pose no new challenges, and they are a natural way to extend the elementary edits to the structured forces. Little more need be said about them, but they remain useful tools in the kit.

⁴⁶Here we use the notation E^+ to denote the transitive closure of a relation E —i.e., the smallest transitive relation containing E . This ensures that the new edge is added in a way that respects the existing subordination relationships. Defining Edge Addition at (v, w) only when $(v, w), (w, v) \notin E_0$ ensures that the output satisfies the antisymmetry property of a partial order.

Next, we generalize the elementary edits to the non-command relationships. The nuances discussed in the run-up to Primitive 3.14 will persist, and the intuitions developed there will carry over.

3.16 Primitive (Elementary Restructuring Operations for Non-Command)

Given a structured force

$$(V_0, E_0, \mathbf{Rel}_0, \text{assign}_0) \in \mathbb{F}_R^\star(\mathbb{M}_L^\star),$$

we define eight elementary non-command restructuring operations on \mathbf{Rel} :

1. Global Addition of \mathcal{R} : set

$$\mathbf{Rel}_1 := \mathbf{Rel}_0 \cup \{\mathcal{R}\};$$

2. Hyperedge Addition of i at r : set

$$\mathbf{Rel}_1 := \{\mathcal{R}_0^s \in \mathbf{Rel}_0 \mid s \neq r\} \cup \{\mathcal{R}_0^r \cup \{(D_i^r, C_i^r)\}\};$$

3. Surgical Domain Addition of v at (r, i) : set

$$\mathbf{Rel}_1 := \{\mathcal{R}_0^s \in \mathbf{Rel}_0 \mid s \neq r\} \cup \{\mathcal{R}_0^r \cup (D_{i0}^r \cup \{v\}, C_{i0}^r)\};$$

4. Surgical Codomain Addition of v at (r, i) : set

$$\mathbf{Rel}_1 := \{\mathcal{R}_0^s \in \mathbf{Rel}_0 \mid s \neq r\} \cup \{\mathcal{R}_0^r \cup (D_{i0}^r, C_{i0}^r \cup \{v\})\};$$

5. Global Deletion of \mathcal{R} : for some $\mathcal{R} \in \mathbf{Rel}_0$, set

$$\mathbf{Rel}_1 := \mathbf{Rel}_0 \setminus \{\mathcal{R}\};$$

6. Hyperedge Deletion of i at r : set

$$\mathbf{Rel}_1 := \{\mathcal{R}_0^s \in \mathbf{Rel}_0 \mid s \neq r\} \cup \{\mathcal{R}_0^r \setminus \{(D_{i0}^r, C_{i0}^r)\}\};$$

7. Surgical Domain Deletion of v at (r, i) : set

$$\mathbf{Rel}_1 := \{\mathcal{R}_0^s \in \mathbf{Rel}_0 \mid s \neq r\} \cup \{\mathcal{R}_0^r \cup (D_{i0}^r \setminus \{v\}, C_{i0}^r)\}; \text{ and}$$

8. Surgical Codomain Deletion of v at (r, i) : set

$$\mathbf{Rel}_1 := \{\mathcal{R}_0^s \in \mathbf{Rel}_0 \mid s \neq r\} \cup \{\mathcal{R}_0^r \cup (D_{i0}^r, C_{i0}^r \setminus \{v\})\}.$$

These are similar in spirit to the edits we considered when deleting a unit, but here the motivations are institutional rather than unit-level. They provide a flexible set of ways to restructure the myriad possible non-command relationships in a structured force, including nuclear, middle-ground, and scalpel-like options. Since these relationships need not follow any pattern—say, the antisymmetry of a partial order—we have more freedom in how we edit them.

Primitives 3.13 to 3.16 provide the tools required to manipulate org charts; one can imagine the force-maker at her pin-board, moving units around, adding and removing relationships, and generally reorganizing the force. Any org chart can be obtained from any other via a finite sequence of these operations. But though the force-maker might be working up a sweat moving pins around, the soldiers have yet to stir from their barracks. Save for mustering up a configuration for a new unit (or firing the staff of a deleted unit), the force-maker has yet to do anything that would affect the soldiers themselves. We now turn our attention to “relabeling” the configurations in a structured force. Where once relabeling represented the change of one element of force to another—say, a spear to a sword—it now represents the change of one configuration to another—say, one combination of personnel, armored vehicles, trucks, and so on to another. At the level of the structured force, *these* are the atoms, and the variety of configurations is elemental in the same way that the variety of elements was.

For all the nuances implicit in such wide varieties of configurational reassignment, the graph-theoretic character of the structured forces provides starting points for the discussion. Indeed, reassignment *could* be quite easy:

3.17 Primitive (Localized Reassignment)

Given a structured force

$$(V_0, E_0, \mathbf{Rel}_0, \text{assign}_0) \in \mathbb{F}_R^*(\mathbb{M}_L^*),$$

we define the localized reassignment operation by:

$$\text{assign}_1 := (\text{assign}_0 \setminus \{(v, \text{assign}_0(v))\}) \cup \{(v, \biguplus \mathcal{M})\}$$

for some $v \in V_0$ and $\biguplus \mathcal{M} \in \mathbb{M}_L^$.*

This is a simple operation that changes the configuration of a single unit. It is “local” because it introduces no new information or resources from elsewhere in the force structure: all we’ve done is gone through the conversion process

$$\begin{aligned} \text{assign}_0(v) &\xrightarrow{\phi} \text{assign}_1(v), \\ \phi &\rightsquigarrow \mathbf{CS}(\text{Hom}(\text{assign}_0(v), \text{assign}_1(v))), \end{aligned}$$

where ϕ is the chosen process from those available for converting $\text{assign}_0(v)$ to $\text{assign}_1(v)$. In this sense, localized reassignment is less about reassignment and more about converting configurations without any configurational input from the rest of the force. It is thus quite useful for small changes, say in preparedness, organization, or equipment, that do not affect the rest of the force. We’ve not yet mentioned the costs of restructuring, but it stands to reason that the cost of such a conversion is simply $(\text{cost} \circ \phi)(\text{assign}_0(v), \text{assign}_1(v))$.

As such, we might consider more nuanced operations involving more than one unit. It will be easier to think things through if we define a few operations on force configurations. First comes a notion of subconfigurationhood, which we think of moleculewise.

3.18 Definition (Subconfigurationhood)

Given two configurations

$$\biguplus \mathcal{M}_1 = \biguplus_{i=1}^{n_1} M_i \quad \text{and} \quad \biguplus \mathcal{M}_2 = \biguplus_{i=1}^{n_2} M_i \in \mathbb{M}_L^{\star},$$

we say that $\biguplus \mathcal{M}_1$ is a subconfiguration of $\biguplus \mathcal{M}_2$, denoted $\biguplus \mathcal{M}_1 \preccurlyeq \biguplus \mathcal{M}_2$, just in case there exists an injection $\iota : \{1, \dots, n_1\} \hookrightarrow \{1, \dots, n_2\}$ such that $M_i \cong M_{\iota(i)}$ for all $i \in \{1, \dots, n_1\}$.⁴⁷

In words, one configuration is a subconfiguration of another if the first can be obtained by removing some of the molecules from the second.⁴⁸ The molecules in the second might not be in the same order as those in the first, but they are the same molecules. Recall from Primitive 2.1 that Ψ is commutative up to isomorphism, so the case where the injection witnessing subconfigurationhood is a bijection is not particularly interesting; we already had $\biguplus \mathcal{M}_1 \cong \biguplus \mathcal{M}_2$ there simply from the commutativity of Ψ .⁴⁹ The definition is essentially a shortcut for the more primitive operation of projecting a configuration onto a subset of its molecules, and it is useful for the more complex operations we will consider.

Predictably, subconfigurationhood is a preorder on \mathbb{M}_L^{\star} .

3.19 Proposition (Subconfigurationhood is a Preorder)

Subconfigurationhood is a preorder on \mathbb{M}_L^{\star} .⁵⁰

[*Proof.*]

This is a simple result that follows from the definition of subconfigurationhood; it is useful to have this result on hand, as (for example) it allows us to speak of monotone functions on \mathbb{M}_L^{\star} with respect to subconfigurationhood.

⁴⁷ Since any function defined on the empty set is (vacuously) injective, the empty configuration is a subconfiguration of any configuration.

⁴⁸ Observe that we require only moleculewise isomorphism (Remark 1.4), meaning that the structure is the same but the order in which the atoms are declared might differ.

⁴⁹ The definition of subconfigurationhood is a bit more restrictive than mere subgraphhood, where the only requirement is that the subgraph is a subgraph of the supergraph. Here we demand that the molecules are unaffected by the reordering of the atoms, which is a stronger condition more in keeping with the spirit of the theory.

⁵⁰ In fact, in the proof it is shown that \preccurlyeq is a partial order on \mathbb{M}_L^{\star} up to permutation of the molecules, which is in keeping with the spirit of \biguplus .

Pressing on will require a few more definitions.

3.20 Construction (Counts and Intersections for Configurations)

For a force molecule

$$\biguplus \mathcal{M} = \biguplus_{i=1}^n M_i \in \mathbb{M}_L^*$$

and a molecule $M \in \mathbb{M}_L$, we define the strict and lax counts of M in $\biguplus \mathcal{M}$ by

$$\begin{aligned} \#_S(M, \biguplus \mathcal{M}) &:= |\{i \in \underline{n} \mid M_i = M\}|, \text{ and} \\ \#_L([M], \biguplus \mathcal{M}) &:= \sum_{M' \in [M]} \#_S(M', \biguplus \mathcal{M}), \end{aligned}$$

where $[M]$ is the equivalence class of M under the isomorphism relation for molecules.

We also define the intersection of $\biguplus \mathcal{M}_1$ and $\biguplus \mathcal{M}_2$ by

$$\biguplus \mathcal{M}_1 \bowtie \biguplus \mathcal{M}_2 := \biguplus_{[M] \in \biguplus \mathbb{M}_L^*/\cong} \overbrace{M \bowtie \dots \bowtie M}^{\min\{\#_L([M], \biguplus \mathcal{M}_1), \#_L([M], \biguplus \mathcal{M}_2)\} \text{ times}}.$$

where M is some representative of $[M]$.

The intersection of two configurations returns a new configuration that contains the molecules that appear in both configurations. A few quick matters of fact:

1. *\bowtie is commutative up to isomorphism of molecules:* since \min is commutative, the same number of representatives from each equivalence class will appear in the intersection, but there are no guarantees that the same representative is chosen across the two calls to \bowtie .
2. *The empty configuration is the identity for \bowtie :* since the “ \min ” operation defining intersection always returns 0, the empty configuration is the identity for \bowtie .
3. *The intersection of a configuration with itself is the configuration:* since the “ \min ” operation defining intersection always returns the exact number of times a molecule appears in the configuration, the intersection of a configuration with itself is the configuration.
4. *Subconfigurationhood and intersection:* if $\biguplus \mathcal{M}_1 \preccurlyeq \biguplus \mathcal{M}_2$, then $\biguplus \mathcal{M}_1 \bowtie \biguplus \mathcal{M}_2 = \biguplus \mathcal{M}_1$.

These all mimic set intersection, with only a little additional care about the ordering of the atoms or of the molecules.

Just as we needed to transport set intersection to its configurational variant, so too must we do the same for relative complement. We want an operation that removes those elements one set shares with another.

3.21 Construction (Relative Complement for Configurations)

We define the relative complement of $\biguplus \mathcal{M}_1$ in $\biguplus \mathcal{M}_2$ by

$$\biguplus \mathcal{M}_2 \setminus \biguplus \mathcal{M}_1 := \biguplus_{[M] \in \biguplus \mathbb{M}_L^*/\cong} \overbrace{M \uplus \dots \uplus M}^{\max\{\#_L([M], \biguplus \mathcal{M}_2) - \#_L([M], \biguplus \mathcal{M}_1), 0\} \text{ times}}.$$

where M is some representative of $[M]$.

Once again, the set complement is transported to the configurational realm, where again the only fidelity lost regards the order of atoms in molecules or molecules in configurations, both of which we've chosen to ignore.

These simple operations suffice to define a more complex operation that allows the force-maker to reassign molecules from one set of units to another. We now introduce the powerful operation of reassignment, which allows the force-maker to move molecules from one set of units to another, neither creating nor destroying them, spreading a fixed amount of force icing across a fixed number of forcecakes (arranged in some fixed order).

3.22 Primitive (Reassignment)

Given a structured force $(V_0, E_0, \mathbf{Rel}_0, \text{assign}_0) \in \mathbb{F}_R^*(\mathbb{M}_L^*)$, we define the the reassignment operation by:

1. Choose the Units Involved: choose some $V_S \subseteq V_0$ and $V_T \subseteq V_0$;
2. Melt the Sources: choose some $(\text{assign}_1(v_s))_{v_s \in V_S}$ such that $\text{assign}_1(v_s) \preccurlyeq \text{assign}_0(v_s)$ for all $v_s \in V_S$;
3. Partition the Residue: choose a sequence $(\text{pad}_1(v_T))_{v_T \in V_T}$ such that

$$\biguplus_{v_T \in V_T} \text{pad}_1(v_T) = \biguplus_{v_s \in V_S} \text{assign}_0(v_s) \setminus \text{assign}_1(v_s),$$

$$i \neq j \implies \text{pad}_1(v_i) \cap \text{pad}_1(v_j) = \emptyset; \text{ and}$$

4. Reassign: set

$$\text{assign}_1(v_s) := \text{assign}_1(v_s) \text{ for all } v_s \in V_S,$$

$$\text{assign}_1(v_T) := \text{assign}_0(v_T) \uplus \text{pad}_1(v_T) \text{ for all } v_T \in V_T.$$

This definitions packs a whallop, so we discuss some details here.

In words, the reassignment operation is simple: we choose some units to donate molecules and some units to receive them, we remove some (or all) of the molecules from the donating units, and we divide the removed molecules among the receiving units. We ignore the path the molecules take from their old home to their new one, focusing only on the before-and-after configurations of the units. In the simplest case, we have $V_S = \{v_S\}$ and $V_T = \{v_T\}$, and we simply transfer some of the molecules in the configuration $\text{assign}_0(v_S)$ to the configuration $\text{assign}_0(v_T)$. In the most complex case, we have $V_S = V_T = V$, and all units both donate and receive new molecules. Any intermediate level of complexity is possible, and this can occur for any subconfigurations of the units, partitions of the residue, and so on. Indeed, the only thing that unifies this wide set of possibilities is the way the resource constraint is derived and enforced. The force-maker can only reassign the molecules that are available, neither creating nor destroying them. This is the essence of reassignment, the bridge linking all these disparate islands, and so often the tension motivating real-world force-makers to restructure their forces in the first place.

Reassignment represents a substantively-distinct sort of operation on the structured forces than its localized counterpart, since it involves the movement of resources from one set of units to another. Localized reassignment is the sort of thing we tinkerers of the laws of physics might do to evaluate a counterfactually-possible world, while reassignment is the sort of thing a real-world force-maker might have to do to reorganize her force. She faces a constraint, namely the equality of the total amount of force in the source and target units, and she must respect this constraint in her reorganization. If we inspected her work, it's true that we tinkerers could generate the same output from the same input via a well-chosen sequence of localized reassessments alone. Nevertheless, the reassignment operation is a more natural way to think about the process of reorganization, and indeed it deserves to be thought of separately from its localized counterpart. Another way of saying the same thing: were we to think about the costs involved in the same reorganization via different operations, would we think in the same terms? Unitwise localized reassessments is a matter of iterative force conversion, so the terms of the conversation would transpire in the structure of Ξ . Genuine reassignment, on the other hand, is a matter of moving resources around, so the terms of the conversation might transpire using new concepts—or might take place between different sets of the force-maker's underlings. In other words, the question goes beyond the mere “but the costs are different” to the more nuanced “but the costs are conceived of differently, and potentially may not be directly comparable.” This is a subtle and rich vein of thought, and so we will leave it to marinate for a while; we will define a new cost calibration device and a linkage between it and Ξ in the next section.

By now it is hoped that the idea is clear: we can define a wide variety of elementary operations on structured forces that allow the force-maker to restructure her force. Surely those listed here are not exhaustive, but they are a good start—at minimum, they are sufficient in the sense that any structured force can be turned into any other by way of a finite sequence of these operations.⁵¹ In other words, we have the tools at our disposal to restructure the force in any way we see fit. But there exist a wide variety of other operations that the force-maker might wish to perform, many of which involve clever combinations of the operations we have already defined. For example:

3.23 Construction (Merging and Splitting)

Given a structured force $(V_0, E_0, \mathbf{Rel}_0, \mathbf{assign}_0) \in \mathbb{F}_R^ (\mathbb{M}_L^*)$ we define unit merging and splitting as follows:*

1. choose some $V_{old} \in V_0$ and V_{new} such that $V_{new} \cap V_0 = \emptyset$;
2. apply Unit Addition at V_{new} ;
3. apply Reassignment from V_{old} to V_{new} in such a way that $\mathbf{assign}_1(v_{old}) = \emptyset$ for all $v_{old} \in V_{old}$;
4. apply Unit Deletion at V_{old} ; and then
5. optionally, apply some combination of Edge Addition and the non-command restructuring operations of Primitive 3.16 on units in V_{new} to achieve some desired structured force.

This is a simple example of a compound restructure, one that involves a combination of the operations we have already defined. In case $|V_{old}| = 1$ and $|V_{new}| > 1$, this is indeed a “split:” the resources in one unit are divided among multiple units. In case $|V_{old}| > 1$ and $|V_{new}| = 1$, this is a “merge:” the resources in multiple units are combined into one. Intermediate cases are also possible, and the force-maker can use this operation to achieve a wide variety of restructures.

⁵¹This is stated without proof, but it is a simple consequence of the fact that the structured forces are finite graphs with decorations hailing from a countable set. As a heuristic proof:

1. repeatedly apply Global Deletion on the source force until $\mathbf{Rel}_0 = \emptyset$;
2. repeatedly apply Local Reassignment on the source force until all units have been assigned the empty configuration;
3. repeatedly apply Vertex Deletion on the source force until $V_0 = \emptyset$, which also ensures $E_0 = \emptyset$;
4. repeatedly apply Vertex Addition until the target force is reached; and
5. repeatedly apply Edge Addition and Global Addition until the target force is reached.

More elaborate procedures are possible, such as the following example: a repopulating operation that, with as much generality as possible, changes the configuration of the force within the confines of a fixed org chart.

3.24 Construction (Repopulating)

Given a structured force $(V_0, E_0, \mathbf{Rel}_0, \text{assign}_0) \in \mathbf{F}_R^*(\mathbf{M}_L^*)$ we define repopulating as follows:

1. apply Unit Addition to include two units, $v_{discard}$ and v_{inject} , equipping $v_{discard}$ with the empty configuration and v_{inject} with some $\biguplus \mathcal{M}_{inject} \in \mathbf{M}_L^*$;
2. apply Reassignment with $V_S = V_0$, $V_T = \{v_{discard}\}$, and $\text{assign}_1(v_{discard}) = \emptyset$;
3. apply Unit Deletion at $v_{discard}$;
4. apply Reassignment with $V_S = \{v_{inject}\}$ and $V_T = V_0$;
5. apply Unit Deletion at v_{inject} ; and
6. optionally, repeatedly apply Localized Reassignment to convert the $\text{assign}_0(v) \uplus \text{pad}_1(v)$ to a target configuration.

The repopulating operation has the potential to perform an exhaustive overhaul of a given org chart's contents. The org chart is augmented with two temporary units, one to discard old resources from the force, the other to inject new ones. Each unit donates resources to be disposed to the discard pile, which is then discarded, and each unit receives resources from the inject pile, which is then discarded, as well. After, the new recruits can be folded into their new home unit's molecule via a round of conversions. This is a powerful operation that allows the force-maker to completely reconfigure her force, and it is a good example of the sort of complex operation that can be built from the elementary operations we have defined.

These two examples illustrate the wide variety of operations that the force-maker can perform on her force. Just as with conversion processes, the graph-theoretic structure of the structured forces provides a framework for thinking about these operations, and the force-maker can use these operations to restructure her force in any way she sees fit. However, the wide array of sequences we can build from these operations suggests that the force-maker might need some way to reason about them—again, this runs parallel with the way she reasons about the conversion of configurations. As such, we will again need a general structure putting the restructuring processes at the fore of the force-maker's mind, one that contains these elementary processes, and those of its ilk, as its atoms—*i.e.*, a categorical structure.

We therefore arrive at the mother of all structured forces.

3.25 Construction (Structured Forces of Restructuring)

We construct the category of structured forces $\mathbf{StructForce}_{\mathbf{F}_R^*(\mathbf{M}_L^*)}$ as follows:

1. the objects are the structured forces $\mathbf{F}_R^*(\mathbf{M}_L^*)$;
2. the morphisms are restructuring operations; these include
 - (a) the primitive operations of Primitives 3.13 to 3.17 and 3.22 and Constructions 3.23 and 3.24; and
 - (b) the impossible restructuring \mathbf{x} , which exists for all pairs of structured forces;
3. the identity morphisms are the do-nothing operations; and
4. composition of morphisms is by concatenation of edits, along with

$$\mathbf{x} \circ \rho = \mathbf{x} = \rho \circ \mathbf{x}$$

for all restructuring operations ρ .

Given two structured forces

$$\mathcal{F}_0 := (V_0, E_0, \mathbf{Rel}_0, \mathbf{assign}_0) \quad \text{and} \quad \mathcal{F}_1 := (V_1, E_1, \mathbf{Rel}_1, \mathbf{assign}_1),$$

the set of all morphisms from \mathcal{F}_0 to \mathcal{F}_1 is denoted $\mathbf{Hom}_{\mathbf{StructForce}_{\mathbf{F}_R^*(\mathbf{M}_L^*)}}(\mathcal{F}_0, \mathcal{F}_1)$. The set of all such morphisms is

$$\mathbf{Hom}_{\mathbf{StructForce}_{\mathbf{F}_R^*(\mathbf{M}_L^*)}}(\mathbf{F}_R^*(\mathbf{M}_L^*)) := \bigcup_{\mathcal{F}_0, \mathcal{F}_1} \mathbf{Hom}_{\mathbf{StructForce}_{\mathbf{F}_R^*(\mathbf{M}_L^*)}}(\mathcal{F}_0, \mathcal{F}_1),$$

where the union is taken over $\mathbf{F}_R^*(\mathbf{M}_L^*) \times \mathbf{F}_R^*(\mathbf{M}_L^*)$.

The category of structured forces is a rich structure that allows the force-maker to reason about the restructuring of her force, be it in terms of new recruits, new equipment, new training, or new organization. The objects of the category are the structured forces, and the morphisms are the restructuring operations that the force-maker can perform. As with the category of force configurations (Construction 2.13), each pair of structured forces comes equipped with the possibility that they cannot be transformed into one another, and this is represented by the impossible restructuring \mathbf{x} . We include the elementary edits just defined as restructuring possibilities, but we also leave open the possibility that the force-maker knows shortcuts beyond sequences of these elementary edits. This allows us to avoid the hubris that we have captured all possible restructuring operations, and it allows the force-maker to reason about the restructuring of her force in a manner faithful to her own understanding of the problem.

3.5 Lifting the Metric

The structured forces living in $\mathbb{F}_R^*(\mathbb{M}_L^*)$ are comprised of configurations, which suggests that we might transport some of the structure of \mathbb{M}_L^* to $\mathbb{F}_R^*(\mathbb{M}_L^*)$. But since these constructions place the configurations into a structured set of buckets, we have more differences to respect whilst lifting the metric. Two force structures may differ not only in the configurations they contain but also in their organization: the number of units, the length of the chain of command, the density of support relationships, and so on. To fully appreciate the structure of the structured forces, we must lift the metric in a way that respects these differences. Thus, we need to consider both the differences between the configurations involved and the manner in which those configurations are wired.

We will continue our political-economic approach to costs, conceiving of them in terms of restructuring. In a sense, we will have to retrace our steps; after all, it is not obvious that the machine we use to calibrate costs for the conversion of configurations will work for restructuring forces—recall that this was the quantale Ξ , which we discussed at length in Sections 2.3 and 2.5, among other locales in the text. It would appear that the way the force-maker reasons about restructuring forces is different from the way the force-maker reasons about the conversion of configurations: one involves institutional development at a very high level, while the other involves equipping, training, manufacturing, and other nitty-gritty details. As such, we introduce a new quantale, Ψ , to calibrate costs for the restructuring of forces.

3.26 Primitive (Cost Labels for Restructuring Forces)

There is a quantale

$$(\Psi, \leqslant_\Psi, \oplus_\Psi, \mathbb{O}_\Psi, \neg_\Psi)$$

of cost labels for the restructuring of forces. Ψ satisfies the properties defined for Ξ Primitives 2.16, 2.18 and 2.19.

Ψ calibrates the evaluation of restructuring costs—i.e., there exists a functor

$$\text{cost}_\Psi : \text{Hom}_{\text{StructForce}_{\mathbb{F}_R^*(\mathbb{M}_L^*)}}(\mathbb{F}_R^*(\mathbb{M}_L^*)) \longrightarrow \Psi$$

assigning a restructuring cost to each morphism in the category of structured forces.

This mimics the construction of Ξ in Section 2.3, but it is a new quantale, Ψ , specifically designated to handle the restructuring of forces.

Just as Ξ is the language in which the force-maker describes the costs of converting configurations, Ψ is the language in which the force-maker describes the costs of restructuring forces. Of course, it could be that $\Xi = \Psi$, suggesting that the two conversations are held in the same terms—or at least, that the languages can be gathered under some common umbrella and share enough overlap in meaning to be useful. But this seems untenable given the differences in the problems at hand and in the sorts of people who play the roles of force-maker at the respective levels. But since the two conversations, however different their syntax, are held by the same force-maker, there ought to be some sort of bridge between the two quantales. For example, consider the Repopulating operation defined in Construction 3.24, where units are added, reassigned, and trained; the operation includes both conversion and restructuring. The force-maker must be able to reason about the costs of this operation in a unified way, even though it involves both conversion and restructuring. This suggests that we need a way to unify the two quantales, Ξ and Ψ , in a way that respects the differences between the two conversations.

We therefore introduce the notion of a bridge between the two quantales. Informally, a bridge is a pair of functions that allow the force-maker to translate between the two languages. We take her ability to do so as primitive.

3.27 Primitive (Bridges of Perspective)

There exists a pair of maps

$$\Phi^* : \Xi \longrightarrow \Psi \quad \text{and} \quad \Phi : \Psi \longrightarrow \Xi$$

providing a bridge of perspective between the quantales

$$\underbrace{(\Xi, \leq_{\Xi}, \oplus_{\Xi}, \mathbb{0}_{\Xi}, \neg_{\Xi})}_{\text{conversion}} \quad \text{and} \quad \underbrace{(\Psi, \leq_{\Psi}, \oplus_{\Psi}, \mathbb{0}_{\Psi}, \neg_{\Psi})}_{\text{restructuring}}.$$

These satisfy the monotonicity condition

$$\begin{aligned} \xi_1 \leq_{\Xi} \xi_2 &\quad \text{implies} \quad \Phi^*(\xi_1) \leq_{\Psi} \Phi^*(\xi_2), \\ \psi_1 \leq_{\Psi} \psi_2 &\quad \text{implies} \quad \Phi(\psi_1) \leq_{\Xi} \Phi(\psi_2). \end{aligned}$$

The bridge of perspective is a pair of functions that allow the force-maker to translate between the two languages of Ξ and Ψ . Since these functions are monotone, higher costs in one domain are mapped to higher costs in the other, and vice versa. This suggests that the force-maker—or pair of force-maker-underlings, one handling conversion and the other restructuring—share unified goals in the conversation, even though they are speaking in different tongues. This equips the force-maker with a powerful capacity of reasoning.

To get a sense of why this matters, consider a large-scale reoutfitting of a force—say, from one sort of rifle to another. To make matters concrete, consider a set of infantry units $V = \{v_1, \dots, v_n\}$. Each of the units houses k infantry; at the beginning of the operation, we have

$$\text{assign}_0(v_k) = \bigcup_{i=1}^{n_k} \left(\underbrace{\text{SOLDIER}^{i,k} \longrightarrow \text{RIFLE}_0^{i,k}}_{\text{call this graph } M_0^{i,k}} \right),$$

where $\text{SOLDIER}^{i,k}$ is the i th soldier in unit v_k and $\text{RIFLE}_0^{i,k}$ is the rifle he is equipped with at the beginning of the operation. The force-maker wishes to reequip the soldiers with a new rifle, $\text{RIFLE}_1^{i,k}$. We have conversion costs

$$\overbrace{\bigoplus_{i,k \in \underline{n} \times \underline{n}_k} \underbrace{(\text{cost}_\Xi \circ \text{CS}) \left(\text{Hom}_{\mathbf{M}_L^\star} \left(M_0^{i,k}, M_1^{i,k} \right) \right)}_{\text{recall} = d_{\mathbf{M}_L^\star} \left(M_0^{i,k}, M_1^{i,k} \right)}}^{\text{in } \Xi},$$

where CS is the choice schedule that picks out the appropriate conversion process for each soldier. But such operations come laden with organizational costs: delivery, coordination, and so on. These costs are not captured by the conversion costs, and thus require their own reasoning.

Assuming (for now) that costs may be decomposed into distinct parts, the bridge allows us to put the two kinds of reasoning into a single expression; suppose $\psi \in \Psi$ is the (fixed) institutional cost of the reoutfitting operation. Then we might write the total cost in terms of Ψ , as in

$$\text{cost}_\Psi(\text{reoutfit}) = \psi \oplus_\Psi \Phi^* \left(\underbrace{\bigoplus_{i,k \in \underline{n} \times \underline{n}_k} d_{\mathbf{M}_L^\star} \left(M_0^{i,k}, M_1^{i,k} \right)}_{\text{in } \Psi} \right).$$

Or, we might write the total cost in terms of Ξ , as in

$$\text{cost}_\Xi(\text{reoutfit}) = \Phi(\psi) \oplus_\Xi \underbrace{\bigoplus_{i,k \in \underline{n} \times \underline{n}_k} d_{\mathbf{M}_L^\star} \left(M_0^{i,k}, M_1^{i,k} \right)}_{\text{in } \Xi}.$$

The quality of the bridge tells us how much is lost when giving ourselves the freedom to consider the costs in either language; in particular, we need decisions made based on the first expression to be consistent with decisions made based on the second. This points us toward properties of the bridge of perspective.

The monotonicity inherent in goal unification is powerful, but it remains that the bridges might lose considerable information and structure in the course of translation. As such, we introduce a series of properties that the bridges might satisfy, each of which captures a different sort of relationship between the two quantales. These are stated in increasing order of strength, all the way up to full-blown isomorphism between the two quantales.

3.28 Definition (Bridges of Perspective)

Given the two cost calibration devices

$$\underbrace{(\Xi, \leq_{\Xi}, \oplus_{\Xi}, \mathbb{0}_{\Xi}, \neg_{\Xi})}_{\text{conversion}} \quad \text{and} \quad \underbrace{(\Psi, \leq_{\Psi}, \oplus_{\Psi}, \mathbb{0}_{\Psi}, \neg_{\Psi})}_{\text{restructuring}},$$

and bridges of perspective

$$\Phi^* : \Xi \longrightarrow \Psi \quad \text{and} \quad \Phi : \Psi \longrightarrow \Xi,$$

we say $\langle \Phi^*, \Phi \rangle$ satisfies:

1. weak monotonicity *just in case*

$$\xi \leq_{\Xi} \Phi(\psi) \quad \text{implies} \quad \Phi^*(\xi) \leq_{\Psi} \psi;$$

2. weak convertibility *just in case*

$$\xi \leq_{\Xi} \Phi(\psi) \quad \text{if and only if} \quad \Phi^*(\xi) \leq_{\Psi} \psi;$$

3. approximate identities *just in case* it satisfies weak convertibility and

$$(\Phi \circ \Phi^*)(\xi) \leq_{\Xi} \xi \quad \text{and} \quad \psi \leq_{\Psi} (\Phi^* \circ \Phi)(\psi);$$

4. strong convertibility *just in case*

$$\Phi^*(\xi) = \psi \quad \text{if and only if} \quad \xi = \Phi(\psi);$$

5. order embeddingness *just in case* Φ and Φ^* are injections and

$$\xi_1 \leq_{\Xi} \xi_2 \quad \text{if and only if} \quad \Phi^*(\xi_1) \leq_{\Psi} \Phi^*(\xi_2),$$

$$\psi_1 \leq_{\Psi} \psi_2 \quad \text{if and only if} \quad \Phi(\psi_1) \leq_{\Xi} \Phi(\psi_2); \text{ and}$$

6. isomorphism *just in case*

$$\Phi^* \circ \Phi = \text{id}_{\Psi} \quad \text{and} \quad \Phi \circ \Phi^* = \text{id}_{\Xi}.$$

These properties navigate the slack between the two quantales. In so doing, they provide us tinkerers with varieties of ways to think about the relationship between the two languages, and thus between the two levels of transformation. The question, then, is how much slack we ought to maintain.

Had we a complete characterization of the morphisms linking two structured forces, we might be on firm ground to answer this question. For example, we might feel comfortable decomposing a given restructuring into its organizational and conversion components, and then reasoning about the costs of each—this was the tack taken in the reoutfitting example above. But the category of structured forces is a complex structure, and we took pains to retain generality in its definition so that the force-maker could consider substitution and complementarity in her restructuring, not to mention generative effects in the combination of costs. But we do not have a complete characterization of the morphisms in this category, and so we must be cautious in our own reasoning; really, all that we know is that restructuring and conversion are both parts of the process and that it stands to reason that their costs are evaluated in different, but related, vocabularies. As such, we content ourselves with a very blunt tool.

3.29 Primitive (Restructuring Costs)

There is a quantale Θ of restructuring costs, and three cost calibration devices

$$\begin{aligned} \text{cost}_\Xi : \text{Hom}_{\text{StructForce}_{\mathbb{F}_R^*(M_L^*)}}(\mathbb{F}_R^*(M_L^*)) &\longrightarrow \Xi, \\ \text{cost}_\Psi : \text{Hom}_{\text{StructForce}_{\mathbb{F}_R^*(M_L^*)}}(\mathbb{F}_R^*(M_L^*)) &\longrightarrow \Psi, \\ \text{cost}_\Theta : \Xi \times \Psi &\longrightarrow \Theta, \end{aligned}$$

where cost_Ξ assigns a conversion component to each morphism, cost_Ψ assigns a restructuring component to each morphism, and cost_Θ combines the two components into a total cost. We assume cost_Θ is monotone in each argument—i.e.,

$$\begin{aligned} \xi_1 \leq_\Xi \xi_2 &\text{ implies } \text{cost}_\Theta(\xi_1, \psi) \leq_\Theta \text{cost}_\Theta(\xi_2, \psi) \text{ for all } \psi \in \Psi, \\ \psi_1 \leq_\Psi \psi_2 &\text{ implies } \text{cost}_\Theta(\xi, \psi_1) \leq_\Theta \text{cost}_\Theta(\xi, \psi_2) \text{ for all } \xi \in \Xi. \end{aligned}$$

The general approach taken here reflects a balance of pragmatic and principled concerns. We leave aside questions of the relationship between the two components of the cost, and we focus instead on the total cost of the operation measured in some unified currency. Whether we require a particular bridge depends on the structure of the quantale Θ and the nature of the costs it contains. For example, if $\Theta = \Xi \times \Psi$, then we might not need a bridge at all, since the total cost is already a combination of the two components. If, on the other hand, Θ is a more complex structure, then we might need a bridge to reason about the relationship between the two components of the cost. The degree to which we need a bridge, and the properties it must satisfy, are questions that depend on the structure of the costs themselves, and so—in the name of analytic humility—we leave them to the force-maker to decide.

3.6 The First Rationality, Part Deux

By now, Section 3 ought to seem an awful lot like Section 2:

1. We've reasoned our way to an appropriate class of structures: in one case the force configurations, in the other the structured forces.
2. We've considered what it means to change these structures, where in both cases we leverage graph-theoretic intuitions to define a category of operations, a considerable generalization in terms of the available transformation operations.
3. We've introduced cost calibration devices to reason about the costs of these operations independently on the two structures, and we've introduced a unification of these devices to reason about the costs of the operations in a single currency.

We will continue to retrace our steps *en route* to a weak metric structure on $\mathbb{F}_R^*(\mathbb{M}_L^*)$. Since the two problems are near-identical in structure, we will be able to reuse much of the machinery we developed in Sections 2.5 and 3.5. We begin by restricting attention to a particular selection of restructuring processes, mimicking Primitive 2.24 in this larger context.

3.30 Primitive (Choice Schedule for Restructuring)

There is a selection of processes

$$\begin{aligned} \mathbf{CS}_R : \mathbb{F}_R^*(\mathbb{M}_L^*) \times \mathbb{F}_R^*(\mathbb{M}_L^*) &\longrightarrow \text{Hom}(\mathbb{F}_R^*(\mathbb{M}_L^*)) , \\ (F_1, F_2) &\longmapsto \mathbf{CS}_R(F_1, F_2) \in \text{Hom}(F_1, F_2) , \end{aligned}$$

representing the force-maker's choice of a restructuring process between all pairs of structured forces.

Just as with the choice schedule for conversion processes, the choice schedule for restructuring processes picks out the restructuring process that the force-maker chooses to apply between any two structured forces. Once again, the impossible process \mathbf{x} is always an option, so that the force-maker need not believe that all pairs of structured forces can be transformed into one another. Since the set of structured forces is so vast and so rich, it stands to reason that the hom sets linking these objects will cover a wide variety of processes, and indeed a wide variety of classes of processes. As such, we again lean hard on the Axiom of Choice to construct the choice schedule, and we leave open the possibility that the force-maker might have a more structured way of choosing between processes. As ever, the blunt instrument will prove more than sufficient for our purposes, and we expect that future work will refine it.

Recall the rationality postulate we imposed in Section 2.4: the force-maker's choice schedule must be internally consistent in the sense that she takes the combination of costs down multi-step paths into account when making her decisions. Suppose that we had

$$\begin{aligned}\mathbf{CS}_C(\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_2) &= f_{12} \in \text{Hom}_{\mathbf{Force}}(\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_2), \text{ and} \\ \mathbf{CS}_C(\biguplus \mathcal{M}_2, \biguplus \mathcal{M}_3) &= f_{23} \in \text{Hom}_{\mathbf{Force}}(\biguplus \mathcal{M}_2, \biguplus \mathcal{M}_3),\end{aligned}$$

and suppose that we now asked the force-maker how she'd convert a $\biguplus \mathcal{M}_1$ to a $\biguplus \mathcal{M}_3$. The rationality postulate tells us that she should consider the costs of the two-step process $f_{23} \circ f_{12}$, and in particular that that her choice $f_{13} = \mathbf{CS}_C(\biguplus \mathcal{M}_1, \biguplus \mathcal{M}_3)$ must satisfy

$$\text{cost}_\Xi(f_{12}) \oplus_\Xi \text{cost}_\Xi(f_{23}) \geq_\Xi \text{cost}_\Xi(f_{13}).$$

Thus, the rationality postulate imposes an upper bound on how costly a process can be, namely no more costly than the multistep process that achieves the same goal. We called this *Compositional Awareness* to reflect the force-maker's need to be aware of the costs of the processes she chooses, and to be aware of how these costs combine in the course of her decision-making. We now extend this postulate to the restructuring of forces.

3.31 Assumption (Compositional Awareness for Restructuring)

The force-maker's choice schedule for restructuring processes satisfies

$$\overbrace{\bigoplus_{i=1}^n (\text{cost}_\Theta \circ \mathbf{CS}_R)(F_{i-1}, F_i)}^{in \Theta} \geq_\Theta (\text{cost}_\Theta \circ \mathbf{CS}_R)(F_0, F_n)$$

for all structured forces \underline{F} and \overline{F} in $\mathbb{F}_R^\star(\mathbb{M}_L^\star)$ and all chains

$$\underline{F} = F_0 \xrightarrow{\mathbf{CS}_R(F_0, F_1)} F_1 \xrightarrow{\mathbf{CS}_R(F_1, F_2)} \dots \xrightarrow{\mathbf{CS}_R(F_{n-1}, F_n)} F_n = \overline{F}.$$

We call this the Compositional Awareness for Restructuring.

Assumptions 2.25 and 3.31 look identical, and indeed they are. Our behavioral imposition is only that the force-maker knows that change can be concatenated. The combination of costs across paths might be nuanced or straightforward, but the force-maker must be aware of how they combine. This is, to repeat, the first rationality, and it is rather weak: in particular, no local optimization is implied.

Compositional Awareness is a desirable property in the present context. Consider a force-maker considering a restructuring operation, say Napoleon considering the introduction of the corps system. In plain terms, the corps system was little more than the introduction of a new level of organization between the division and the army, but it had profound effects on the organization of the French army. But for the present purposes, the more important point is that the introduction of the corps system was a big job—not the sort of thing that could be done all at once. Instead, it was a process that took place over a number of years, with the introduction of the corps system in the Grande Armée in 1805, and its full implementation in 1807. During that time, a broad swath of restructuring operations took place, including the development of new command structures, adjustments to the logistics and supply chains, training of new officers, and so on. In the context of our theory, the introduction of the corps system was a restructuring operation that took place over a number of years, and it was composed of a number of smaller restructuring operations. Each of these operations is encoded as a morphism linking one intermediate structured force to another, and the total restructuring operation is encoded as a chain of morphisms from the initial structured force to the final one:

$$\begin{array}{ccc}
 \text{Force} & \xrightarrow{\text{Develop Command Structures}} & \text{Force with Command,} \\
 & \xrightarrow{\text{Adjust Logistics}} & \text{Force with Command, Logistics,} \\
 & \xrightarrow{\text{Train Officers}} & \text{Force with Command, Logistics, Officers.}
 \end{array}$$

Each of these operations incurred costs, and the total cost of the restructuring was the sum of the costs of the individual operations. (Recall that \oplus_Θ , the operation combining restructuring costs, need not be the literal sum; linearity might be undesirable for such long-term projects, where accumulated fatigue or path dependencies might make the total cost of a restructuring operation more than the sum of its parts, and likewise for happier tales.) Napoleon did not simply wake up one day and say *le corps, c'est chic*; instead, he reasoned about the restructuring of his army in a way that was consistent with the costs of the individual operations. His mind's eye had a target structured force in view, but he also had to think about the processes that would get him there, not to mention the costs of those processes. Compositional Awareness does not demand that Napoleon chose the cheapest restructuring operation, nor that he chose the restructuring operation that would get him to his target structured force in the fewest steps. Rather, it requires that, among the paths of which he was aware, he chose one with an aggregate cost no higher than any other path to the same target structured force.

We have again introduced Compositional Awareness in the service of generating a metric structure, this time for the category of structured forces rather than of force configurations. We define the candidate metric just as with Construction 2.29.

3.32 Construction (Candidate Weak Distance for Structured Forces)

We define the function

$$d_{\mathbb{F}_R^*(\mathbb{M}_L^*)} : \mathbb{F}_R^*(\mathbb{M}_L^*) \times \mathbb{F}_R^*(\mathbb{M}_L^*) \longrightarrow \Theta, \\ (F_1, F_2) \longmapsto \theta(F_1, F_2) \ominus_\Theta \theta(F_1, F_1),$$

where \ominus_Θ is the subtraction operation in Θ and $\theta := \text{cost}_\Theta \circ \mathbf{CS}_R$.

Once again, the candidate weak distance is the difference in cost between moving to a target structured force and staying put in the initial structured force. Just as with force configurations themselves, structured forces are subject to slingshot effects and the urge to decay thanks to a lack of discipline. In the static case where all maintenance processes are costless, the candidate weak distance is simply the cost of moving to the target structured force; in that case, Compositional Awareness is necessary and sufficient for the candidate weak distance to be a weak metric on $\mathbb{F}_R^*(\mathbb{M}_L^*)$ in terms of the quantale Θ . But in the general case, we again need to strengthen Compositional Awareness to Dynamic Awareness, which requires that the force-maker consider the costs of the maintenance processes that keep her structured force in place.

3.33 Assumption (Dynamic Awareness for Restructuring)

We say a choice schedule \mathbf{CS}_R satisfies the Dynamic Awareness for Restructuring just in case

$$(\theta_{12} \ominus_\Theta \theta_{11}) \oplus_\Theta (\theta_{23} \ominus_\Theta \theta_{22}) \geq_\Theta (\theta_{13} \ominus_\Theta \theta_{11})$$

where $\theta_{ij} = (\text{cost}_\Theta \circ \mathbf{CS}_R)(F_i, F_j)$, and where this obtains for all structured forces F_1, F_2, F_3 in $\mathbb{F}_R^*(\mathbb{M}_L^*)$.

Dynamic Awareness is the extension of Compositional Awareness to the dynamic case, where the force-maker must consider the costs of the maintenance processes that keep her structured force in place. She may not choose maintenance processes more costly than the restructuring processes that would achieve the same goal: for the purposes of this theory, *maintenance is necessarily cheaper than transformation*. This is a strong assumption, but it is still weaker than the assumption that maintenance costs are zero. The chaos inherent in a world where maintenance is more costly than transformation is too great for us to consider here, but future work might consider the implications of such a world.

We now have all the pieces in place to deliver the main result of this section.

3.34 Proposition (Weak Metric on Structured Forces)

If cost_Θ and CS_R satisfy Dynamic Awareness for Restructuring, then the candidate weak distance $d_{\mathbb{F}_R^(\mathbb{M}_L^*)}$ is a weak metric on $\mathbb{F}_R^*(\mathbb{M}_L^*)$ calibrated by Θ .*

Proof. This is likely obvious, but given the centrality of the result, we will provide a proof. Recall the requirements were given in Definition 2.28.

1. *Weak Codomain:* naturally, $d_{\mathbb{F}_R^*(\mathbb{M}_L^*)}$ takes values in Θ , because it is a function with Θ as its codomain.
2. *Identity of Disceribles:* we need to show that, for all structured forces $F \in \mathbb{F}_R^*(\mathbb{M}_L^*)$, we have $d_{\mathbb{F}_R^*(\mathbb{M}_L^*)}(F, F) = \mathbb{0}_\Theta$. The distance is given by

$$d_{\mathbb{F}_R^*(\mathbb{M}_L^*)}(F, F) = \theta \Theta_\Theta \theta,$$

where $\theta = (\text{cost}_\Theta \circ \text{CS}_R)(F, F)$. Recall the definition of Θ_Θ :

$$\theta_1 \Theta \theta_2 := \begin{cases} \mathbb{0}_\Theta & \text{if } \theta_2 \geq_\Theta \theta_1, \\ \theta_2 \multimap \theta_1 & \text{otherwise,} \end{cases}$$

where \multimap_Θ is the hom object in Θ . We don't need to consider it, since \geq_Θ is reflexive; thus, we have $\theta \Theta \theta = \mathbb{0}_\Theta$.

3. *Triangle Inequality:* the inequality takes the form

$$d_{\mathbb{F}_R^*(\mathbb{M}_L^*)}(F_1, F_2) \oplus_\Theta d_{\mathbb{F}_R^*(\mathbb{M}_L^*)}(F_2, F_3) \geq_\Theta d_{\mathbb{F}_R^*(\mathbb{M}_L^*)}(F_1, F_3)$$

for all structured forces F_1, F_2, F_3 in $\mathbb{F}_R^*(\mathbb{M}_L^*)$. This is precisely Dynamic Awareness for Restructuring, and so we are done.

We conclude that the candidate weak distance $d_{\mathbb{F}_R^*(\mathbb{M}_L^*)}$ is a weak metric on $\mathbb{F}_R^*(\mathbb{M}_L^*)$ calibrated by Θ . ■

The proof is straightforward, suggesting that the hard work lay in the construction of our theory rather than in the verification of its properties. The result is a weak metric on the category of structured forces, calibrated by the quantale Θ of restructuring costs. The force-maker's reasoning is essentially spatial in character: she can think of the status quo as a point in a space of structured forces. All other possibilities laid out around her, some closer and others farther away. The distances are defined by the costs of the changes she would need to make to get from one structured force to another, the intermediate steps presenting themselves as paths. For all its ugliness, it is, without doubt, a beautiful conceptual world she inhabits.

The final reprise from Section 2 will be to topologize the set of structured forces that reflects the weak metric structure we have just defined. We will follow the same path as before, defining the open sets of the topology to be the sets of structured forces within a certain distance of a given structured force.

3.35 Construction (Topology on Structured Forces)

We define the topology $\mathcal{T}_{\mathbb{F}_R^*(\mathbb{M}_L^*)}$ on $\mathbb{F}_R^*(\mathbb{M}_L^*)$ to be the topology generated by the basis

$$\mathcal{B}_{\mathbb{F}_R^*(\mathbb{M}_L^*)} = \{B_\Theta(F, \varepsilon) \mid F \in \mathbb{F}_R^*(\mathbb{M}_L^*), \varepsilon \in \Theta\},$$

where $B_\Theta(F, \varepsilon) = \{G \in \mathbb{F}_R^*(\mathbb{M}_L^*) \mid \varepsilon >_\Theta d_{\mathbb{F}_R^*(\mathbb{M}_L^*)}(F, G)\}.$

The topology on structured forces is generated by the open sets $B_\Theta(F, \varepsilon)$, which are the sets of structured forces within a distance ε of a given structured force F . Continuous functions on the space of structured forces are those that preserve these open sets, and the force-maker can reason about the continuity of her operations in terms of the open sets of the topology.

Summary of Section 3. In this section, we have developed a theory of structured forces. The theory ran parallel to the theory of force configurations, but with a focus on the organization of forces rather than on the resources they contain. Where a force configuration is a graph with vertices representing force atoms and edges representing their relationships, a structured force is a graph with vertices representing force units and edges representing their relationships. The sorts of atoms accommodated by the theory are the elements of force, and the sorts of force units are the force configurations. But structured forces bring their own special problems to the table: command relationships necessarily form a hierarchy, but there are myriad non-hierarchical relationships that must be accommodated as well. Moreover, reasoning about change among structured forces requires some sort of conceptual bridge between the two languages of conversion and restructuring. We have proceeded by making the weakest possible impositions about the force-maker’s reasoning, and we have shown that these impositions are sufficient to generate a weak metric structure on the category of structured forces. Owing to the similarity between the two theories, we have been able to reuse much of the machinery we developed in the previous section, and we have been able to deliver the main result with relative ease.⁵²

⁵²Your humble author apologizes in case the second treatment has moved a bit too quickly; the similarity between the two theories is such that it is difficult to avoid redundancy, and so the author has taken the liberty of assuming that the reader is comfortable with the concepts and the machinery developed in the previous section. Such assumptions are fraught with peril. Asses like you, me, and the force-maker—we’re all in this together.

4 Force is a Concept

We have climbed the mountain of the concept of force, and it's come time to descend. We have collected many souvenirs along the way, most notably:

1. We have several important sets, including:
 - (a) the elements of force, L ;
 - (b) the force molecules, \mathbb{M}_L ;
 - (c) the force configurations, \mathbb{M}_L^* ; and
 - (d) the structured forces, $\mathbb{F}_R^*(\mathbb{M}_L^*)$.
2. The last two of these come equipped with useful constructions:
 - (a) We can see them as categories. The category of force configurations takes configurations as objects and conversion processes as morphisms, whereas the category of structured forces takes structured forces as objects and force restructurings as morphisms. The respective hom sets are meant to capture real-world possibilities, and surely both categories include conversion processes and force restructurings—not to mention force configurations and structured forces—about which we humans are not yet aware. As such, each hom set also includes the impossible process \mathbb{X} .
 - (b) We can see them as enriched categories under particular choice schedules with appropriate rationality impositions. To do so, we attach to each morphism a cost evaluated in terms of some quantale, and we require that the cost of a composite morphism be the sum (in that quantale) of the costs of its constituent morphisms. Further, we take a snapshot of the force-maker's means of navigating the category, which we call a choice schedule. We require that the force-maker's choices be rational in the sense that they minimize the cost of the force configuration or structured force they produce given the processes of which they are aware.
 - (c) We can see them as weak metric spaces with distances calibrated by the quantale in question. The objects of the category become the skeleton of a sort of map, and the selected morphisms become edges encoding how far one object is from another.

By page 125, the reader is on firm footing in asking whether these souvenirs are worth the journey. Many such readers are motivated by practical concerns, and so we must ask: what can we do with these souvenirs? And, can we store them in our luggage for the next leg of our journey?

This section is dedicated to showing that the souvenirs we have collected are not just trinkets, nor even one-off tools used to strengthen our understanding of force. They are quite capable of doing applied work, too.

1. First, we will take on a particular applied problem in international relations: force emulation. States often mimic one another's force structures, both in how they outfit and organize their forces. We will show that the spaces laid out here can be used to model this process. This is a valuable contribution, since emulation is both common and understudied in the literature—perhaps because there has not been a good way to model it, or perhaps due to bad luck.
2. Second, we will take a more general approach to the question of the desirability of a given structured force—after all, there are many reasons to think that a given force is good or bad, not just the fact that it is or is not a copy of another force. We will show that the structure of the space of structured forces provides the richness we need to bring classical utility theory to bear on the question of force desirability. This, too, is a valuable contribution, as it links our constructions to the real numbers used by game theorists to model strategic arming. Whereas the first two parts of this section are about the *structure* of force, this part is about the *value* of force, which reinforces its conceptual nature: force plays the same part utility does in the theory of rational choice, both a crystal and a residue.
3. Third, we will demonstrate the usefulness of the general approach by consider another applied problem in international relations: strategic arming. There is a robust literature on the topic, including a large set of formal models that use real numbers to model the value of force. This suggests a means of calibrating doctrine to particularized strategic needs, and we will show that our constructions can be used to model this process. Moreover, we will be able to show what is lost by taking a unidimensional approach, namely that equilibria are mod a force level.

In both cases, we will show that the souvenirs we have collected are not just useful, but essential. What is more, we will find that an important aspect of force is not just the variable that takes on the value of a structured force, but the way force-makers navigate the space of structured forces, compare them, and choose among them. Force is not just a variable, but a variable with a particular structure, and it is precisely because we have done justice to this structure that we can do applied work with the concept of force. Of course, it is hoped future readers are more creative than is your humble author, and that they will find even more uses for the souvenirs we have collected.

4.1 Emulation

It was in August of 1885 that Chilean President Domingo Santa María did bring into his employ a Prussian military educator by the name of Emil Körner. Körner, a decorated veteran of the Franco-Prussian war and a well-respected professor of military history, tactics, and ballistics at the Artillery and Engineering School in Charlottenburg, was to lead the Chilean military into the modern era. “Surely,” one might think, “such reforms must have seemed necessary because the Chilean army seemed too weak and not just because it seemed like the wrong sort of army.” But one would be wrong: the Chilean army had just won a valuable victory against Peru and Bolivia in the War of the Pacific, and was widely regarded as one of the most powerful in Latin America. This was not a matter of modernization so much as prussianization: Körner was to bring the Chilean military into the image of the prestigious and recently-successful Prussian army.⁵³ Through a series of reforms, Körner remade the Chilean army in the image of the Prussian.

Körner wasn’t the first choice for the job: that honor went to another Prussian officer, Jakob Meckel (Nunn, 1970). Meckel, however, was already on loan, having just been retained by the Japanese government in a similar capacity; his recommendation for the position came from no less an authority than the Chief of the German General Staff, Helmuth von Moltke (Dogauchi, 2008). Having observed the titanic success of the Prussian army in the Franco-Prussian war, the Japanese strategic establishment—previously under the tutelage of the now-disgraced French—sought to change their guiding star. Though less ingratiated with the Japanese than Körner would be with the Chileans, Meckel nevertheless made a lasting impression on the Japanese military through a series of reforms that brought the Japanese army more in line with the Prussian model.

The Chilean and Japanese cases are not unique, nor is Prussia the only model for emulation, nor is there anything special about the end of the 19th century, nor are military forces the only things that states emulate. It happens to be the case that South American states were particularly likely to emulate western-European models in the 19th century, but this is not a general rule. Resende-Santos (2007, p. 3) records twelve emulation projects by South American states from 1870–1930, five emulating Germany and four emulating France.⁵⁴

⁵³Details throughout this section are drawn from a few useful sources, including Frederick M. Nunn’s *Yesterday’s Soldiers* (1983), Bernd Martin’s *Japan and Germany in the Modern World* (1995), William F. Sater and Holger H. Herwig’s *The Grand Illusion* (1999), Emily O. Goldman and Leslie C. Eliason’s *The Diffusion of Military Technology and Ideas* (2003), and João Resende-Santos’s *Neorealism, States, and the Modern Mass Army* (2007).

⁵⁴The others all emulate Chile, who by then was already emulating Germany. Japan emulated France from 1866–1878 and Germany from 1878–onward, which poses interesting problems

Körner and Meckel's reforms covered a wide range of topics, some of which comport well with our understanding of force, others suggesting promise for future work. Regardless, they suggest that the emulation problem is a rich one, and that it is worth our time to model it.

1. *The org chart:* The Prussian army defeated the French, in large part, because of its ability to deliver force to the battlefield more quickly. Mobility lay at the core of the Prussian army's success, and this was in large part due to its organizational structure. Meckel formally introduced the divisional and regimental level of organization to the Japanese army, which had previously been organized only at lower levels despite influences from the French. Improved mobility is credited as an important factor in the Japanese victory in the First Sino-Japanese war (Dogauchi, 2008), where the Japanese army was able to move more quickly than the Chinese.⁵⁵ Less successful were the changes to org charts in the Chilean army, where Körner originally placed himself (as Inspector General) at the top of a centralized org chart with direct reporting from a broad, high-level staff. This invited political pushback from the civilian government, and eventually the compromise was a set of units determined by the army but understaffed due to insufficient support (Nunn, 1970). Sater and Herwig (1999) argue that these changes were too superficial in nature and that the attendant lack of change in the more subtle support structures of the Chilean army led to diminished effectiveness. But, such changes were part of second- and third-wave reforms (Resende-Santos, 2007, p. 147), which also included the introduction of the divisional and corps-based structure to the Chilean army.

It should be noted that, for both the Japanese and Chilean cases, the changes to the force structure were implemented reasonably quickly: it was among Meckel's first changes the Japanese army and Körner's first changes in the second-wave of reforms after the Chilean Civil War. This suggests that org charts are indeed a good place to start when asking a force-maker what "change" looks like. Happily, our construction is well-suited to accommodate these decisions (and occasional non-decisions) about

because France started emulating Germany in 1870. *Mon dieu! Dios mío! Mein Gott!* 神様!

⁵⁵The Chinese army had itself undergone extensive renovation, but this renovation was not emulation-based. The so-called *Self-Strengthening Movement* was a series of reforms that sought to modernize the Chinese army without reference to any particular model. Included in the movement was an attempt to send Chinese cadets to the United States Military Academy, but no formal changes were made to the Chinese army's org chart, nor the way it drilled (Fung, 1996). Resende-Santos (2007) would call this a case of *innovation*, not emulation.

force structure, not to mention deeper aspects of organization. Suddenly our sandbox feels more like a box full of fertile soil.

2. *Equipment.* Prior to Hölder’s arrival in Chile, the Chilean army’s small arms were largely French (with Belgian and American influences) and its artillery largely German; after his arrival, the Chilean army was outfitted with German small arms and artillery exclusively, with purchasing and testing also among the highest early priorities (Resende-Santos, 2007, p. 134). This, too, suggests a move in the direction of the Prussian model. Meckel’s reforms suggest that reoutfitting is a matter of considerable detail: Japanese uniforms were recut to match Prussian styles rather than the prevailing French ones (Martin, 1995, p. 40).

Our construction is well-suited to accommodate these decisions, too: we have extreme flexibility in the choice of force molecules, and we can easily model the costs of retooling via appropriate force conversions or restructurings. The costs of such processes are not trivial, so the flexibility of the quantales Ξ , Ψ , and Θ are most useful here.

3. *The officer corps and general staff:* Both Körner and Meckel expanded the officer corps and general staff of their respective armies, and both introduced a system of military education that was more in line with the Prussian model. This is a more subtle change than the previous two, but it is no less important. We can proceed quite naively regarding the size of the officer corps, since we merely need to load up larger configurations into the respective headquarters units. The general staff takes its character from a mixture both of the resources required and the place occupied in both the command partial order and the non-command hypergraphs; while it is more nuanced than the size of the officer corps, it seems quite tractable nonetheless. Training is more difficult to model, though non-command relationships in the hypergraph might be a good place to start, along with “catalyst” units that point to more desirable configurations.
4. *Conscription:* Now here emerge more serious difficulties, as we have not thought very hard about the origin of the materials that make up the force molecules. Körner and Meckel both introduced mandatory conscription to their respective armies, and this is a change that is not easily modeled in our current framework. Now, it could well be that any one of the cost calibration quantales includes some dimension—or dimensions—for various aspects of conscription, and that we could model the costs of conscription in terms of these dimensions. Or, the molecules might differentiate a soldier who’s been conscripted from one who hasn’t, keeping them the same in all other respects. These provide partial coverage of the problem, but not complete solutions; we’d at least have to think through a state’s labor market.

5. *Tactics*: Meckel brought with him an abiding love of infantry-based tactics, again in line with the Prussian model. This is a change that is not easily modeled in our current framework, either: we have studied the structure of force, but not the way it is used. Here it is hard even to salvage partial coverage, so instead we must admit that our construction is incomplete. But let it be suggested here, if the reader will indulge, that it would be most difficult indeed to talk very deeply about the way force is used without knowing just what it is that's being used in the first place. Naturally, it is not hard to imagine applications involving constructions from a given force representing deployments and tactics thereof.
6. *Other practices, like promotion and retirement*: When a state hires a foreign officer to reform its military, it is not just hiring a foreign officer. It is hiring a foreign officer who has been successful in their own military, and who has a particular set of practices that they bring with them. These foreign officers often come with their own staffs from their home countries, and these staffs often bring with them their own practices. The changes these foreign officers impose often refer more to internally-focused, non-force-bearing practices that nevertheless remain altogether indispensable to the character of the force. These are changes that are not easily modeled in our current framework, either.

Suffice it to say that the changes were deep and wide-ranging, intentional but subject to the vagaries of politics and the limits of the possible, oft-understandable in the language we've developed but imperfectly so. The emulation problem is a rich one, and it is worth our time to model it. What is more, we will find that an important aspect of force is not just the variable that takes on the value of a structured force, but the way force-makers navigate the space of structured forces, compare them, and choose among them. The force space is not just a variable, but a variable with a particular structure, and it is precisely because we have done justice to this structure that we can do applied work with the concept.

Given the incomplete coverage of our construction, our modeling goals are relatively modest: the goal is more to show that the problem can be appreciated in the language of force than to get every detail right. Indeed, the most immediate goal is to show that a first-cut encoding of the emulation problem is well-posed in the language of force—*i.e.*, it can be written down in a way that makes sense and that can be solved. Only then can we begin to think about how to refine our construction to better model the emulation problem, though the thoughts below are more applicable than they are applied. However, it is hoped that the reader will find the thoughts below to be a useful starting point for thinking about the emulation problem, and similar applied problems, in the language of force.

The emulation problem takes a rather simple form:

4.1 Construction (Emulation Problem)

An emulation problem *takes the form*

$$\min_{\mathcal{F} \in \mathcal{D}_S(\beta)} d_\Theta(\mathcal{F}, \mathcal{T}),$$

where

1. $d_\Theta : \mathbb{F}_R^*(\mathbb{M}_L^*) \times \mathbb{F}_R^*(\mathbb{M}_L^*) \rightarrow \Theta$ is the emulator's weak metric on structured forces calibrated by Θ ;
2. $\mathcal{S} \in \mathbb{F}_R^*(\mathbb{M}_L^*)$ is a structured force representing the emulator's status quo;
3. $\mathcal{T} \in \mathbb{F}_R^*(\mathbb{M}_L^*)$ is a structured force representing the target force;
4. $\beta \in \Theta \setminus ([0_\Theta] \cup \{\infty\})$ is the budget for the emulation process; and
5. the set

$$\mathcal{D}_S(\beta) = \{\mathcal{F} \in \mathbb{F}_R^*(\mathbb{M}_L^*) \mid \beta \geq_\Theta d_\Theta(\mathcal{F}, \mathcal{S})\} \subseteq \mathbb{F}_R^*(\mathbb{M}_L^*)$$

is the domain of the emulation problem.

The emulation problem is quite simple: minimize the distance between the *status quo* force and some target force, subject to a budget constraint based on that *status quo* and an acceptable amount of change.

The most immediate question is whether the emulation problem is well-posed—*i.e.*, does it always possess at least one solution? This is a common question in optimization theory, and social scientists usually write down problems where simple versions of the Weierstrass extreme value theorem apply. The Weierstrass theorem(s) are a set of results that guarantee the existence of a solution to an optimization problem under certain conditions. Ok's (2007, p. 67) “Baby” Weierstrass Theorem reads: for any $a, b \in \mathbb{R}$ such that $a \leq b$ and any continuous function $f : [a, b] \rightarrow \mathbb{R}$, there exists some $x \in [a, b]$ such that $f(x) \leq f(y)$ for all $y \in [a, b]$. His full-blown Weierstrass Theorem (p. 225) reads: if X is a compact metric space and $\varphi : X \rightarrow \mathbb{R}$ is continuous, then there exists some $x \in X$ such that $\varphi(x) = \inf_{y \in X} \varphi(y)$. Thus, Ok's most general version of the Weierstrass theorem requires that the objective function be real-valued and the domain be a compact metric space. We have neither of these at our disposal. Aliprantis and Border (2006, Corollary 2.35, p. 40) help by removing the metric requirement, but they continue to work in the real-valued case. Our objective function takes values in the quantale Θ , and it seems unlikely that emulators will have appreciably-simpler Θ s than other force-makers. We therefore must look elsewhere for help.

We first want to show that the function

$$d_\Theta : \mathbb{F}_R^*(M_L^*) \times \mathbb{F}_R^*(M_L^*) \rightarrow \Theta$$

is continuous. This raises a natural question: what is continuity in abstract spaces like these? After all, it seems unlikely that we can appeal to you-never-lift-your-pencil continuity, since we are working in highly-abstract spaces. However, this is precisely what topology is for: it is the study of continuity in highly-abstract spaces. The most general definition of continuity is in terms of open sets, and the most general definition of an open set is in terms of a topology. A function $f : X \rightarrow Y$ between topological spaces is continuous if the preimage of every open set in Y is an open set in X . Now, we've worked very hard to identify a suitable topology for our X —in this case, $\mathbb{F}_R^*(M_L^*)$ —but we've yet to provide a topology for our Y —in this case, Θ . This means that we don't yet know what tests our function must pass to be continuous.

We have a few options here.

1. At one end of the spectrum, we could equip Θ with the trivial topology, where the only open sets are \emptyset and Θ . The preimage of any open set in Θ is either \emptyset or $\mathbb{F}_R^*(M_L^*)$, both of which are open in $\mathbb{F}_R^*(M_L^*)$. Thus, *any* function from $\mathbb{F}_R^*(M_L^*)$ to Θ is continuous, including d_Θ . This is akin to curing a headache by amputation.
2. At the other end of the spectrum, we could equip Θ with the discrete topology, where every subset of Θ is open. Continuity therefore involves checking an inordinate number of preimages for subsets of Θ about which we know very little (by design). This is akin to curing a headache by working through a 12-step program.
3. We instead focus on intermediate possibilities that leverage what we know about Θ while still providing a reasonable test for continuity. We know, for example, that Θ comes equipped with a preorder \geq_Θ , so we could equip Θ with the order topology, generated by the open intervals

$$\begin{aligned} \text{Up}^{\text{Order}}(\theta) &:= \{\theta' \in \Theta \mid \theta' >_\Theta \theta\} \quad \text{and} \\ \text{Lo}^{\text{Order}}(\theta) &:= \{\theta' \in \Theta \mid \theta' <_\Theta \theta\}, \end{aligned}$$

which winds up including all the “invervals” and “rays” of Θ , where the scare quotes reflect the fact that these are not intervals or rays in the usual sense but rather are an abstractified set of projected rectangles and half-planes. The lack of completeness on \geq_Θ makes this an unwieldy choice, but it is a choice that is at least somewhat reasonable—akin to curing a headache by taking flu medicine.

4. We could try something more tailored, like a topology defined only by the half-planes $\text{Lo}^{\text{Order}}(\theta)$. This is related, but not identical to, a well-known topology: the *dual Alexandrov topology*.⁵⁶ The dual Alexandrov topology is defined by the open sets

$$\text{Lo}^{\text{Alex}}(\theta) := \{\theta' \in \Theta \mid \theta' \leq_{\Theta} \theta\},$$

which are the half-planes of Θ . However, the inclusivity of the Alexandrov topology makes it a poor choice for our purposes, as it is not clear that the preimage of an open set in the Alexandrov topology is open in $\mathbb{F}_R^*(\mathbb{M}_L^*)$. This is akin to curing a headache by taking a nap. But by defining the open sets as

$$[\theta_\Theta, \theta) = \{\theta' \in \Theta \mid \theta' <_{\Theta} \theta\},$$

we can define a topology that is both more natural for our purposes and more likely to be useful.⁵⁷ This will serve as our choice for the topology on Θ , curing a headache by taking an appropriate dose of aspirin. We call it the *strict-bound topology*.

This list is surely not exhaustive, but it is enough to get a sense about things. No one decision is perfect, but the strict-bound topology is a reasonable choice that is likely to be useful. Other applied problems will likely require other topologies.

Beyond a topology for Θ , we also introduce an assumption ensuring that continuity is a meaningful property with respect to our intuitions. We have not assumed much about how many elements are in Θ , and it could well be that it is prosaic in the extreme—say, $\Theta = \{\text{FALSE}, \text{TRUE}\}$, where this represents impossibility and possibility, respectively. While continuity is a meaningful property in this case, it is not particularly interesting. To that end, we introduce a density assumption that ensures that if one element is strictly above another, then there is some element that is strictly above the lower element but strictly below the higher element. This property is often called *order density*, and it is a reasonable assumption to make here. We will go one step further to assume that the “gap” can be filled, or partially filled, by two copies of a gap-filling element.

⁵⁶Here “dual” refers to the fact that the open sets are defined by the lower sets of the preorder, rather than the upper sets. In the standard Alexandrov topology, the open sets are defined by the upper sets of the preorder. In our context, nothing is lost by taking the dual.

⁵⁷The only difference here lies in the use of $<_{\Theta}$ rather than \leq_{Θ} . This precludes the usual problems one might encounter by using \leq on \mathbb{R} , but also situations where the lack of antisymmetry in Θ introduces isomorphism classes of bounds. Given the similarity, it seems like we should call our topology something related to Alexandrov, but this would insult the latter, which includes the important feature that arbitrary intersections remain open. No such properties emerge here, meaning the strict-bound topology lacks the property that makes Alexandrov Alexandrov.

We formalize the idea as follows.

4.2 Assumption (Double Separation Property)

For the purposes of Section 4.1, we assume

$$\theta_1 >_\Theta \theta_2 \implies (\theta_1 \geq_\Theta \theta_3 \oplus_\Theta \theta_2 \oplus_\Theta \theta_3 \text{ for some } \theta_3 >_\Theta 0_\Theta).$$

This is a property that is often assumed in the study of order topologies, and it is a reasonable assumption to make here. In essence, it asserts that each hom element $\theta_2 \rightsquigarrow \theta_1$ is dense in Θ , and it requires this *only* of the hom elements. This assumption is slightly stronger than is absolutely necessary for the current purposes, but it is a reasonable assumption to make, and it makes the proof of the claim below particularly straightforward.

We now turn to the question of continuity. To repeat, continuity is important because it ensures that the function d_Θ is well-behaved with respect to the topology on Θ . It means that small changes in the forces between which we measure distance result in small changes in the distance between them: if your hands are one foot apart and then you move them each by a small amount, they will still be close to one foot apart. Here we mean that if two forces can be converted into one another with a small amount of restructuring, then this will remain the case even if each of them is restructured by a small amount first.

It turns out that our care in choosing topologies and issuing assumptions has paid off, as we can show that the function d_Θ is continuous rather painlessly.

4.3 Lemma (Continuity of the Weak Metric)

Under the double separation property, the function

$$d_\Theta : \mathbb{F}_R^*(M_L^*) \times \mathbb{F}_R^*(M_L^*) \longrightarrow \Theta, \\ (\mathcal{F}, \mathcal{G}) \longmapsto (\text{cost}_\Theta \circ \mathbf{CS}_R)(\mathcal{F}, \mathcal{G})$$

is continuous, where the domain $\mathbb{F}_R^(M_L^*) \times \mathbb{F}_R^*(M_L^*)$ has the product topology obtained from the weak metric topology on $\mathbb{F}_R^*(M_L^*)$ representing restructuring costs, and the codomain Θ has the strict-bound topology.* *[Proof.]*

This generalizes the well-known fact that traditional—*i.e.*, real-valued—metrics are continuous on their domain. This suggests that this well-understood continuity does not rely on all of the properties of traditional metrics: not real-valuedness, nor symmetry, nor the identity of indiscernibles matters for continuity. On the other hand, the triangle inequality plays an important role in the proof, as it does in the proof of continuity for traditional metrics; this is the essence of distance, and it provides us with the basic property of continuity that we need. This result represents one passed test for traditional approaches to optimization.

We now turn our attention to the compactness the set

$$\mathcal{D}_S(\beta) = \{\mathcal{F} \in \mathbb{F}_R^*(\mathbb{M}_L^*) \mid \beta \geq_{\Theta} d_{\Theta}(\mathcal{F}, S)\} \subseteq \mathbb{F}_R^*(\mathbb{M}_L^*),$$

where S is the emulator's status quo force. Compactness is a topological property that is a generalization of the notion of a set being closed and bounded in \mathbb{R}^n . The set $[0, 1] \subseteq \mathbb{R}$ is compact, for example, as is the closed unit ball in \mathbb{R}^n . $[0, \infty)$ is not compact, as it is closed but not bounded, and $(0, 1)$ is not compact, as it is bounded but not closed. In general topological spaces, one defines compactness by way of open covers and finite subcovers, as in the following definition.

4.4 Definition (Compact Set)

Let (X, \mathcal{T}) be a topological space. We say a set $K \subseteq X$ is compact just in case every open cover of K has a finite subcover—i.e., if there exists

$$\{O_{\alpha}\}_{\alpha \in A} \subseteq \mathcal{T} \text{ such that } K \subseteq \bigcup_{\alpha \in A} O_{\alpha},$$

then there exists some $n \in \mathbb{N}$ and some function $\alpha : \underline{n} \rightarrow A$ such that

$$K \subseteq \bigcup_{i=1}^n O_{\alpha(i)}.$$

Essentially, compactness is a property that ensures that a set is not too big to study things like continuity and convergence; this is quite important in the optimization setting laid out in Construction 4.1. It generalizes our ideas of closedness, boundedness, and—more fundamentally—finiteness. We would very much like to show that $\mathcal{D}_S(\beta)$ is compact, as it would provide intuitions from a host of optimization problems that are not immediately available in our setting. However, compactness can be difficult to show, and a host of related properties are often required to prove it.

We therefore count our blessings when we find that the set $\mathcal{D}_S(\beta)$ is compact.

4.5 Lemma (Emulation Problems Have Compact Domain)

Under the topologies from Lemma 4.3, the set

$$\mathcal{D}_S(\beta) = \{\mathcal{F} \in \mathbb{F}_R^*(\mathbb{M}_L^*) \mid \beta \geq_{\Theta} d_{\Theta}(\mathcal{F}, S)\} \subseteq \mathbb{F}_R^*(\mathbb{M}_L^*)$$

is compact for any $S \in \mathbb{F}_R^*(\mathbb{M}_L^*)$ and $\beta \in \Theta$.

[[Proof](#).]

Thus, we have the tools at our disposal required to apply Weierstrassian thinking to the emulation problem.

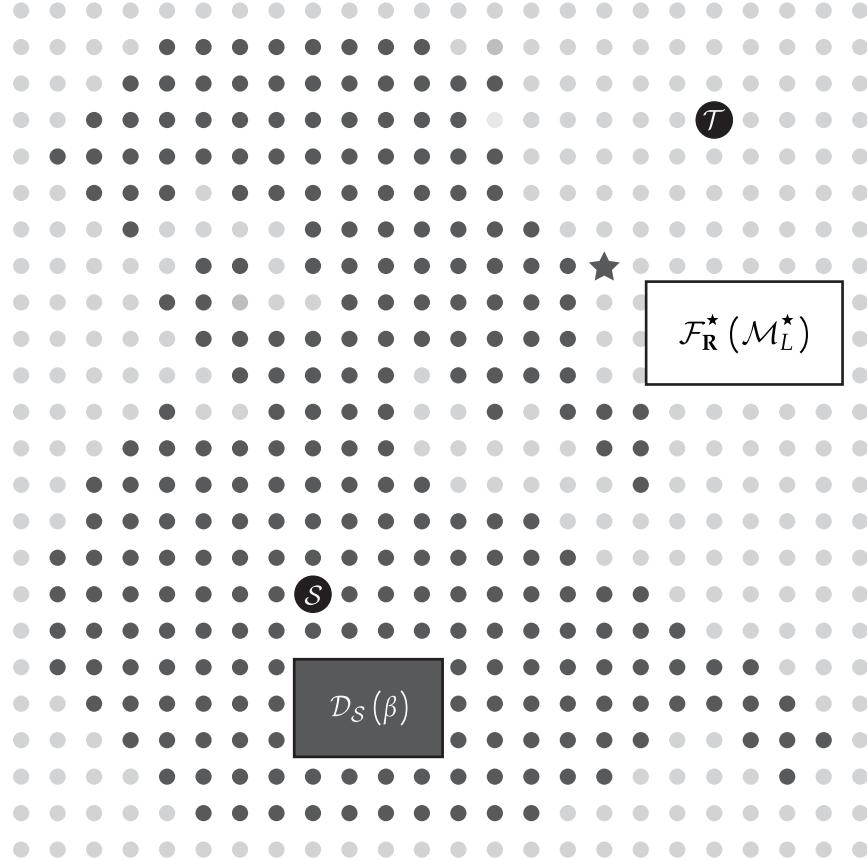


Figure 12: Solving the emulation problem.

We now state the main result about the emulation problem.

4.6 Proposition (Emulation Problems Have Solutions)

Under the topologies from Lemma 4.3, the emulation problem has a solution. [Proof.]

Figure 12 depicts the solution to the emulation problem. The choice set $\mathcal{D}_S(\beta)$ is a compact subset that includes the source force \mathcal{S} , and the target force \mathcal{T} is some point in the space of structured forces. The solution, marked with a star, is the point in $\mathcal{D}_S(\beta)$ that is closest to \mathcal{T} in the weak metric. Remarkably, this process works no matter where our source force lives in the vast sea that is $\mathcal{F}_R^*(\mathcal{M}_L^*)$, no matter where the target force lives, and no matter the size of the budget, which merely defines the dark gray region.⁵⁸

⁵⁸Note also that neighborhoods need not look like neighborhoods in the usual sense.

Based on what we learned *en route* to the proof of Proposition 4.6, we can make a few more observations to enrich our understanding of the emulation problem. We begin by observing that Θ , being merely a preorder, can make it difficult to understand the senses in which a given budget is “large” or “small.” The problem demands that the amount spent making changes to the force be less than some budget $\beta \in \Theta$, but it is not clear how much information is involved in this statement. We can envision something like a “multidimensional” budget

$$B = \{\beta_1, \dots, \beta_k\} \subseteq \Theta,$$

where each β_i represents a different aspect of the budget.⁵⁹ These aspects might include purely-financial considerations, but they might also include considerations of political feasibility, social acceptability, or military necessity. In case $B = \{\beta_1, \beta_2\} \subseteq \Theta$, four cases emerge:

1. $\beta_1 \lessdot_\Theta \beta_2$, so that β_2 is the only constraint on the emulation problem;
2. $\beta_2 \lessdot_\Theta \beta_1$, so that it is β_1 that is the sole constraint;
3. $\beta_1 \doteq \beta_2$, so that both represent the same constraint; or
4. β_1 and β_2 are incomparable, so that both are distinct constraints.

Naturally, as the number of dimensions in the budget increases, the number of cases increases. The common structure of Θ ensures that the complications all take similar forms, so that we study finite “strands” within the vast “web” of Θ .

Happily, this enrichment creates no new difficulties.

4.7 Corollary (Multidimensional Emulation Problems Have Solutions)

Let $B = \{\beta_1, \dots, \beta_k\} \subseteq \Theta$ be a multidimensional budget, and define the domain

$$\mathcal{D}_S(B) = \{\mathcal{F} \in \mathbb{F}_R^*(\mathbb{M}_L^*) \mid \beta_i \geq_\Theta d_\Theta(\mathcal{F}, \mathcal{S}) \text{ for all } i \in \underline{k}\} \subseteq \mathbb{F}_R^*(\mathbb{M}_L^*).$$

The problem of minimizing $d_\Theta(\mathcal{F}, \mathcal{T})$ subject to $\mathcal{F} \in \mathcal{D}_S(B)$ has a solution.

The key here is that $|B| = k$ is finite: each aspect of the budget induces a compact set in Θ , and their union is compact there, too. We will stick with multidimensional budgets throughout the rest of the section.

⁵⁹The scare quotes on “multidimensional” are meant to indicate that we are not talking about a vector space, but rather a set of elements from a preorder. It’s not as if the elements of B can be added or subtracted, nor that they are in any sense orthogonal. In a vector space, budgets might look like orthogonal vectors—*i.e.*,

$$(\beta_1, 0_\Theta, 0_\Theta, \dots, 0_\Theta), \quad (0_\Theta, \beta_2, 0_\Theta, \dots, 0_\Theta), \quad \dots, \quad (0_\Theta, 0_\Theta, 0_\Theta, \dots, \beta_k),$$

or if you prefer, a diagonal matrix $\text{diag}(\beta_1, \dots, \beta_k)$. But we need no such structure here.

Naturally, the budget is not the only thing that might be multidimensional, nor is it the most interesting thing. For example, it could well be that the emulator wishes to emulate multiple target forces; it was common, for example, for 19th-century states to emulate German armies and British navies; nowadays, one might emulate the United States in some respects and China in others. This is a more complicated problem, and since we know less about the structure of $\mathbb{F}_R^*(\mathbb{M}_L^*)$ than we do about Θ , we must proceed delicately. Of course, in choosing a delicate course of action, we learn quite a bit about the problem itself, which at least ensures that the rumination is of substantive value.

Multitargeted Emulation via \oplus_Θ . Suppose we started from some finite set of emulation targets $\mathcal{T}_1, \dots, \mathcal{T}_m \in \mathbb{F}_R^*(\mathbb{M}_L^*)$, all of which we wished to emulate. One way to handle this is to consider the problem

$$\min_{\mathcal{F} \in \mathcal{D}_S(B)} \bigoplus_{i=1}^m d_\Theta(\mathcal{F}, \mathcal{T}_i).$$

This is a straightforward generalization of the single-target case, where the idea is to penalize the emulator for being far from any of the targets. This is a good place to start, as it seems to introduce minimal modifications to the single-target case. But, this approach has a downside: under the topologies we have been using, it works only if Θ has more structure than we've been assuming:

4.8 Lemma (Continuity of \oplus_Θ)

The function $\oplus_\Theta : \Theta \times \Theta \rightarrow \Theta$ is continuous in the strict-bound topology if, and only if, for any pair of labels $\theta_1, \theta_2 \in \Theta$, there exists some $\theta_1 \leftharpoonup \theta_2 \in \Theta$ such that

$$\theta_3 \oplus_\Theta \theta_2 \leq_\Theta \theta_1 \quad \text{if and only if} \quad \theta_3 \leq_\Theta \theta_1 \leftharpoonup \theta_2,$$

called the strict left adjoint of \oplus_Θ .

[*Proof.*]

We've taken great pains to impose as little structure as possible on Θ , and thus prefer not to introduce new structure if we can avoid it. In some circumstances, the juice is worth the squeeze, such as when a problem naturally allows for Euclidean-like operations on Θ . But in general, it is better to proceed without requiring new machinery. What's more, it is not obvious that this is how the emulator actually reasons: it seems more likely that the emulator would consider the targets separately and then make a decision based on all of them, instead of considering the sum of the costs it would take to get to each target. Consequently, we will leave this option aside and consider more general approaches.

Multitargeted Idealization. Now suppose you were tasked with emulating the Prussian army and the British navy. Would your reasoning look like the reasoning inherent in the approach just discussed? It seems unlikely: when evaluating a feasible option, would you say to yourself “and I could get this close to the Prussian army and this close to the British navy?” Probably not. Instead, you might use the Prussian army and British navy as inputs to an idealization process, where you envision a structured force similar to the Prussian force in army respects and similar to the British force in navy respects. In other words, you would turn two (real) targets into one (ideal) target.

One could have more than two targets, which need not be Prussian or British. The same goes for the reality of the targets: for example, if you handed Machiavelli a structured force and told him to make changes based on multitargeted emulation, his idealization process would include a structured force that modern in some senses and ancient in others. Indeed, our structure allows the emulator to conjure up any number of ideal targets from very different parts of $\mathbb{F}_R^*(\mathbb{M}_L^*)$, a set that contains every force that ever was or could be. So, we introduce the *idealization map*:

4.9 Primitive (Idealization)

There is a map

$$\begin{aligned} \text{Ideal} : \mathbb{P}(\mathbb{F}_R^*(\mathbb{M}_L^*)) \times \mathbb{F}_R^*(\mathbb{M}_L^*) &\longrightarrow \mathbb{F}_R^*(\mathbb{M}_L^*) , \\ (\{\mathcal{T}_1, \dots, \mathcal{T}_m\}, \mathcal{F}) &\longmapsto \text{Ideal}(\{\mathcal{T}_1, \dots, \mathcal{T}_m\}, \mathcal{F}) \end{aligned}$$

called the idealization map. We refer to $\text{Ideal}(\{\mathcal{T}_1, \dots, \mathcal{T}_m\}, \mathcal{F})$ as the ideal target of \mathcal{F} with respect to $\{\mathcal{T}_1, \dots, \mathcal{T}_m\}$. As a matter of course, we impose $\text{Ideal}(\emptyset, \mathcal{F}) = \mathcal{F}$.

An idealized emulation problem takes the form

$$\min_{\mathcal{F} \in \mathcal{D}_S(B)} d_\Theta(\mathcal{F}, \text{Ideal}(\{\mathcal{T}_1, \dots, \mathcal{T}_m\}, \mathcal{S})) ,$$

where $\{\mathcal{T}_1, \dots, \mathcal{T}_m\} \subseteq \mathbb{F}_R^(\mathbb{M}_L^*)$ is a set of targets.*

The idealization map is a powerful tool, as it allows the emulator to turn any number of targets into a single target, which then is input to the single-target emulation problem as we’ve already discussed.⁶⁰ The idealization process is with respect to the status quo, \mathcal{S} , ensuring conceptual continuity with the single-target case. The nuances introduced are precisely the nuances of the map Ideal , which we now examine in more detail.

⁶⁰Nothing precludes a stubborn emulator from having a constant-valued ideal target, but this then reduces to the single-target case with some exogenously-determined target.

Since d_Θ is continuous in its two arguments, our main goal is to show that Ideal is continuous in its two arguments—this would mean that d_Θ is continuous in the idealized emulation problem. This requires setting a topology for the domain of Ideal , which is $\mathbb{P}(\mathbb{F}_R^*(\mathbb{M}_L^*)) \times \mathbb{F}_R^*(\mathbb{M}_L^*)$, where $\mathbb{P}(\mathbb{F}_R^*(\mathbb{M}_L^*))$ is the power set of $\mathbb{F}_R^*(\mathbb{M}_L^*)$. We will use the strict-bound topology on $\mathbb{P}(\mathbb{F}_R^*(\mathbb{M}_L^*))$, where the order comes from supersethood, \supseteq . In other words, for any two subsets $A, B \subseteq \mathbb{F}_R^*(\mathbb{M}_L^*)$, we have $A \supseteq_B B$ if $A \supseteq B$. Substantively, this means that if two target sets $T = \{\mathcal{T}_1, \dots, \mathcal{T}_{m_1}\}$ and $T' = \{\mathcal{T}'_1, \dots, \mathcal{T}'_{m_2}\}$ are similar to one another by way of one being a superset of the other—*i.e.*, if we add a few more targets to T to get T' —then the ideal targets,

$$\text{Ideal}(T, \mathcal{F}) \quad \text{and} \quad \text{Ideal}(T', \mathcal{F}),$$

should be easy to convert from one to the other.

Hierarchical Targeting. Third, we consider the case where the emulator has a set of targets that are related to one another in a hierarchical fashion. Rather than amalgamating the targets into a single ideal target, the emulator might instead use the hierarchy to guide the emulation process. For example, given a finite set of targets $\mathcal{T}_1, \dots, \mathcal{T}_m \in \mathbb{F}_R^*(\mathbb{M}_L^*)$, the emulator might go:

$$\begin{aligned} \mathcal{F}_1 &:= \underset{\mathcal{F} \in \mathcal{D}_{\mathcal{S}}(B_1)}{\operatorname{argmin}} d_\Theta(\mathcal{F}, \mathcal{T}_1), \\ \mathcal{F}_2 &:= \underset{\mathcal{F} \in \mathcal{D}_{\mathcal{F}_1}(B_2)}{\operatorname{argmin}} d_\Theta(\mathcal{F}, \mathcal{T}_2), \\ &\vdots \\ \mathcal{F}_m &:= \underset{\mathcal{F} \in \mathcal{D}_{\mathcal{F}_{m-1}}(B_m)}{\operatorname{argmin}} d_\Theta(\mathcal{F}, \mathcal{T}_m), \end{aligned}$$

where $B_1, \dots, B_m \subseteq \Theta$ are the budgets for each step.⁶¹ At each step, the emulator minimizes the distance to the target force, subject to the constraint that the restructuring costs are less than the budget for that step. The next step begins with the force from the previous step, rather than the status quo force. We again think of movement throughout the space in a more dynamic way, this time via a sequence of well-structured targets, rather than a single ideal target. Since emulators often change courses, this is an appealing option.

⁶¹There are other ways one could proceed—perhaps one might impose

$$\bigoplus_{i=1}^{m-1} d_\Theta(\mathcal{F}_i, \mathcal{F}_{i+1}) \leq_\Theta \beta_k \text{ for all } \beta_k \in B$$

so that the total restructuring costs are less than some total budget.

4.2 The Second Rationality

The force space being vast, a force-maker must choose one trick or another to navigate it. Emulation—choosing some target force and making changes to the source force to get closer to it—is one such trick. It is intuitive, flexible, and powerful, and the historical sketches above suggest that it is at least somewhat realistic. But, it is not the only trick in the book.

The target force in an emulation problem represents a concrete goal that might not always be available. It could well be that, in the absence of such concrete goals, the force-maker must rely on more abstract considerations. The word “abstract” suggests that we will be discussing something esoteric, but the opposite is true: it could well be that the force-maker uses sentences like

“My force is stronger than the enemy’s force.”

to guide her decisions. We need not bring the Prussian armies of yesteryear into the picture to understand this sentence, and yet we do not yet have a way to model it. Purpose-driven logic spoken in the language of force requires thinking through what steps are required to become *stronger* than the enemy or some potential enemy. The abstract consideration is a simple matter of comparing forces and issuing a judgment.

Remarkably, one can use an emulation problem to arrive at such binary comparisons—say, one force is stronger than another if it is closer to some ideal target. This logic has been at work throughout the previous section, albeit in the background. However, emulation targets are not necessary for the development of such a means of comparison, and surely they are not the only origins of such comparisons. The force-maker might have more immediate concerns in mind than the ideal types lionized by force-theorists. Though Machiavelli might have had Cæsar in mind when developing thoughts about which sorts of forces are best, Cæsar himself likely had no such template. Thus, the introduction of binary comparisons allows us to expand the set of motivations under study, albeit at a loss of fidelity to any given motivation.

Of course, binary comparisons play a central role in the theory of rational choice, where one of the main rationality fables takes as primitive a binary relation encoding either weak preference (“— is at least as good as —”) or strict preference (“— is better than —”). It turns out that some of the basic results one derives from the theory of rational choice can be applied to the problem of force-making, where the force-maker is the rational agent and the forces are the alternatives. In this sense, Sections 1 to 3 set the stage for traditional decision-theoretic thinking by developing an elaborate, comprehensive, structured set of alternatives, creating a traditional decision problem.

What should one call the means of comparison between two structured forces? If a force-maker prefers infantry-based forces to artillery-based forces, *ceteris paribus*, then what is the name we'd assign to the impetus behind this preference? Similar questions arise for highly-detailed org charts versus flexible org charts, or large officer corps versus small officer corps, or speed versus firepower, or navies versus armies. We use the following terminology.

4.10 Primitive (Doctrine)

There is a binary relation

$$\geq \subseteq \mathbb{F}_R^*(M_L^*) \times \mathbb{F}_R^*(M_L^*)$$

representing the force-maker's doctrine. We read " $\mathcal{F}_1 \geq \mathcal{F}_2$ " as " \mathcal{F}_1 is at least as forceful as \mathcal{F}_2 ."

From here one derives

$$\begin{aligned} \mathcal{F}_1 > \mathcal{F}_2 &\iff \mathcal{F}_1 \geq \mathcal{F}_2 \text{ and } \mathcal{F}_2 \not\geq \mathcal{F}_1, \\ \mathcal{F}_1 \sim \mathcal{F}_2 &\iff \mathcal{F}_1 \geq \mathcal{F}_2 \text{ and } \mathcal{F}_2 \geq \mathcal{F}_1. \end{aligned}$$

These are read " \mathcal{F}_1 is more forceful than \mathcal{F}_2 " and " \mathcal{F}_1 is as forceful as \mathcal{F}_2 ," respectively.

We can imagine handing the force-maker two structured forces and asking her to compare them, and she would respond with a statement comparing the forcefulness of the two. In calling this "doctrine," we run the risk of using a multifaceted term in a narrow way; for example, in his *Sources of Military Doctrine*, Barry M. Posen (1984, p. 13) uses the term to refer both to what means of force are employed and how they are employed. We are in good position to speak to his first sense—he specifically refers to the technologies uses, the structures of forces, and so on—but not the second, which (to repeat a theme from above) refers more to tactics than to the strategy of structure. Posen's examples often involve *combinations* of technologies and structures, and given the time we spent on force configuration and conversion, our theory is well-suited to encode these combinations. For example, Posen's exemplary offensive doctrine involves "the method of combining tanks, motorized infantry, and combat aircraft to achieve rapid victory invented by the Germans in the 1930s, and called Blitzkrieg ever since" (p. 14); one can imagine a \geq that tends to favor forces with these components to the detriment of forces with slow-moving infantry. Naturally, other doctrines might induce other \geq relations with different proclivities for different components: say, Swiss preferences for defensive specialization or American preferences for forces with a great deal of high-tech equipment. These preferences, and those of their kind, are what we mean by "doctrine."

The symbols just introduced are already familiar to political scientists given their central place in the theory of rational choice.⁶² But despite this being their natural home, it would be a mistake to think that they convey any sense of objectivity. The words “— is at least as forceful as —” use the word “forceful,” a value-laden term that is not neutral. That which is forceful in the eyes of one force-maker might not be forceful in the eyes of another, and this is the case for various reasons, some of which include:

1. *Strategic Context*. In *On War*, Carl von Clausewitz advocates for actions tailored for disarming the enemy as quickly and decisively as possible. This is a strategic matter, but it also has implications for the structure of the force: stronger organization to ensure constant pressure in the desired direction, more firepower to ensure that the enemy is disarmed quickly, and so on. Those who adopt this strategy use the word “forceful” differently than those who adopt a strategy of containment or a strategy of attrition. The strategic context is a key determinant of what is forceful.
2. *Branches*. In *The Influence of Sea Power Upon History*, Alfred Thayer Mahan (1890) argues that naval power is the key to global dominance. His work was influential in the development of the United States Navy, which has since been a key part of the United States’ military strategy. Mahan, of course, was not alone: for example, Giulio Douhet developed similar theories in favor of air power in *The Command of the Air* (1921). It is not hard to imagine similar arguments in favor of land, cyber, or space power.
3. *Logistics and Supply*. One force-maker might prefer a smaller structured force with dense support relationships, wanting only those units she knew she could keep well-supplied. Another might prefer a larger force with more units and fewer formal support relationships, wanting to be able to draw on a larger pool of resources or to live off the land in the name of maintaining speed and flexibility. Again, differences in doctrine lead to differences in what is forceful.
4. *Technology*. The introduction of a new technology can change the calculus of what is forceful. When the word “artillery” points to catapults and trebuchets, it is not the same as when it points to cannons and howitzers, which is not the same as when it points to missiles and drones. The same goes for infantry, cavalry, and so on. The force-maker’s doctrine will be influenced by the technologies available to her.

This list is not exhaustive, but it provides several key sources of subjectivity.

⁶²See, for example, David M. Kreps’s *Notes on the Theory of Choice* (1988).

The introductory theory of rational choice takes as primitive a collection of these subjective judgments stored in the relation \geq . The rationality part lies not in the objectivity of the judgments but rather in their consistency. The traditional postulates in this introductory theory are as follows.

4.11 Definition (Regular Doctrine)

We say that the doctrine \geq is weakly regular just in case it satisfies:

1. Reflexiveness: for all $\mathcal{F} \in \mathbb{F}_R^*(\mathbb{M}_L^*)$, we have $\mathcal{F} \geq \mathcal{F}$; and
2. Transitivity: for all $\mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_3 \in \mathbb{F}_R^*(\mathbb{M}_L^*)$,
$$\mathcal{F}_1 \geq \mathcal{F}_2 \text{ and } \mathcal{F}_2 \geq \mathcal{F}_3 \text{ imply } \mathcal{F}_1 \geq \mathcal{F}_3.$$

We say that the doctrine \geq is regular just in case it is weakly regular and satisfies:

3. Completeness: for all $\mathcal{F}_1, \mathcal{F}_2 \in \mathbb{F}_R^*(\mathbb{M}_L^*)$, we have $\mathcal{F}_1 \geq \mathcal{F}_2$ or $\mathcal{F}_2 \geq \mathcal{F}_1$.

Reflexiveness is the idea that a force is at least as forceful as itself, which is simply a check on the concept our language points to. The somewhat-awkward terminology “— is at least as forceful as —” naturally means to include the case where a force is compared to a copy of itself: you are at least as tall as yourself, at least as old as yourself, and so on. There is no real rationality bet here, but it remains an important and necessary condition. The main rationality bet is transitivity, which says that if one force is at least as forceful as another, and that force is at least as forceful as a third, then the first force is at least as forceful as the third. This is the consistency condition that makes the doctrine a useful tool for the force-maker, the thing that allows her to put the forces into some kind of order. One can construct a theory of rational choice from reflexivity and transitivity alone, but (as we will see) this falls short of some of the more interesting results in the theory. Completeness is the condition that ensures that the force-maker can always make a comparison between two forces, even if she is no more than indifferent between them. She may not respond with apathy, only indifference—and in particular, she may not say the forbidden words “I don’t know” when asked to compare two forces. Completeness is therefore a more stringent condition, one you and I fall short of when picking one tomato from a pile of tomatoes at the grocery store. But whereas you or I do not compare each pair of tomatoes due to constraints on our time, attention, and dexterity, it seems that the force-maker (and her team) are more likely of thinking through many more comparisons than we would be willing to endure. Of course, the sheer size of the force space makes this a daunting task, but it also stands to reason that she is capable of abiding by completeness in subspaces of the force space, where the number of comparisons is more manageable.

Emulation problems as studied above give rise to doctrines, and it is natural to ask what properties these doctrines have.

4.12 Lemma (Emulation Doctrine)

Consider an emulation problem parameterized by budget $\beta \in \Theta$, status quo force $\mathcal{S} \in \mathbb{F}_R^*(\mathbb{M}_L^*)$, and target force $\mathcal{T} \in \mathbb{F}_R^*(\mathbb{M}_L^*)$. Define the doctrine

$$\mathcal{F}_1 \geq_{\mathcal{T}} \mathcal{F}_2 \text{ if and only if } d_{\Theta}(\mathcal{F}_1, \mathcal{T}) \leq_{\Theta} d_{\Theta}(\mathcal{F}_2, \mathcal{T}).$$

Then, the doctrine $\geq_{\mathcal{T}}$ is weakly regular.

The doctrine $\geq_{\mathcal{T}}$ is regular if, and only if, \leq_{Θ} is complete on $\text{im } d_{\Theta}(\cdot, \mathcal{T})$.

Proof. This is straightforward.

1. *Reflexiveness:* choose any $\mathcal{F} \in \mathbb{F}_R^*(\mathbb{M}_L^*)$. Since $d_{\Theta}(\mathcal{F}, \mathcal{F}) = \mathbb{0}_{\Theta}$ by construction, we have $\mathcal{F} \geq_{\mathcal{T}} \mathcal{F}$.

2. *Transitivity:* choose any $\mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_3 \in \mathbb{F}_R^*(\mathbb{M}_L^*)$ such that

$$\underbrace{d_{\Theta}(\mathcal{F}_1, \mathcal{T}) \leq_{\Theta} d_{\Theta}(\mathcal{F}_2, \mathcal{T})}_{\mathcal{F}_1 \geq_{\mathcal{T}} \mathcal{F}_2} \quad \text{and} \quad \underbrace{d_{\Theta}(\mathcal{F}_2, \mathcal{T}) \leq_{\Theta} d_{\Theta}(\mathcal{F}_3, \mathcal{T})}_{\mathcal{F}_2 \geq_{\mathcal{T}} \mathcal{F}_3}.$$

Since \leq_{Θ} is transitive, we infer that $d_{\Theta}(\mathcal{F}_1, \mathcal{T}) \leq_{\Theta} d_{\Theta}(\mathcal{F}_3, \mathcal{T})$, which implies $\mathcal{F}_1 \geq_{\mathcal{T}} \mathcal{F}_3$.

3. *Completeness:* choose any $\mathcal{F}_1, \mathcal{F}_2 \in \mathbb{F}_R^*(\mathbb{M}_L^*)$. We handle both implications.

(a) Suppose \leq_{Θ} is complete on $\text{im } d_{\Theta}(\cdot, \mathcal{T})$. Then $d_{\Theta}(\mathcal{F}_1, \mathcal{T}) \leq_{\Theta} d_{\Theta}(\mathcal{F}_2, \mathcal{T})$ or $d_{\Theta}(\mathcal{F}_2, \mathcal{T}) \leq_{\Theta} d_{\Theta}(\mathcal{F}_1, \mathcal{T})$. From the definition of $\geq_{\mathcal{T}}$, we have $\mathcal{F}_1 \geq_{\mathcal{T}} \mathcal{F}_2$ or $\mathcal{F}_2 \geq_{\mathcal{T}} \mathcal{F}_1$.

(b) Suppose $\geq_{\mathcal{T}}$ is complete. Then we have $\mathcal{F}_1 \geq_{\mathcal{T}} \mathcal{F}_2$ or $\mathcal{F}_2 \geq_{\mathcal{T}} \mathcal{F}_1$. From the definition of $\geq_{\mathcal{T}}$, we have $d_{\Theta}(\mathcal{F}_1, \mathcal{T}) \leq_{\Theta} d_{\Theta}(\mathcal{F}_2, \mathcal{T})$ or $d_{\Theta}(\mathcal{F}_2, \mathcal{T}) \leq_{\Theta} d_{\Theta}(\mathcal{F}_1, \mathcal{T})$. Since \mathcal{F}_1 and \mathcal{F}_2 were arbitrary, this must be the case across all of $\text{im } d_{\Theta}(\cdot, \mathcal{T})$, which implies that \leq_{Θ} is complete on that set.

This completes the proof. ■

So, emulation problems generate weakly regular doctrines, but their ability to generate (fully) regular doctrines depends on the richness of the language of costs. Moving from real-world emulation to the wider class of abstract doctrines has not cost us anything, and the same problems apply in the two settings.

As it is a subset of $\mathbb{F}_R^*(M_L^*) \times \mathbb{F}_R^*(M_L^*)$ —which is to say, a subset of a product of two uncountable sets—the doctrine \geq is a large object. If we are to study the force-maker by observing her decisions, we must be able to make sense of this large object. One way to do this is to study the properties of the doctrine, which we have begun to do. Another way is to determine conditions under which the doctrine can be represented by a function, which is a more manageable object. Ideally, we would send the doctrine to a function that would output a real number, which we could then use to study the force-maker’s decisions. Having endured all of the pains of a non-complete, potentially-complicated order in the form of the quantales Ξ , Ψ , and Θ , the reals are a welcome respite.

Now, in the case of decision theory, we call a real-valued function that summarizes the comparisons from a binary preference a *utility function*. Utility functions are powerful tools, the very backbone of applied rational choice theory. But here, we are not talking about the desirability of an alternative—the thing pointed at by the word “good” in “— is at least as good as —”—but rather the forcefulness of a decorated org chart—the thing pointed at by the word “forceful” in “— is at least as forceful as —.”⁶³ We adopt the following terminology.

4.13 Definition (Force Scale)

By a force scale we mean a function

$$m_{\geq} : \mathbb{F}_R^*(M_L^*) \longrightarrow \mathbb{R}$$

that represents the doctrine \geq —i.e., for which we have

$$\mathcal{M}_1 \geq \mathcal{M}_2 \quad \text{if and only if} \quad m_{\geq}(\mathcal{M}_1) \geq m_{\geq}(\mathcal{M}_2).$$

In case such a function exists, we say that the doctrine \geq is scalable. In case \geq is scalable by a continuous function, we say that it is continuously scalable.

A force scale captures all relevant ordinal information contained in a doctrine while sending the humiliating complexity of $\mathbb{F}_R^*(M_L^*) \times \mathbb{F}_R^*(M_L^*)$ to the friendly confines of \mathbb{R} . It is hard to imagine $\mathbb{F}_R^*(M_L^*)$ on either the left- or right-hand side of a regression equation, whereas \mathbb{R} is the regressionist’s natural habitat. It would be good news indeed—though by no means a panaceum for issues of

⁶³Of course, life is not all that easy even in decision theory, as the word “good” is quite vague. Camerer (2006) distinguishes “wanting” from “liking,” which are two different ways of thinking about what is good. The path from either, or some mixture of the two, to “choosing” is fraught with peril, and the decision theorist is reminded of the importance of the distinction between “is” and “ought”—not to mention the importance of humility.

complexity—if the doctrine were scalable. Of course, continuous scalability is even better, as it allows us to transport even more information from the force space to the real line and back again.

Of course, the idea of using a single real number to capture a state's relative strength is not new. The Correlates of War project has long included the Composite Index of National Capability (CINC) as a measure of state capabilities; given our material focus with the concept of force, these capabilities are similar in spirit to the forces we have been discussing. Of course, quantitative researchers do not have the entire force space at their disposal, and so they must make do with the data they have. In a classic treatment, [Singer, Bremer and Stuckey \(1972\)](#) generate what remains the canonical measure of state capabilities by combining measures of six indicators: total population, urban population, iron and steel production, energy consumption, military personnel, and military expenditure. Since their domain is not the same as ours, it would be a mistake to say that the CINC score is a force scale as in Definition 4.13. However, one might use the numbers for similar reasons, reading $CINC_1 > CINC_2$ as "state 1 is at least as capable as state 2." The pin-dancing required to delineate "forcefulness" from "capability" is a task for another day, but it is clear that the two are related.⁶⁴

Remarkably, the six-dimensional real-valued space used by [Singer, Bremer and Stuckey \(1972\)](#) is similar to the force space in an important respect: it does not naturally have a total order. This would be the case even if capabilities were two-dimensional in their inputs: for example, if we looked only at military personnel and military expenditure, we would still have a two-dimensional space that does not naturally have a total order.⁶⁵ The force space is similar, but with (potentially) many more dimensions: we could, perhaps, record the number each type of molecule occurs in the force, the number of each type of unit, the number of each type of officer, and so on. Though highly-reductionist, this manner of thinking makes the linkages to CINC logic clear: we would have a spreadsheet with a countably infinite number of columns, each containing some real number. And, we could devise a function that reads in these numbers and outputs a real number, just as CINC reads in six numbers and outputs a real number.⁶⁶ The force-maker's doctrine is a more complex object than the CINC function, but the logic is the same. Whereas we cannot appeal to a scientific-looking formula to describe the relationship between the high-dimensional space and the scale, we can still use the same logic to understand the relationship between the two.

⁶⁴Moreover, the CINC score is not a force scale because (1) it is not derived from a doctrine, and (2) it is a function of proxies for capability rather than the forces themselves. The similarity here is more in terms of the rôle the numbers play in the decision-making process, their interpretations, and the reasonable operations one might perform on them.

⁶⁵By this we mean that there is no natural way to say that one point in the space is at least as capable as another. (1, 1) and (2, 2) might be easy to compare, but what about (1, 2) and (2, 1)? This is precisely the problem that pointed us to quantales in previous sections.

⁶⁶The CINC function is more nuanced than this, but not by too much.

In the absence of such a formula, we have no choice but to instead evaluate the scalability of a doctrine in terms of its properties. The following is a consequence of the regularity of the doctrine and the countability of the force space.

4.14 Proposition (Scalability of Regular Doctrines)

If \geq is a regular doctrine on $\mathbb{F}_R^(\mathbb{M}_L^*)$, then it is scalable.*

[*Proof.*]

This is a powerful result, as it means that the force-maker's doctrine can be represented by a real-valued function. For all of the complexity of the force space, the doctrine reduces it to a single dimension, which is a manageable object. Proposition 4.14 suggests that there is no inherent flaw in attempting to measure the forcefulness of a structured force, and that the force-maker's judgments can be captured by a single number.

Of course, there is a cost to this simplification, namely the behavioral assumptions required to ensure that the doctrine is scalable. As we were just discussing, any attempt to reduce the complexity of the force space to a real number requires a set of postulates that ensure that the real number is a faithful representation of the force-maker's judgments. Whereas the first rationality of the manuscript—compositional awareness, the ability to reason in terms of sequences of processes—is rather weak, this second rationality is far more demanding. The force-maker must be able to make consistent judgments about the forcefulness of structured forces, and these judgments must be complete. This must be accomplished not only in regions of the force space where the force-maker has some experience, but also in regions where she has no experience. This is just what unidimensional concepts like force compel us to do.

Proposition 4.14 is a powerful result, but it is not the end of the story. After all, we do not know whether small changes in a structured force—by which we mean changes that induce low restructuring and/or conversion costs—will lead to small changes in the force scale. To get at this question, we need to unify the order structure of the doctrine with the topology on the force space.

4.15 Definition (Continuous Regularity of a Doctrine)

We say \geq is continuously regular just in case it is regular and satisfies:

4. Continuity: for all structured forces $\mathcal{F} \in \mathbb{F}_R^*(\mathbb{M}_L^*)$, the sets

$$\text{Up}(\mathcal{F}) := \{\mathcal{G} \in \mathbb{F}_R^*(\mathbb{M}_L^*) \mid \mathcal{G} \geq \mathcal{F}\},$$

$$\text{Down}(\mathcal{F}) := \{\mathcal{G} \in \mathbb{F}_R^*(\mathbb{M}_L^*) \mid \mathcal{F} \geq \mathcal{G}\}$$

are closed with respect to the restructuring cost topology.

We will spend a few moments on this definition.

In words, the continuity condition says that the sets of structured forces that are at least as forceful as a given structured force and those that are less forceful than a given structured force are closed with respect to the restructuring cost topology. If we think about this in terms of sequential intuitions, it is not hard to see why this is a reasonable condition, at least in a broad sense. Suppose we had a sequence of pairs of forces, where for each pair, the force-maker judged the first to be at least as forceful as the second. And suppose these sequences of pairs converged to a pair of forces—*i.e.*, we eventually got to a point where the restructuring costs as we moved through the sequence got very small. The idea behind the continuity condition is that the force-maker’s judgment should be consistent with the limit of the sequence—if the force-maker judges the first force to be at least as forceful as the second in the limit, then she should judge the first force to be at least as forceful as the second in the limit. This makes sense in many relevant contexts, say if we were making infinitesimal changes to an org chart, adding soldiers here or there, or reloading a few cannons. If the force-maker’s judgment is not consistent with the limit, then we might worry that the force-maker is not making consistent judgments, which would be a problem for the doctrine. However, one can imagine reasonable cases where continuity is less acceptable, say if the second sequence eventually converges to a force with, say, fully-capable nuclear weapons, and the first sequence converges to a force with no nuclear weapons. The question, then, would be whether it is reasonable to suspect that small changes in the restructuring inputs could possibly yield such a large change in the force. After all, nuclear weapons are the product of many costly processes, and it is not clear that they could be added to a force without incurring large restructuring costs. This is a question for the force-maker, and it is not clear that the doctrine should be held to a standard that would require her to make consistent judgments in such cases.

Apologies made, we can now state the main result of the manuscript.

4.16 Proposition (Continuity of a Doctrine Implies Scalability)

If \geq is a continuously regular doctrine on $\mathbb{F}_R^ (\mathbb{M}_L^*)$, then it is continuously scalable.*

This result obtains from a very famous result due to Gerard Debreu (1964), the most general of the traditional representation theorems in the theory of rational choice. It relies on little more than the countability of the force space (for which we worked so hard in Sections 2 and 3) and the continuity of the doctrine. It asserts that small changes in the force space will lead to small changes in the force scale, which is a comforting result. Adding a few soldiers here or there, or changing the structure of a few units, will not lead to large changes in the force scale, so long as the force-maker is making consistent judgments.

Something that wouldn't have made sense had it been said in the beginning.

Your humble author has not been trying to keep things opaque, but rather to build up to a point. Make no mistake: Propositions 4.14 and 4.16 are not the end of the story, but they certainly are the climax of the manuscript.

The constructions of this manuscript are not the only way to think about the force space: we could have applied a different topology to the force space, for example, or just used one that came off the shelf. But the way we have done things adds a particular interpretation to Proposition 4.16. Suppose the force-maker has some preference relation over the force space, and that you and I were tasked with scaling this preference relation into a real-valued function. We would assign each structured force a real number, hoping to do so in a way that respects the force-maker's judgments. But, every now and again, we might announce a number the force-maker found to be absurd. "How could you have assigned \mathcal{F}_1 a 0 when you gave \mathcal{F}_2 a 1?" she might ask, insinuating that \mathcal{F}_1 is at least as forceful as \mathcal{F}_2 . This is, of course, a reasonable question for her to ask, and it is not up to her to abide by our mathematical formalism—indeed, *quite the opposite*. It is *we* who must ensure that our formalism respects her judgments, not the other way around.

This is precisely why the topology on the force space must be constructed in the terms of the costs *the force-maker pays* to restructure her forces. Our response to the force-maker might be something of the form "your points are valid, but at the very least you could change the forces a little bit to arrive at our conclusion, approximately." But approximately *to whom*? It's easy for us to say that the restructuring costs are small, but it is the force-maker who must pay them. The force-maker's judgments are the only ones that matter, and it is her judgments that must be respected. This is why:

1. We have not required that the force-maker is aware of every way of getting from one configuration to another, nor from one structured force to another—instead, we have only asked for compositional awareness. To have done otherwise would have been to set the force-maker up for failure. Imagine how she would have felt if we stated "you could change your forces a little bit to get from \mathcal{F}_1 to \mathcal{F}_2 " and she had no idea how to do so. The glibness of our statement would have been a slap in the face.
2. We have not imposed a language of cost, instead allowing the force-maker's reasoning to be in any language that satisfies the minimal requirements of a quantale. We even allowed her to speak in different languages about conversion versus restructuring costs. To have done otherwise would have been to impose our own understanding of the world on her, rather than allowing her to speak for herself.

3. We have constructed the weak metric in terms of her understanding of things spoken in the language of her understanding of costs. It is not up to us to discuss what is near and what is far in the context of a force-space we, quite frankly, will never navigate. It is up to the force-maker to decide what is near and what is far, and it is up to us to respect her judgments. Our task is to ensure that the force scale is a faithful representation of her judgments, not to impose our own understanding of the world on her.
4. We have constructed the force scale in terms of her judgments, rather than some absolute standard. There is no one doctrine that is correct, and while surely some are more useful than others, it is not up to us to decide which is which. It is up to the force-maker to decide what is forceful and what is not, and it is up to us to respect her judgments. Our task is to ensure that the force scale is a useful tool for understanding the force-maker's decisions, not to impose our own understanding of the world on her.

In other words, we have constructed the theory in terms of the force-maker's understanding of the world, rather than our own. The decisions we have made have been in the service of analytic humility, the idea that we should not impose our own understanding of the world on the force-maker, but rather allow her to speak for herself. This is the only way to ensure that the force scale is a faithful representation of the force-maker's judgments, and that the theory is a useful tool for understanding the force-maker's decisions. Of course, these decisions have made our life somewhat difficult, as the concepts at work are not the most familiar. Their introductions have been slow and deliberate, but it would have been impossible to justify them in the terms just enumerated *ex ante*. It is only now that we can see the full picture.

To be sure, we have occasionally proceeded as if we know things the force-maker does not. For example, we took as primitive the set of conversion morphisms between any two force configurations before compelling the force-maker to choose one element from each set. But, never did we say that we knew anything about these morphisms save for their existence, and never did we construct anything in terms of conversion morphisms save for those selected by the force-maker. It is an assumption—and not an empirically-verifiable one—that the morphisms can be stored in a set, and the Axiom of Choice plays no small role in the execution of the choice schedules. But, the force-maker is not required to know this to "discover" new technologies that generate new conversion possibilities, nor new ways of organizing her forces that generate new restructuring possibilities. We have respected her limitations while constructing the theory, but we have also respected her ability to transcend them. It is hoped this aside will clarify the humble spirit in which the theory was constructed.

4.3 Strategic Calibration

Before we wrap up, it might help to offer a dessert course to demonstrate the usefulness of the previous few results. After all, now that we have a real-valued scale representing the forcefulness of structured forces, we can use it to study those decisions the force-maker makes with forcefulness in mind. There is a robust literature in game theory on *strategic arming*, where two or more players decide how to arm themselves in the face of a potential conflict.

Many models in this literature employ contest models where the players' likelihood of military victory depends on the relative strength of their forces, which is endogenously determined.⁶⁷ For example, consider the following set-up from a subgame in [Beviá and Córchón \(2010\)](#).

4.17 Game (Contest Model of War)

Two states, $i \in \{1, 2\}$, decide how much to invest in their military forces, $m_i \in \mathbb{R}_+$. The probability that State i wins a war is given by

$$\pi_i(m_1, m_2) = \begin{cases} \frac{m_i}{m_1 + m_2} & \text{if } m_1 + m_2 > 0, \\ \frac{1}{2} & \text{if } m_1 + m_2 = 0. \end{cases}$$

The payoffs are given by

$$U_i(m_1, m_2) = \pi_i(m_1, m_2) V - \kappa(m_1 + m_2),$$

where $V \in \mathbb{R}_{++}$ is the value of winning the war and where $\kappa \in [0, 1]$ is the proportion of effort that cannot be recovered by the winner.

The game has a unique Nash equilibrium where

$$m_1^* = \frac{V}{4\kappa} = m_2^*.$$

Simple and elegant, the model captures the essence of strategic arming in a contest model. The players invest in their forces, and the player with the larger force wins the war with probability equal to the proportion of the total force she has. This maps rather naturally to a von-Neumann-Morgenstern utility function,

⁶⁷See, for example: [Garfinkel \(1990\)](#); [Grossman \(1991\)](#); [Hirshleifer \(1991\)](#); [Skaperdas \(1992\)](#); [Powell \(1993\)](#); [Meiowitz and Sartori \(2008\)](#); [Jackson and Morelli \(2009\)](#); [Beviá and Córchón \(2010\)](#); [Hodler and Yektaş \(2012\)](#); [Fearon \(2018\)](#). Other strategic-arming games exist where the decision is binary—say, to arm or not to arm; see, for example: [Baliga and Sjöström \(2004, 2008\)](#); [Debs and Monteiro \(2014\)](#); [Bas and Coe \(2016, 2018\)](#); [Coe and Vaynman \(2020\)](#).

where the player's utility is the probability of winning the war times the value of winning the war, minus the cost of investing in the force. The game's Nash equilibrium reflects the delicate balance between the marginal costs (κ) and marginal benefits (V) of arming as a strategic investment.

But what is m_i ? In the model, it is a non-negative real number, but what does it mean? What information does it encode? In a similar (albeit dynamic) model, [Fearon \(2018\)](#) continually refers to these as “force levels,” and this seems like a reasonable interpretation.⁶⁸ After all, for fixed m_2 , it is the case that

$$m_1 \geq m'_1 \quad \text{if and only if} \quad \pi_1(m_1, m_2) \geq \pi_1(m'_1, m_2),$$

suggesting that, in a very real sense, m_1 is at least as forceful as m'_1 —after all, it has a better chance of winning the war. Thus, it seems reasonable to interpret m_i as a measure of the forcefulness of State i ’s military forces, even if we don’t have any information about what those forces look like.

Given the symmetric set-up of the game,⁶⁹ it is unsurprising that the two states choose the same force level in equilibrium. But, does this mean that we should expect the two sides to arrive at the battlefield with *exactly* the same force? The same number of units, the same number of soldiers, the same number of tanks? The same uniforms, the same support relationships, the same training? Of course not, and it would be uncharitable to suggest that the model implies this. Well then, what *does* the model imply? It seems reasonable to argue that

$$\begin{aligned} \mathcal{F}_1^* &\in \left\{ \mathcal{F} \in \mathbb{F}_R^*(\mathbb{M}_L^*) \mid m_{\geq_1}(\mathcal{F}) = \frac{V}{4\kappa} \right\} = m_{\geq_1}^{-1}\left(\frac{V}{4\kappa}\right), \\ \mathcal{F}_2^* &\in \left\{ \mathcal{F} \in \mathbb{F}_R^*(\mathbb{M}_L^*) \mid m_{\geq_2}(\mathcal{F}) = \frac{V}{4\kappa} \right\} = m_{\geq_2}^{-1}\left(\frac{V}{4\kappa}\right), \end{aligned}$$

where \mathcal{F}_1^* and \mathcal{F}_2^* are the structured forces that State 1 and State 2, respectively, choose in equilibrium; where \geq_i is the doctrine representing State i ’s judgments; and where $m_{\geq_i}^{-1}$ is the preimage of the force scale under the force-maker’s doctrine. In other words, the model implies that the two states will choose structured forces that are at least as forceful as the force level that maximizes their utility. However, the subjective nature of the force scale—not the mention the varieties available within a given preimage of a force level—suggests that the two states will not arrive at the battlefield with exactly the same force.

⁶⁸ [Meiowitz and Sartori \(2008\)](#) specifically refer both to levels of force and other factors like the technologies available and the tactics employed. The first of these has a natural interpretation in the force space, whereas the latter is more difficult but manifests in things like organization and technologies. We cannot guarantee full coverage of the myriad concepts at work here, but it seems like the first few cuts are reasonable.

⁶⁹ [Beviá and Córchón \(2010\)](#) do not stop at the symmetric case, allowing for differences both in the relative technologies of the two states and the decisiveness of the war. Relative technologies seem best captured by the force space rather than the real-valued force scale, though it is not hard to imagine linkages between the two. Decisiveness of the war—the degree to which the side with the larger force can expect to win—is more definitively outside of the force space, and best left to applied modeling after the fact.

Can we go further? It would be good news indeed if the preimages derived above were nonempty, and we currently have no such guarantee. Indeed, because the force space is countable, the force scale cannot be surjective: its domain has strictly fewer elements than its codomain. Even if we restricted attention to the rational numbers \mathbb{Q} , there would be no guarantees that each element of the codomain would have a nonempty preimage. We could instead think in terms of approximations, where the force-maker's judgments are consistent with the force scale to within some small error—say,

$$m_i^{-1}\left(\left(\frac{V}{4\kappa} - \varepsilon, \frac{V}{4\kappa} + \varepsilon\right)\right) = \left\{ \mathcal{F} \mid \frac{V}{4\kappa} - \varepsilon < m_i(\mathcal{F}) < \frac{V}{4\kappa} + \varepsilon \right\}$$

which would be nonempty for large enough ε . This would suggest that the two states would choose structured forces that are at least as forceful as the force level that maximizes their utility, to within some small error. As a corollary to Proposition 4.16, we have the following.

4.18 Corollary (Preimages of Jittered Force Levels)

If \geq is a continuously regular doctrine on $\mathbb{F}_R^*(\mathbb{M}_L^*)$, then for all $\varepsilon > 0$, the preimages

$$m_i^{-1}\left(\left(\frac{V}{4\kappa} - \varepsilon, \frac{V}{4\kappa} + \varepsilon\right)\right) = \left\{ \mathcal{F} \mid \frac{V}{4\kappa} - \varepsilon < m_i(\mathcal{F}) < \frac{V}{4\kappa} + \varepsilon \right\}$$

are open in the restructuring cost topology.

Since they are open in the restructuring cost topology, the preimages can be written as unions of open balls in the restructuring cost topology—*i.e.*, small margin for error in the force scale implies small margin for error in the force space. This is a fresh manifestation of our humility comment above: in case the force-maker scoffed at being told to construct a force with a scale of *exactly* $V/4\kappa$, we can now say that she can construct a force with a scale of *approximately* $V/4\kappa$, where the terms of the word “approximately” are hers to define.

Now, the mere fact that we can cover the preimages with open balls does not mean that the preimages are nonempty unless ε is large enough. But, this is a entry point to the logic at work here, the logic of strategic calibration. The force-maker's judgments are not perfect, and the force scale is not a perfect representation of her judgments. However, thanks to the continuity of that scale under mild impositions on her doctrine, we can begin to form linkages between the simple, applicable world of force levels and the nuanced, complex world of structured forces. Moreover, the substantive terms of the construction bode well for our ability to determine what to expect from the force-maker's decisions, so long as we dedicate ourselves to understanding her understanding of costs and doctrine. Of course, that is a difficult task, but at least it is known.

So, how do we know that the situation is amenable to strategic calibration? The existence of Nash equilibria in games like Game 4.17 usually depends on one of two factors:

1. *Diminishing returns in the contest*: observe that the contest success function π_i is concave in m_i , so that increasing the force level continuous to improve one's chances, but at a decreasing rate. The costs of arming, however, are linear in the force level, so that the marginal cost of arming is constant. Eventually, the marginal cost of arming will exceed the marginal benefit, and the player will stop arming; this is the equilibrium force level. This is how the model in Game 4.17 works, and it is a common feature of contest models throughout the literature.
2. *Compactness of the strategy space*: in the absence of such structural headwinds, the existence of Nash equilibria can be guaranteed by the compactness of the strategy space. Force levels have a natural lower bound at 0, and it is natural—though by no means substantively necessary—to include it as an available strategy. Upper bounds represent resource constraints, whether they be financial, material, or political. The idea then goes that the force level is chosen from $[0, \bar{m}]$, where \bar{m} is the upper bound on the force level. The resulting strategy space, being a closed and bounded subset of the reals, is compact, and existence results follow from the Brouwer fixed-point theorem.

The first of these is very much real-numbers oriented: the concavity of a function, the linearity of costs, the additivity of costs and benefits—these all exploit the rich structure afforded by the real numbers. As such, we will not lean on them too heavily in our analysis, as they are not the most general of results. Compactness, of course, has been a recurring theme throughout our article, and it applies in general topological settings, not just in the real case.

To make the linkages clear, we must introduce what resource constraints look like in the force space.

4.19 Definition (Limited Capacities)

We say the force-maker has limited capacities if she may only choose forces out of some compact subset $K \subseteq \mathbb{F}_R^* (\mathbb{M}_L^*)$.

Alternatively, limited capacities might be interpreted that, outside of a compact subset, the costs of conversion and restructuring go to ∞ , so that the force-maker is effectively constrained to choose forces from a compact subset. This interpretation allows for the compactness to arise not just due to resources, but also the basic understanding of converting one military resource to another.

It is now straightforward to apply the results of the previous section to the strategic calibration of the force-maker's judgments. First, let us observe how the preimages of the force levels are affected by the limited capacities.

4.20 Proposition (Compact Force Preimages)

If \succsim is a continuously regular doctrine on $\mathbb{F}_R^(\mathbb{M}_L^*)$ and the force-maker has limited capacities at $K \subseteq \mathbb{F}_R^*(\mathbb{M}_L^*)$, then for all $x \in \mathbb{R}$, the preimage*

$$m_i^{-1}(x) = \{\mathcal{F} \in K \mid m_i(\mathcal{F}) = x\} \subseteq K \subseteq \mathbb{F}_R^*(\mathbb{M}_L^*)$$

is compact in the restructuring cost topology.

Proof. In case $m_i^{-1}(x)$ is empty, it is vacuously compact, so suppose it is nonempty. Being the preimage of the closed set $\{x\}$ under the continuous function m_i , $m_i^{-1}(x)$ is closed. Being a closed subset of a compact set, $m_i^{-1}(x)$ is compact. \blacksquare

This suggests that we retain a degree of analytic control over force scaling and strategic calibration: asking for the set of all forces with a given force level is a well-defined question, and the answer is a well-behaved object. The fact that the force space is so nuanced compared to the scaling space—despite the fact that the scaling space includes more elements—does not mean that the dual enterprises are incompatible.

The proof of Proposition 4.20 is a simple application of the fact that the preimage of a closed set under a continuous function is closed and that a closed subset of a compact set is compact. It is therefore important for us to write out the more general result, which we state without further proof.

4.21 Proposition (Compact Force Preimages II)

Under the terms of Proposition 4.20, for all $x \in \mathbb{R}$ and all closed $C \subseteq \mathbb{R}$, the preimage

$$m_i^{-1}(C) = \{\mathcal{F} \in K \mid m_i(\mathcal{F}) \in C\} \subseteq K \subseteq \mathbb{F}_R^*(\mathbb{M}_L^*)$$

is compact in the restructuring cost topology.

Again, the continuity of the force scale under the force-maker's doctrine allows us to make strong statements about the preimages of the force levels, even when the force-maker has limited capacities. This, in turn, makes it easier for us to think through what it would look like to unpack results obtained at the strategic level into the force space. The ideas at work beneath the surface are complex, but the results are clear and actionable. To repeat, it is the terms of the relationship between the two spaces that matter, not the spaces themselves.

We can write down precisely what is meant by strategic calibration.

4.22 Construction (Strategic Calibration)

A strategic calibration problem *takes the form*

$$\min_{\mathcal{T} \in m_i^{-1}(C)} d_{\Theta}(\mathcal{S}, \mathcal{T}),$$

where

1. $d_{\Theta} : \mathbb{F}_R^*(\mathbb{M}_L^*) \times \mathbb{F}_R^*(\mathbb{M}_L^*) \rightarrow \Theta$ is the calibrator's weak metric on structured forces calibrated by Θ ;
2. $\emptyset \neq C \subseteq \mathbb{R}$ is a closed set representing the calibrator's strategic target;
3. m_i is the force scale representing the calibrator's judgments; and
4. $\mathcal{S} \in m_i^{-1}(C)$ is the calibrator's status quo.

Notice the layers of cost at work here: in a game like Game 4.17, there was already a cost at work, namely the cost of destruction associated with the actual use of force. The marginal cost term κ plays a key role in determining the equilibrium force level, which in turn would be used to define the strategic target C . However, the costs of getting around the force space from a given status quo to a strategically-motivated target remain relevant, just as in the emulation problem. We therefore have distinguished the costs of the strategic decision from the more tactical costs associated with making the decision a reality. The strategic target might be a single force level—say, $V/4\kappa$, or any other equilibrium arming level—or it might be a closed set of them.⁷⁰

The following result is a direct application of emulation logic to the present problem, taking advantage of the compact structure of the preimages of the force levels—and, of course, the continuity of the weak metric.

4.23 Proposition (Strategic Calibration Problems Have Solutions)

Under the terms of Proposition 4.20, every strategic calibration problem has a solution.

⁷⁰In many games like this, the theorem of the maximum delivers the useful result that the correspondence sending parameters to equilibrium arming levels has a closed graph. This entire graph could be fed into a strategic calibration problem, where the calibrator's status quo is the current equilibrium arming level.

Thus, we have unearthed yet another application of the force space, this time in the realm of strategic calibration. This one works hand-in-glove with the idea of a doctrine, but its structure is identical to the non-doctrinal calibration problem. This dessert course, then, is a fitting end to the meal, a demonstration of the utility of the force space in understanding the force-maker's decisions.

5 What is Force?

[...] the previous submission of force to methods of reason. Civilisation is nothing else than the attempt to reduce force to being the ultima ratio. We are now beginning to realise this with startling clearness, because "direct action" consists in inverting the order and proclaiming violence as *prima ratio*, or strictly as *unica ratio*. It is the norm which proposes the annulment of all norms, which suppresses all intermediate process between our purpose and its execution. It is the *Magna Charta of barbarism*.

José Ortega y Gasset, *The Revolt of the Masses* (1930)

This manuscript began with a simple question: What is force? Given the centrality of the question in the introduction, it is fitting that the conclusion should return to it. Of course, given the nature of the answer proffered in this manuscript, simple summaries are exceedingly difficult. We have explored the concept not through definitions, nor even through historical examples, but rather by developing a theory of force and an attendant mathematical model. In your humble author's view, that theory and that model *are* the answer to the question. To describe force is to employ the concepts named by the theory and encoded in the model, nothing more and nothing less.

The section headings of this manuscript have been chosen to reflect basic stylized facts that appear to be true of force. These facts are exceedingly simple:

- § Force is atomic, arriving in discrete packets.
- § Force is a resource, capable of being collected into a setting and capable of being converted from one form to another.
- § Force is organized, consisting not just of collections of force but of collections of force that are organized into structures.
- § Force is a concept, something that can be manipulated, reasoned about, and used to simplify aspects of the world.

The model developed in this manuscript reflects these facts, which—despite their apparent simplicity—gives rise to a rich and complex theory of force. We have used a wide variety of techniques to explore the model, borrowing concepts from order theory, category theory, graph theory, algebra, topology, and the theory of metric spaces, not to mention a smattering of the usual suspects from the social sciences: decision theory, game theory, and the theory of organizations. This has not been done for the sake of it, but rather because the concept of force has pointed us in these directions. It is remarkable what one must learn in order to understand something as simple as force.

We have taken a particular approach to answering the ontological question of force, and this deserves some reflection. We took inspiration from Quine's criterion that "to be is to be the value of a variable" and applied it to the concept of force—*i.e.*, we have answered the question of what force is by asking what class of things accommodates it. In a sense, our answer to what force is

A *force* is a finite partial order representing an organizational chart, along with a countable set 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.

This is essentially what we have encoded as a structured force, and the set of all such things— $\mathbb{F}_R^*(\mathbb{M}_L^*)$ —is the variable that takes on the value of force. Thus, to answer the ontological question of force reduces to answering questions about how one ought to structure the information one uses when reasoning about force—literally, what do you need to know to describe a force in its most essential, most basic form? What are the columns in the data-set you would use to study force, and what values does each column take?

But is that all we've done, is constructed a set of partial orders with some fancy décor? Perhaps so. And yet, the facts just enumerated encouraged us to look into particular properties of these structures, and in so doing we unearthed a variety of insights that seem more than just the sum of their parts. Much of this was done in the service of constructing a reasonable topology for the set of things called force, which (as we've just seen) was done in the service of understanding the conditions under which force can be used as a self-simplifying concept. And because we wanted to know when this was the case, the topology we constructed needed to bear properties that made sense, given the stylized facts we had identified. We needed the space between two forces to represent the difficulty of transforming one into the other, since otherwise we would run the risk of committing analytic hubris. This pointed us toward thinking about force like a resource, which in turn committed us toward particular formal structures. The continual cycle between the facts and the theory, the theory and the facts—not to mention the model and the facts and the theory and the model and the...—is the fountainhead of those few insights we have uncovered. Precisely because we wanted to know what it would take to know something about force, we have learned something about force, and perhaps about knowing. All the while, we have felt the continual pull toward humility, making sure not to impose too much on the concept of force, but rather to let it guide us in our thinking. We trace the contours of force, but we do not dictate them.

This manuscript, then, does not contain an answer to the question of what force is.⁷¹ Instead it is meant to be a guide to thinking about force, an experience of mind that points toward the concept, rather than pinning it down in a box. But that experience of mind, the structures one encounters and the themes that emerge, is itself a kind of answer. If the reader's experience has been anything like the author's, one of these themes will have been the sense of beauty that emerges from the theory of force. Its essential discreteness, its natural fractal-like structure, the hills and valleys of the metric space linking all forces, the processes sending one force to another, the continuity with judgment, the parallel between conversion and restructuring, the unexpected relationship between rationality and distance, the structure of the force-maker's reasoning—all of these have a kind of beauty to them. Had we been theorizing a different concept, we might have found other things to be beautiful, but they would not be the same as the beautiful things we found here. Absurd as it might seem, there is something humiliatingly beautiful *about force* that invites further reflection. And though that beauty has not been defined by the abscissas and ordinates we have used to describe it, it has been illuminated by them in the course of this reflection. It is hoped that the ignorance recorded in this manuscript will help others to bear their own ignorance of fact, and to find beauty even in the ugliest of things.

⁷¹Whoops.

A Proofs

This section contains those proofs not included in the main text.

2.5 Lemma (Graph Unions Configure Molecules)

*Graph union satisfies the properties of Primitive 2.1.*⁷²

[*Proof.*]

Proof. We address each property in turn.

1. *Unitality.* Choose any $M = (n, E, \ell : \underline{n} \rightarrow L) \in \mathbb{M}_L$. We have

$$\begin{aligned} M \uplus (0, \emptyset, \emptyset) &= (n + 0, E \sqcup \emptyset, \ell : \underline{n+0} \rightarrow L), \\ &= (n, E, \ell : \underline{n} \rightarrow L) = M. \end{aligned}$$

The proof for the other side is similar.

2. *Associativity.* Choose any

$$\begin{aligned} M_1 &= (n_1, E_1, \ell_1 : \underline{n_1} \rightarrow L), \\ M_2 &= (n_2, E_2, \ell_2 : \underline{n_2} \rightarrow L), \\ M_3 &= (n_3, E_3, \ell_3 : \underline{n_3} \rightarrow L) \in \mathbb{M}_L. \end{aligned}$$

There is no loss of generality in assuming that each of these molecules is nonempty, as the empty molecule is the identity element for the \uplus operator. We handle each molecular component in turn:

- (a) *Size.* Molecule $(M_1 \uplus M_2) \uplus M_3$ has size $(n_1 + n_2) + n_3$, and molecule $M_1 \uplus (M_2 \uplus M_3)$ has size $n_1 + (n_2 + n_3)$. Addition of natural numbers being associative, these sizes are equal.
- (b) *Edges.* We first construct $(E_1 \sqcup E_2) \sqcup E_3$ from its components:

$$\begin{aligned} X_{11} &:= \{(i, j) \mid i, j \in \underline{n_1}\}, \\ X_{12} &:= \{(i + n_1, j + n_1) \mid i, j \in \underline{n_2}\}, \\ X_{13} &:= \{(i + (n_1 + n_2), j + (n_1 + n_2)) \mid i, j \in \underline{n_3}\}, \end{aligned}$$

⁷²The proof of Lemma 2.5 is straightforward, but it is also a bit tedious and not very illustrative; as such, it is our first proof relegated to Appendix A. The reader ought to be able to prove Lemma 2.5 with a little effort, and indeed it is a good exercise.

where we have grouped the terms in X_{13} to keep the order of operations clear. Then we have

$$\begin{aligned}(E_1 \sqcup E_2) \sqcup E_3 &= (X_{11} \cup X_{12}) \cup X_{13}, \\ &= X_{11} \cup X_{12} \cup X_{13},\end{aligned}$$

where we may drop the parentheses because set union is associative. We now construct $E_1 \sqcup (E_2 \sqcup E_3)$ from its components:

$$\begin{aligned}X_{21} &:= \{(i, j) \mid i, j \in \underline{n}_1\}, \\ X_{22} &:= \{(i + n_1, j + n_1) \mid i, j \in \underline{n}_2\}, \\ X_{23} &:= \{(i + n_2) + n_1, (j + n_2) + n_1 \mid i, j \in \underline{n}_3\},\end{aligned}$$

where we have again grouped the terms in X_{23} to keep the order of operations clear. Then we have

$$\begin{aligned}E_1 \sqcup (E_2 \sqcup E_3) &= X_{21} \cup (X_{22} \cup X_{23}), \\ &= X_{21} \cup X_{22} \cup X_{23},\end{aligned}$$

where again we may drop the parentheses because set union is associative. The question therefore reduces to showing that $X_{11} \cup X_{12} \cup X_{13} = X_{21} \cup X_{22} \cup X_{23}$; since $X_{11} = X_{21}$ and $X_{12} = X_{22}$ literally, we need only show that $X_{13} = X_{23}$. Addition of natural numbers being both associative and commutative, we have the required

$$\begin{aligned}i + (n_1 + n_2) &= (i + n_2) + n_1, \\ j + (n_1 + n_2) &= (j + n_2) + n_1.\end{aligned}$$

Thus, the edge sets are equal. ■

(c) *Labels*. We define

$$\begin{aligned}(\ell_1 \sqcup \ell_2)(i) &= \begin{cases} \ell_1(i) & \text{if } i \leq n_1, \\ \ell_2(i - n_1) & \text{if } i > n_1, \end{cases} \\ ((\ell_1 \sqcup \ell_2) \sqcup \ell_3)(i) &= \begin{cases} (\ell_1 \sqcup \ell_2)(i) & \text{if } i \leq n_1 + n_2, \\ \ell_3(i - (n_1 + n_2)) & \text{if } i > n_1 + n_2. \end{cases}\end{aligned}$$

Similarly, we define

$$\begin{aligned}(\ell_2 \sqcup \ell_3)(i) &= \begin{cases} \ell_2(i) & \text{if } i \leq n_2, \\ \ell_3(i - n_2) & \text{if } i > n_2, \end{cases} \\ (\ell_1 \sqcup (\ell_2 \sqcup \ell_3))(i) &= \begin{cases} \ell_1(i) & \text{if } i \leq n_1, \\ (\ell_2 \sqcup \ell_3)(i - n_1) & \text{if } i > n_1. \end{cases}\end{aligned}$$

We need to show that $(\ell_1 \sqcup \ell_2) \sqcup \ell_3 = \ell_1 \sqcup (\ell_2 \sqcup \ell_3)$, which reduces to showing that they have the same output for all inputs. Choose any $i \in \underline{n_1 + n_2 + n_3}$. We study three cases.

i. In case $i \leq n_1$, we have

$$\begin{aligned} ((\ell_1 \sqcup \ell_2) \sqcup \ell_3)(i) &= (\ell_1 \sqcup \ell_2)(i), \\ &= \ell_1(i), \\ &= (\ell_1 \sqcup (\ell_2 \sqcup \ell_3))(i), \text{ as required.} \end{aligned}$$

ii. In case $n_1 < i \leq n_1 + n_2$, we have

$$\begin{aligned} ((\ell_1 \sqcup \ell_2) \sqcup \ell_3)(i) &= (\ell_1 \sqcup \ell_2)(i), \\ &= \ell_2(i - n_1), \\ &= (\ell_2 \sqcup \ell_3)(i - n_1), \\ &= (\ell_1 \sqcup (\ell_2 \sqcup \ell_3))(i), \text{ as required.} \end{aligned}$$

iii. In case $n_1 + n_2 < i \leq n_1 + n_2 + n_3$, we have

$$\begin{aligned} ((\ell_1 \sqcup \ell_2) \sqcup \ell_3)(i) &= \ell_3(i - (n_1 + n_2)), \\ &= (\ell_2 \sqcup \ell_3)(i - n_2), \\ &= (\ell_1 \sqcup (\ell_2 \sqcup \ell_3))(i), \text{ as required.} \end{aligned}$$

These cases being exhaustive, we conclude the labels are equal. \square

We have shown that all three components of the force molecules are equal, and we conclude that the graph union satisfies Associativity. \square

3. *Commutativity.* Choose any

$$\begin{aligned} M_1 &= (n_1, E_1, \ell_1 : \underline{n_1} \rightarrow L), \\ M_2 &= (n_2, E_2, \ell_2 : \underline{n_2} \rightarrow L) \in \mathbb{M}_L. \end{aligned}$$

We must demonstrate that $M_1 \uplus M_2 \cong M_2 \uplus M_1$. We again handle each molecular component in turn:

(a) *Size.* $n_1 + n_2 = n_2 + n_1$ by commutativity of $+$ on \mathbb{N} . \square

(b) *Edges.* The first edge set is

$$E_1 \sqcup E_2 = \{(i, j) \mid i, j \in \underline{n_1}\} \cup \{(i + n_1, j + n_1) \mid i, j \in \underline{n_2}\},$$

whereas the second edge set is

$$E_2 \sqcup E_1 = \{(i, j) \mid i, j \in \underline{n}_2\} \cup \{(i + n_2, j + n_2) \mid i, j \in \underline{n}_1\}.$$

These sets are not literally equal, but we only need them to be isomorphic. Being two sets, they are isomorphic just in case they have the same cardinality. Since the sets are disjoint by construction, we begin our counting like so:

$$\begin{aligned} |E_1 \sqcup E_2| &= |E_1| + |E_2|, \\ &= n_1^2 + n_2^2, \\ &= n_2^2 + n_1^2, \\ &= |E_2| + |E_1|, \\ &= |E_2 \sqcup E_1|, \end{aligned}$$

where the second equality follows from commutativity of $+$ on \mathbb{N} . Thus, the edge sets are isomorphic.

(c) *Labels*. We have

$$\begin{aligned} \tilde{\ell}_1(i) &= \begin{cases} \ell_1(i) & \text{if } i \leq n_1, \\ \ell_2(i - n_1) & \text{if } i > n_1, \end{cases} \\ \tilde{\ell}_2(i) &= \begin{cases} \ell_2(i) & \text{if } i \leq n_2, \\ \ell_1(i - n_2) & \text{if } i > n_2, \end{cases} \end{aligned}$$

and we must show that $\tilde{\ell}_1 = \tilde{\ell}_2$. Choose any $i \in \underline{n}_1 + \underline{n}_2$. We study two cases.

i. In case $i \leq n_1$, we have

$$\begin{aligned} \tilde{\ell}_1(i) &= \ell_1(i), \\ &= \ell_2(i), \\ &= \tilde{\ell}_2(i), \text{ as required.} \end{aligned}$$

ii. In case $n_1 < i \leq n_1 + n_2$, we have

$$\begin{aligned} \tilde{\ell}_1(i) &= \ell_2(i - n_1), \\ &= \ell_1(i - n_2), \\ &= \tilde{\ell}_2(i), \text{ as required.} \end{aligned}$$

These cases being exhaustive, we conclude the labels are equal. \blacksquare

We have shown that all three components of the force molecules are equal, and we conclude that the graph union satisfies Commutativity. \blacksquare

We have shown that the graph union satisfies the properties in Primitive 2.1, so we are done. [[Back to the text.](#)]

2.8 Lemma (Graph Union and Deconfiguration)

For all force configurations $\biguplus \mathcal{M} \in \mathbb{M}_L^\star$, there exists a set $\{\pi_i : \biguplus \mathcal{M} \rightarrow M_i\}_{i=1}^n$ as in Primitive 2.7. [[Proof.](#)]

Proof. Choose a force configuration $\biguplus \mathcal{M}$; there is no loss of generality in assuming that this configuration is nonempty and no real loss of generality in assuming it is disconnected—*i.e.*, it is made up of more than one molecule. The whole idea of the lemma is that we know neither the number of molecules in the configuration nor the identities of the molecules. However, we do know that $\biguplus \mathcal{M}$ is the output of graph union on a sequence of force molecules $\mathcal{M} = M_1, \dots, M_k$ for some k . We will refer to the “actual” molecules, numbers, and labels with undecorated notations, whereas the versions obtained from deconfiguration will be decorated with hats—*i.e.*, k is the actual number of force molecules (which we do not know), and \hat{k} is the number of molecules obtained from the deconfiguration. We will refer to the “undecomposed” version of the configuration with bars, so that we are initially given

$$\biguplus \mathcal{M} = (\bar{n}, \bar{E}, \bar{\ell} : \{1, \dots, \bar{n}\} \rightarrow L).$$

We know that this can be written

$$\biguplus \mathcal{M} = \left(\sum_{i=1}^k n_i, \bigsqcup_{i=1}^k E_i, \ell : \sum_{i=1}^k n_i \rightarrow L \right),$$

but we do not know the value of k , nor any of the n_i s, nor any of the E_i s, nor how to label the vertices.

Step 1: Graph decomposition. Let us first consider the unlabeled graph:

$$G_{\biguplus \mathcal{M}} = (V_{G_{\biguplus \mathcal{M}}} = \{1, \dots, \bar{n}\}, E_{G_{\biguplus \mathcal{M}}} = \bar{E}),$$

where the first part is the set of vertices and the second part is the set of edges. Our first task is to identify its set of maximal connected subgraphs, which (as it happens) is a well-studied problem in graph theory.⁷³

⁷³The reader is encouraged to consult Reinhard Diestel’s *Graph Theory* (1998), Douglas B. West’s

We proceed with a *depth-first search* algorithm to decompose the graph into its maximal connected subgraphs. The algorithm proceeds as follows:

0. *Initialization*: we define two sets and fix their starting values:

VISITED $\leftarrow \emptyset$,
COMPONENTS $\leftarrow \emptyset$.

The VISITED set will store the vertices we have visited, and the COMPONENTS set will store the subgraphs we have found.

1. *Selection of Starting Vertex*: we select some $v \in V_{G \sqcup \mathcal{M}} \setminus \text{VISITED}$.
2. *Initialization of Traversal*: we define two sets and fix their starting values:

STACK $\leftarrow \{v\}$, CURRENT $\leftarrow \{v\}$.

The STACK will store the vertices on our to-explore list, and the CURRENT set will store the vertices in the current connected component.

3. *Traversal*: WHILE $\text{STACK} \neq \emptyset$, we do the following:

- (a) *Pop Vertex*: we pop a vertex u from the stack:

$u \leftarrow \text{POP}(\text{STACK})$,

where POP removes and returns the first element of a set.⁷⁴

- (b) *Mark the Popped Vertex*: we mark the vertex as visited and as part of the current connected component:

VISITED $\leftarrow \text{VISITED} \cup \{u\}$, CURRENT $\leftarrow \text{CURRENT} \cup \{u\}$.

- (c) *Push Neighbors*: we push all unvisited neighbors of u onto the STACK:

$$\begin{aligned} \text{ADJ}_u &= \left\{ z \in V_{G \sqcup \mathcal{M}} \mid (u, z) \in E_{G \sqcup \mathcal{M}} \text{ or } (z, u) \in E_{G \sqcup \mathcal{M}} \right\}, \\ \widetilde{\text{ADJ}}_u &= \text{ADJ}_u \setminus \text{VISITED}, \\ \text{STACK} &\leftarrow \text{STACK} \cup \widetilde{\text{ADJ}}_u. \end{aligned}$$

Introduction to Graph Theory (1996), or Béla Bollobás's *Modern Graph Theory* (1998) for thorough treatments of the subject; we will only use basic concepts from introductory chapters in this text. One can write entire books on the subject of the algorithms, and indeed many have been written. The clearest introduction is likely Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein's *Introduction to Algorithms* (1990), recent editions of which include a multi-chapter part on graphs and graph algorithms.

⁷⁴Since the vertices are natural numbers, we can use the fact that the first element of a set is the smallest element. So here, POP picks out the smallest element of the set whilst removing it.

The vertex set being finite, this process will terminate.

4. *Record Component*: we record the output of the traversal:

$\text{COMPONENTS} \leftarrow \text{COMPONENTS} \cup \text{CURRENT}.$

5. *Repeat*: WHILE $\text{VISITED} \neq V_{G \cup \mathcal{M}}$, we repeat Steps 1–4.

The vertex set being finite, this process will terminate.

Step 2: The output of the graph decomposition is a set of connected subgraphs. Choose any $\widehat{V} \in \text{COMPONENTS}$, where this is the vertex set from a “candidate molecule” obtained from the graph decomposition. Choose any $u, v \in \widehat{V}$; we must show that there is a path from u to v . Two vertices are in the same \widehat{V} only when they are part of the same traversal; this is by construction. Without loss of generality, we suppose u was put into the STACK before v . By construction, there exists a sequence of vertices $u = u_1, \dots, u_m = v$ such that u_i was POPped from the STACK before u_{i+1} . Thus, the sequence of edges $(u_1, u_2), \dots, (u_{m-1}, u_m)$ is a path from u to v ; we conclude that \widehat{V} is connected.

Step 3: In fact, the output of the graph decomposition is a set of maximal connected subgraphs. Suppose, for sake of contradiction, that there is some $\widehat{V} \in \text{COMPONENTS}$ that is not maximal. Then there is some $u \in \widehat{V}$ and some $v \in V_{G \cup \mathcal{M}} \setminus \widehat{V}$ such that there is a path from u to v . The associated sequence of edges is written

$(u_1, u_2), \dots, (u_{m-1}, u_m),$

where $u = u_1$ and $v = u_m$. Without loss of generality, we suppose $u_i \in \widehat{V}$ for all $i < m$, so that v is the first vertex in the sequence not in \widehat{V} . But since v is adjacent to u_{m-1} , it would have been part of $\text{ADJ}_{u_{m-1}}$ and therefore would have been added to the STACK when u_{m-1} was processed. Once in the STACK, v would have been POPped and added to the CURRENT set, and thus to \widehat{V} . This is a contradiction, so we conclude that all $\widehat{V} \in \text{COMPONENTS}$ are maximal.

Step 4: Formatting the output of the graph decomposition. We have a set of maximal connected vertex sets $\text{COMPONENTS} = \{\widehat{V}_1, \dots, \widehat{V}_{\hat{k}}\}$, where

$$\bigcup_{i=1}^{\hat{k}} \widehat{V}_i = V_{G \cup \mathcal{M}},$$

$$i \neq j \implies \widehat{V}_i \cap \widehat{V}_j = \emptyset.$$

To get to correctly-formatted molecules, we need to assign numbers to the vertices in each \widehat{V}_i . We will write $\widehat{V}_i = \{v_{i,1}, \dots, v_{i,\widehat{n}_i}\}$, where \widehat{n}_i is the number of vertices in \widehat{V}_i ; this is clearly isomorphic to $\{1, \dots, \widehat{n}_i\}$, and we call the associated bijection $\psi_i : V_i \rightarrow \{1, \dots, \widehat{n}_i\}$. We can write the correctly-formatted vertex set as

$$\underline{\widehat{n}_i} = (\psi_i(v_{i,1}), \dots, \psi_i(v_{i,\widehat{n}_i})).$$

We can now write the correctly-formatted edge set as

$$\widehat{E}_i = \{(\psi_i(u), \psi_i(v)) \mid (u, v) \in E_{G_{\mathcal{U}, \mathcal{M}}}\}.$$

As for the labels, we define each $\widehat{\ell}_i : \underline{\widehat{n}_i} \rightarrow L$ by restricting the original $\bar{\ell}$ to the vertices in \widehat{V}_i :

$$\widehat{\ell}_i = \bar{\ell} \circ \psi_i|_{\widehat{V}_i}.$$

We can now write the correctly-formatted molecule as

$$\widehat{M}_i = (\widehat{n}_i, \widehat{E}_i, \widehat{\ell}_i).$$

This is our candidate molecule, and since i was arbitrary, we have a set of candidate molecules $\{\widehat{M}_1, \dots, \widehat{M}_{\widehat{k}}\}$.

Step 5: For all identified molecules \widehat{M}_i , we have $\widehat{M}_i \cong M_i$. Finally, we must show that the molecules we have identified are isomorphic to the original molecules in the sense of Remark 1.4. We have $\widehat{M}_i = (\widehat{n}_i, \widehat{E}_i, \widehat{\ell}_i)$ and $M_i = (n_i, E_i, \ell)$. We must show that there exists a bijection $\varphi_i : \underline{\widehat{n}_i} \rightarrow \underline{n_i}$ such that $\widehat{E}_i = \{(\varphi_i(u), \varphi_i(v)) \mid (u, v) \in E_i\}$ and $\widehat{\ell}_i = \bar{\ell} \circ \psi_i|_{\widehat{V}_i}$. We have already constructed such a bijection ψ_i and shown that it is an isomorphism of graphs, so the isomorphism of molecules follows.

Step 6: Conclusion. We have shown that the graph union operator is deconfigurable, and we have constructed a set of deconfigurations that recover the original molecules from a configuration. [[Back to the text](#).] ■

3.6 Proposition (The Set of All Structured Forces)

$\mathbb{F}^*(\mathbb{M}_L^*)$ is countable.

[[Proof](#).]

Proof. $\mathbb{F}^*(\mathbb{M}_L^*)$ is a subset of a set isomorphic to

$$\bigcup_{n \in \mathbb{N}} \mathbf{Part}(\underline{n}) \times (\mathbb{M}_L^*)^{\underline{n}},$$

where $(\mathbb{M}_L^*)^{\underline{n}}$ is the set of all functions from \underline{n} to \mathbb{M}_L^* . Choose any $n \in \mathbb{N}$.

1. The set of all partial orders over \underline{n} is a subset of the set of all relations on $\underline{n} \times \underline{n}$, which has cardinality $2^{\binom{\underline{n}^2}{2}}$; this provides a finite upper bound on the cardinality of $\mathbf{Part}(\underline{n})$.
2. The set of ways to send \underline{n} to \mathbb{M}_L^* has cardinality $|\mathbb{M}_L^*|^{\underline{n}} \leq \aleph_0^n = \aleph_0$.

Thus, the set we take unions over is the product of a finite set and a countable set, which is countable. The countable union of countable sets is countable, so we conclude that $\mathbb{F}^*(\mathbb{M}_L^*)$ is countable. [*Back to the text.*] ■

3.10 Lemma (Countability of Non-Command Relations)

For all finite V , the set of non-command relations on V is countable.

Proof. We will show that $\mathbb{F}_R^*(\mathbb{M}_L^*)$ is countable by showing that it is a Cartesian product of countable sets. Choose any $k \in \mathbb{N}$, and let

$$\mathcal{V}_k := \{R^k : \underline{k} \rightarrow V\} = V^{\underline{k}}.$$

Evidently, $|\mathcal{V}_k| = |V|^{\underline{k}}$; since V is a finite set and k a finite number, $|\mathcal{V}_k|$ is finite. Define

$$\mathcal{V} := \bigcup_{k \in \mathbb{N}} \mathcal{V}_k,$$

which (being a countable union of finite sets) is countable. Thus, the set of all ways to assign relationships to the force structure is countable.

Observe that

$$\mathbb{F}_R^*(\mathbb{M}_L^*) := \mathbb{F}^*(\mathbb{M}_L^*) \times \mathcal{V},$$

which is the Cartesian product of two countable sets (the countability of $\mathbb{F}^*(\mathbb{M}_L^*)$ was shown in Proposition 3.6). Thus, $\mathbb{F}_R^*(\mathbb{M}_L^*)$ is countable. [*Back to the text.*] ■

3.19 Proposition (Subconfigurationhood is a Preorder)

Subconfigurationhood is a preorder on \mathbb{M}_L^ .*⁷⁵

[*Proof.*]

Proof. We check the properties in turn.

1. *Reflexivity:* Let $\biguplus \mathcal{M} = \biguplus_{i=1}^n M_i \in \mathbb{M}_L^*$, and set $\iota = \text{id}_{\{1, \dots, n\}}$; the identity, being an injection, satisfies the conditions of the definition. Then $M_i = M_{\iota(i)}$ for all $i \in \{1, \dots, n\}$, so $\biguplus \mathcal{M} \preccurlyeq \biguplus \mathcal{M}$. \blacksquare
2. *Transitivity:* Suppose $\biguplus \mathcal{M}_1 \preccurlyeq \biguplus \mathcal{M}_2$ and $\biguplus \mathcal{M}_2 \preccurlyeq \biguplus \mathcal{M}_3$. We need to show that $\biguplus \mathcal{M}_1 \preccurlyeq \biguplus \mathcal{M}_3$. Let $\iota_{1 \rightarrow 2} : \{1, \dots, n_1\} \hookrightarrow \{1, \dots, n_2\}$ and $\iota_{2 \rightarrow 3} : \{1, \dots, n_2\} \hookrightarrow \{1, \dots, n_3\}$ be the injections witnessing the subconfigurationhood. Then the composition $\iota_{1 \rightarrow 3} = \iota_{2 \rightarrow 3} \circ \iota_{1 \rightarrow 2} : \{1, \dots, n_1\} \hookrightarrow \{1, \dots, n_3\}$ is an injection witnessing $\biguplus \mathcal{M}_1 \preccurlyeq \biguplus \mathcal{M}_3$. \blacksquare
3. *Antisymmetry:* Let

$$\begin{aligned}\biguplus \mathcal{M}_1 &= \biguplus_{i=1}^{n_1} M_i \in \mathbb{M}_L^*, \quad \text{and} \\ \biguplus \mathcal{M}_2 &= \biguplus_{i=1}^{n_2} M_i \in \mathbb{M}_L^*\end{aligned}$$

be two configurations such that $\biguplus \mathcal{M}_1 \preccurlyeq \biguplus \mathcal{M}_2$ and $\biguplus \mathcal{M}_2 \preccurlyeq \biguplus \mathcal{M}_1$. We must show that $\biguplus \mathcal{M}_1 = \biguplus \mathcal{M}_2$. Since there exist injections sending the units of $\biguplus \mathcal{M}_1$ to the units of $\biguplus \mathcal{M}_2$ and vice versa, the two configurations have the same number of units—*i.e.*, we have $n_1 = n_2$. Then there exists a bijection from $\{1, \dots, n_1\}$ to $\{1, \dots, n_2\}$, and so the two configurations are the same up to permutation. \blacksquare

Having satisfied the three properties, we conclude that subconfigurationhood is a preorder. [*Back to the text.*]

⁷⁵In fact, in the proof it is shown that \preccurlyeq is a partial order on \mathbb{M}_L^* up to permutation of the molecules, which is in keeping with the spirit of \biguplus .

4.3 Lemma (Continuity of the Weak Metric)

Under the double separation property, the function

$$d_\Theta : \mathbb{F}_R^*(M_L^*) \times \mathbb{F}_R^*(M_L^*) \longrightarrow \Theta, \\ (\mathcal{F}, \mathcal{G}) \longmapsto (\text{cost}_\Theta \circ \mathbf{CS}_R)(\mathcal{F}, \mathcal{G})$$

is continuous, where the domain $\mathbb{F}_R^*(M_L^*) \times \mathbb{F}_R^*(M_L^*)$ has the product topology obtained from the weak metric topology on $\mathbb{F}_R^*(M_L^*)$ representing restructuring costs, and the codomain Θ has the strict-bound topology. [Proof.]

Proof. Two matters of set-up:

1. In terms of what we have, let $\mathcal{U} \subseteq \Theta$ be open in the strict bound topology. It takes the form

$$\mathcal{U} = \bigcup_{\alpha \in A} [\theta_\Theta, \theta_\alpha),$$

where A is some index set and $\theta_\alpha \in \Theta$ for all $\alpha \in A$. We have all of the θ_α s at our disposal.

2. In terms of what we need, we must show that the preimage of \mathcal{U} is open in the product topology; we will now characterize openness in that topology. Begin with $\mathbb{F}_R^*(M_L^*)$, which has the topology generated by the basis

$$\{B(F, \varepsilon) \mid F \in \mathbb{F}_R^*(M_L^*), \varepsilon \in \Theta \setminus [\theta_\Theta]\},$$

where $B_\Theta(F, \varepsilon) = \{G \in \mathbb{F}_R^*(M_L^*) \mid \varepsilon >_\Theta d_\Theta(F, G)\}$ and where $[\theta_\Theta]$ is the equivalence class of θ_Θ . We therefore generate the product topology on $\mathbb{F}_R^*(M_L^*) \times \mathbb{F}_R^*(M_L^*)$ with the basis

$$\{B_\Theta(F, \varepsilon) \times B_\Theta(G, \varepsilon) \mid F, G \in \mathbb{F}_R^*(M_L^*), \varepsilon \in \Theta \setminus [\theta_\Theta]\}.$$

We may use different ε s for the two sets in the product when necessary, but often we will use the same ε for both. So, we need to show that the preimage of \mathcal{U} is open in $\mathbb{F}_R^*(M_L^*) \times \mathbb{F}_R^*(M_L^*)$ —i.e., it is an arbitrary union of sets of this form.

This concludes our set-up.

We now proceed to the proof. Choose and fix any $\alpha \in A$; we need to show that the preimage of $[\mathbb{0}_\Theta, \theta_\alpha]$ is open in $\mathbb{F}_R^*(\mathbb{M}_L^*) \times \mathbb{F}_R^*(\mathbb{M}_L^*)$. The preimage in question is

$$d_\Theta^{-1}(U_i) = \{(F, G) \in \mathbb{F}_R^*(\mathbb{M}_L^*) \times \mathbb{F}_R^*(\mathbb{M}_L^*) \mid \theta_\alpha >_\Theta d_\Theta(F, G)\},$$

which we need to show can be written as a union of sets of the form $B_\Theta(F, \varepsilon) \times B_\Theta(G, \varepsilon)$. Fix any $(F, G) \in d_\Theta^{-1}(U_\alpha)$; then we know that $\theta_\alpha >_\Theta d_\Theta(F, G)$. To construct our neighborhood, we need to find some $\varepsilon \in \Theta$ such that $B_\Theta(F, \varepsilon) \times B_\Theta(G, \varepsilon) \subseteq d_\Theta^{-1}(U_\alpha)$.

Since $\theta_\alpha >_\Theta d_\Theta(F, G)$, the double separation property of Θ entails the existence of some $\varepsilon >_\Theta \mathbb{0}_\Theta$ such that

$$\theta_\alpha >_\Theta \varepsilon \oplus_\Theta d_\Theta(F, G) \oplus_\Theta \varepsilon.$$

Choose any $(F', G') \in B_\Theta(F, \varepsilon) \times B_\Theta(G, \varepsilon)$. Now, from the Triangle Inequality, we have

$$d_\Theta(F', G') \leq_\Theta d_\Theta(F', F) \oplus_\Theta d_\Theta(F, G) \oplus_\Theta d_\Theta(G, G').$$

By construction, we have $d_\Theta(F', F) <_\Theta \varepsilon$ and $d_\Theta(G, G') <_\Theta \varepsilon$; since \oplus_Θ is monotone and \leq_Θ is transitive, we infer

$$d_\Theta(F', G') <_\Theta \varepsilon \oplus_\Theta d_\Theta(F, G) \oplus_\Theta \varepsilon.$$

A final application of transitivity of \leq_Θ gives

$$d_\Theta(F', G') <_\Theta \theta_\alpha.$$

This shows that $B_\Theta(F, \varepsilon) \times B_\Theta(G, \varepsilon) \subseteq d_\Theta^{-1}(U_i)$.

Now, let $\varepsilon(F, G)$ be the ε that we constructed for each $(F, G) \in d_\Theta^{-1}(U_i)$. Then we have

$$d_\Theta^{-1}(U_i) = \bigcup_{(F,G) \in d_\Theta^{-1}(U_i)} B_\Theta(F, \varepsilon(F, G)) \times B_\Theta(G, \varepsilon(F, G)),$$

which—as a union of sets of the form $B_\Theta(F, \varepsilon) \times B_\Theta(G, \varepsilon)$ —is open in the product topology. [\[Back to the text.\]](#) ■

Before proving Lemma 4.5, we need a bit of set-up. First, let us generalize the notion of a sequence to that of a net. We do so because our problem lacks structure in many key ways; the two most problematic are the lack of completeness of \geq_Θ on Θ and the fact that d_Θ is not a full-blown metric (in particular, the lack of symmetry is a problem). So, let us give a general definition of a net.

A.1 Definition (Net)

Let (X, \leq) be a preorder. A net in X is a function $\varphi : \Lambda \rightarrow X$ from a directed set Λ to X . We write φ_λ for $\varphi(\lambda)$.

Given a net, we can define the oddly-tricky notion of a subnet.

A.2 Definition (Subnet)

Let $\varphi : \Lambda \rightarrow X$ be a net. A subnet of φ is a net $\psi : \Lambda' \rightarrow X$ such that there exists a function $\lambda : \Lambda' \rightarrow \Lambda$ such that, for every $\lambda' \in \Lambda'$, there exists some $\lambda_{\lambda'} \in \Lambda$ such that $\lambda(\lambda') \leq \lambda_{\lambda'}$ and $\varphi_{\lambda_{\lambda'}} = \psi_{\lambda'}$.

Finally, we introduce the idea of a net having converging subnets.

A.3 Definition (Convergent Subnets)

Let (X, \mathcal{T}) be a topological space. We say a set $K \subseteq X$ has convergent subnets just in case every net in K has a convergent subnet—i.e., if $\{X_\lambda\}_{\lambda \in \Lambda}$ is a net in K , then there exists some $x \in X$ and some function $\lambda : \Lambda' \rightarrow \Lambda$ such that, for every open set $O \in \mathcal{T}$ containing x , there exists some $\bar{\lambda} \in \Lambda'$ such that $X_{\lambda(\lambda')} \in O$ for all $\lambda' \geq_{\Lambda'} \bar{\lambda}$.

This is important thanks to the following result; see for example Willard (1970, Theorem 17.4, p. 118).

A.4 Theorem (Compactness is Equivalent to Convergent Subnets)

Let (X, \mathcal{T}) be a topological space. A set $K \subseteq X$ is compact iff it has convergent subnets.

Now we can prove Lemma 4.5.

4.5 Lemma (Emulation Problems Have Compact Domain)

Under the topologies from Lemma 4.3, the set

$$\mathcal{D}_S(\beta) = \{\mathcal{F} \in \mathbb{F}_R^*(\mathbb{M}_L^*) \mid \beta \geq_\Theta d_\Theta(\mathcal{F}, \mathcal{S})\} \subseteq \mathbb{F}_R^*(\mathbb{M}_L^*)$$

is compact for any $\mathcal{S} \in \mathbb{F}_R^*(\mathbb{M}_L^*)$ and $\beta \in \Theta$.

[*Proof.*]

Proof. Suppose Θ has the strict-bound topology, and fix any $\mathcal{S} \in \mathbb{F}_R^*(\mathbb{M}_L^*)$ and $\beta \in \Theta$. We will proceed in a few steps, including a bit of setup.

Step 1: Introducing the \mathcal{S} -map. We define

$$\begin{aligned} d_\Theta(\mathcal{S}, \cdot) : \mathbb{F}_R^*(\mathbb{M}_L^*) &\longrightarrow \Theta, \\ \mathcal{F} &\longmapsto d_\Theta(\mathcal{S}, \mathcal{F}). \end{aligned}$$

This sends a structured force \mathcal{F} to the distance between \mathcal{S} and \mathcal{F} . We know d_Θ is continuous (Lemma 4.3) when Θ is equipped with the strict-bound topology, so this map is continuous. Put a pin in this for now.

Step 2: Introducing the target set. We also define

$$[\mathbb{0}_\Theta, \beta] := \{\theta \in \Theta \mid \mathbb{0}_\Theta \leq_\Theta \theta \leq_\Theta \beta\} \subseteq \Theta.$$

We immediately observe that $[\mathbb{0}_\Theta, \beta]$ is compact under the strict-bound topology. To see why, let $\{O_\alpha\}_{\alpha \in A}$ be an open cover of $[\mathbb{0}_\Theta, \beta]$. Since $\beta \in [\mathbb{0}_\Theta, \beta]$, there exists some $\alpha_0 \in A$ such that $\beta \in O_{\alpha_0}$ —i.e., some θ_{α_0} such that $\theta_{\alpha_0} >_\Theta \beta$. But then, since \geq_Θ is transitive, we can cover all of $[\mathbb{0}_\Theta, \beta]$ with O_{α_0} . This shows that $[\mathbb{0}_\Theta, \beta]$ is compact.

Step 3: The nets. Now consider any net $\{\mathcal{F}_\lambda\}_{\lambda \in \Lambda}$ in $\mathcal{D}_S(\beta)$ where Λ is a directed set—that is, for any two $\lambda_1, \lambda_2 \in \Lambda$, there exists some $\lambda_3 \in \Lambda$ such that $\lambda_1 \leq_\Lambda \lambda_3$ and $\lambda_2 \leq_\Lambda \lambda_3$. It is attached to the net

$$\{d_\Theta(\mathcal{S}, \mathcal{F}_\lambda)\}_{\lambda \in \Lambda} \subseteq [\mathbb{0}_\Theta, \beta]$$

by way of the function $d_\Theta(\mathcal{S}, \cdot)$. Since $[\mathbb{0}_\Theta, \beta]$ is compact, this net has a convergent subnet $\{d_\Theta(\mathcal{S}, \mathcal{F}_{\lambda'})\}_{\lambda' \in \Lambda'}$ —this is Theorem A.4. This means that there exists some $\theta^* \in [\mathbb{0}_\Theta, \beta]$ and some function $\lambda : \Lambda' \rightarrow \Lambda$ such that, for every open set $O \in \mathcal{T}$ containing θ^* , there exists some $\bar{\lambda} \in \Lambda'$ such that $d_\Theta(\mathcal{S}, \mathcal{F}_{\lambda(\bar{\lambda})}) \in O$ for all $\lambda' \geq_{\Lambda'} \bar{\lambda}$.

Step 4: Pulling the subnet back. Since $d_\Theta(\mathcal{S}, \cdot)$ is continuous, the fact that the subnet $\{d_\Theta(\mathcal{S}, \mathcal{F}_{\lambda'})\}_{\lambda' \in \Lambda'}$ converges to θ^* means that the subnet $\{\mathcal{F}_{\lambda'}\}_{\lambda' \in \Lambda'}$ converges to \mathcal{F}^* , where \mathcal{F}^* is the preimage of θ^* under $d_\Theta(\mathcal{S}, \cdot)$. Moreover, by construction we have

$$d_\Theta(\mathcal{S}, \mathcal{F}^*) = \theta^* \leq_\Theta \beta,$$

so that we have $\mathcal{F}^* \in \mathcal{D}_S(\beta)$.

Step 5: Conclusion. We have shown that every net in $\mathcal{D}_S(\beta)$ has a convergent subnet; by Theorem A.4, this means that $\mathcal{D}_S(\beta)$ is compact. [[Back to the text.](#)] ■

4.6 Proposition (Emulation Problems Have Solutions)

Under the topologies from Lemma 4.3, the emulation problem has a solution. [[Proof.](#)]

Proof. Choose any status quo force $\mathcal{S} \in \mathbb{F}_R^*(\mathbb{M}_L^*)$, any target force $\mathcal{T} \in \mathbb{F}_R^*(\mathbb{M}_L^*)$, and any budget $\beta \in \Theta$. Define the map

$$\begin{aligned} d_\Theta(\cdot, \mathcal{T}) : \mathbb{F}_R^*(\mathbb{M}_L^*) &\longrightarrow \Theta, \\ \mathcal{F} &\longmapsto d_\Theta(\mathcal{F}, \mathcal{T}). \end{aligned}$$

This map is continuous by Lemma 4.3. The set $\mathcal{D}_S(\beta)$ is compact by Lemma 4.5, so its image under this map,

$$d_\Theta(\mathcal{D}_S(\beta), \mathcal{T}) = \{d_\Theta(\mathcal{F}, \mathcal{T}) \mid \mathcal{F} \in \mathcal{D}_S(\beta)\} \subseteq \Theta,$$

is also compact. Moreover, since Θ has all joins, there exists some

$$\bigvee d_\Theta(\mathcal{D}_S(\beta), \mathcal{T})$$

such that

$$\begin{aligned} \theta &\geq_\Theta \bigvee d_\Theta(\mathcal{D}_S(\beta), \mathcal{T}) \text{ for all } \theta \in d_\Theta(\mathcal{D}_S(\beta), \mathcal{T}), \\ \bigvee d_\Theta(\mathcal{D}_S(\beta), \mathcal{T}) &\geq_\Theta \theta \text{ for all other such } \theta. \end{aligned}$$

Since $\bigvee d_\Theta(\mathcal{D}_S(\beta), \mathcal{T})$ is compact, we know that

$$\bigvee d_\Theta(\mathcal{D}_S(\beta), \mathcal{T}) \in d_\Theta(\mathcal{D}_S(\beta), \mathcal{T}),$$

so there exists some $\mathcal{F}^* \in \mathcal{D}_S(\beta)$ such that

$$d_\Theta(\mathcal{F}^*, \mathcal{T}) = \bigvee d_\Theta(\mathcal{D}_S(\beta), \mathcal{T}) \in (d_\Theta(\mathcal{D}_S(\beta), \mathcal{T})).$$

By construction, this \mathcal{F}^* is a solution to the emulation problem. [[Back to the text.](#)] ■

4.8 Lemma (Continuity of \oplus_Θ)

The function $\oplus_\Theta : \Theta \times \Theta \rightarrow \Theta$ is continuous in the strict-bound topology if, and only if, for any pair of labels $\theta_1, \theta_2 \in \Theta$, there exists some $\theta_1 \circ \theta_2 \in \Theta$ such that

$$\theta_3 \oplus_\Theta \theta_2 \leq_\Theta \theta_1 \text{ if and only if } \theta_3 \leq_\Theta \theta_1 \circ \theta_2,$$

called the strict left adjoint of \oplus_Θ .

[*Proof.*]

Proof. Suppose that, for any pair of labels $\theta_1, \theta_2 \in \Theta$, there exists some $\theta_1 \circ \theta_2 \in \Theta$ such that

$$\theta_3 \oplus_\Theta \theta_2 \leq_\Theta \theta_1 \text{ if and only if } \theta_3 \leq_\Theta \theta_1 \circ \theta_2.$$

We will first show that this condition ensures that \oplus_Θ is continuous. Begin from an open set on the right-hand side,

$$O = \bigcup_{\alpha \in A} [\theta_\Theta, \theta_\alpha),$$

which is a union of upper sets in the strict-bound topology in Θ . Its preimage under the join operation, $(\oplus_\Theta)^{-1}(O)$, is the set of all pairs of cost labels whose join is less than some θ_α for some $\alpha \in A$:

$$\{(\theta_1, \theta_2) \in \Theta \times \Theta \mid \theta_1 \oplus_\Theta \theta_2 \leq_\Theta \theta_\alpha \text{ for some } \alpha \in A\} \subseteq \Theta \times \Theta.$$

We need to show that this set can be written as a union of open sets in $\Theta \times \Theta$, which means we need to be explicit about the product topology under our strict-bound topology. The product topology is generated by the sets

$$\{(\theta'_1, \theta'_2) \in \Theta \times \Theta \mid \theta'_1 \leq_\Theta \theta_1 \text{ and } \theta'_2 \leq_\Theta \theta_2\},$$

where $\theta_1, \theta_2 \in \Theta$.

Now, choose and fix any $\alpha \in A$, and consider the set

$$\bigcup_{\theta_1 \leq_\Theta \theta_\alpha} \bigcup_{\theta_2 \leq_\Theta \theta_\alpha \circ \theta_1} \{(\theta_1, \theta_2)\},$$

which is a union of open sets in $\Theta \times \Theta$, the first upper bound being θ_α and the second upper bound being $\theta_\alpha \circ \theta_1$. I claim this is the preimage for $[\theta_\Theta, \theta_\alpha)$ under the join operation. We must prove two inclusions.

1. First, suppose that $(\theta_1, \theta_2) \in (\oplus_\Theta)^{-1}([\theta_\Theta, \theta_\alpha])$, meaning that $\theta_1 \oplus_\Theta \theta_2 \lessdot_\Theta \theta_\alpha$. From the definition of \lessdot , we know that $\theta_2 \lessdot_\Theta \theta_\alpha \circ \theta_1$, which gets us halfway there. However, since \oplus_Θ is monotone, it is immediate that $\theta_1 \lessdot_\Theta \theta_\alpha$; otherwise, we could not have $\theta_1 \oplus_\Theta \theta_2 \lessdot_\Theta \theta_\alpha$ to begin with. We conclude that $(\theta_1, \theta_2) \in \bigcup_{\theta_1 \lessdot_\Theta \theta_\alpha} \bigcup_{\theta_2 \lessdot_\Theta \theta_\alpha \circ \theta_1} \{(\theta_1, \theta_2)\}$. \square
2. Second, suppose that $(\theta_1, \theta_2) \in \bigcup_{\theta_1 \lessdot_\Theta \theta_\alpha} \bigcup_{\theta_2 \lessdot_\Theta \theta_\alpha \circ \theta_1} \{(\theta_1, \theta_2)\}$. Since $\theta_2 \lessdot_\Theta \theta_\alpha \circ \theta_1$, the definition of \circ tells us that $\theta_1 \oplus_\Theta \theta_2 \lessdot_\Theta \theta_\alpha$. Thus, $(\theta_1, \theta_2) \in (\oplus_\Theta)^{-1}([\theta_\Theta, \theta_\alpha])$. \square

Since the preimage of $[\theta_\Theta, \theta_\alpha]$ is open in $\Theta \times \Theta$ for all $\alpha \in A$, we have shown that \oplus_Θ is continuous.

For the other direction, we observe that the requirement for continuity is that the preimage of any open set in Θ is open in $\Theta \times \Theta$. This means the preimage must be a union of sets of the form

$$\bigcup_{\theta_1 \lessdot_\Theta \bar{\theta}_1} \bigcup_{\theta_2 \lessdot_\Theta \bar{\theta}_2} \{(\theta_1, \theta_2)\},$$

for some $\bar{\theta}_1, \bar{\theta}_2 \in \Theta$. It is without loss of generality to observe that full coverage of this set requires $\bar{\theta}_1 = \theta_\alpha$, or else the pair $(\theta_\alpha, \theta_\Theta)$ would not be covered even though $\theta_\alpha \oplus_\Theta \theta_\Theta = \theta_\alpha \lessdot_\Theta \theta_\alpha$. But what does this mean for $\bar{\theta}_2$? The two inclusion arguments above show that the construction works only when $\bar{\theta}_2$ satisfies the condition

$$\theta_2 \oplus_\Theta \theta_1 \lessdot_\Theta \theta_\alpha \quad \text{if and only if} \quad \theta_2 \lessdot_\Theta \bar{\theta}_2,$$

where the biconditional is because we used the definition of \circ in both directions, one for each inclusion argument. We conclude that $\bar{\theta}_2 = \theta_\alpha \circ \theta_1$ must exist for the preimage to be open, which completes the proof. [[Back to the text](#).] \blacksquare

4.14 Proposition (Scalability of Regular Doctrines)

If \geq is a regular doctrine on $\mathbb{F}_R^(\mathbb{M}_L^*)$, then it is scalable.*

[[Proof](#).]

Proof. Let us first show the standard result that \sim is an equivalence relation.

1. *Reflexiveness.* Let $\mathcal{F} \in \mathbb{F}_R^*(\mathbb{M}_L^*)$. We must show that $\mathcal{F} \sim \mathcal{F}$. Since \geq is complete, we have $\mathcal{F} \geq \mathcal{F}$ or $\mathcal{F} \geq \mathcal{F}$. This reduces to $\mathcal{F} \sim \mathcal{F}$, as desired. \square

2. *Symmetry.* Let $\mathcal{F}_1, \mathcal{F}_2 \in \mathbb{F}_R^*(\mathbb{M}_L^*)$ be such that $\mathcal{F}_1 \sim \mathcal{F}_2$. We must show that $\mathcal{F}_2 \sim \mathcal{F}_1$. From the definition of \sim , we have $\mathcal{F}_1 \geq \mathcal{F}_2$ and $\mathcal{F}_2 \geq \mathcal{F}_1$. Logical conjunction is symmetric, so we have $\mathcal{F}_2 \geq \mathcal{F}_1$ and $\mathcal{F}_1 \geq \mathcal{F}_2$. Again from the definition of \sim , we conclude that $\mathcal{F}_2 \sim \mathcal{F}_1$, as desired. \square
3. *Transitivity.* Let $\mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_3 \in \mathbb{F}_R^*(\mathbb{M}_L^*)$ be such that $\mathcal{F}_1 \sim \mathcal{F}_2$ and $\mathcal{F}_2 \sim \mathcal{F}_3$. We must show that $\mathcal{F}_1 \sim \mathcal{F}_3$. From the first relation, we have $\mathcal{F}_1 \geq \mathcal{F}_2$ and $\mathcal{F}_2 \geq \mathcal{F}_1$. From the second, we have $\mathcal{F}_2 \geq \mathcal{F}_3$ and $\mathcal{F}_3 \geq \mathcal{F}_2$. By transitivity of \geq , we have $\mathcal{F}_1 \geq \mathcal{F}_3$ and $\mathcal{F}_3 \geq \mathcal{F}_1$, and the definition of \sim gives us $\mathcal{F}_1 \sim \mathcal{F}_3$, as desired. \square

We conclude that \sim is an equivalence relation, and we use the standard terminology for its equivalence classes:

$$[\mathcal{F}]_\sim := \{\mathcal{G} \in \mathbb{F}_R^*(\mathbb{M}_L^*) \mid \mathcal{G} \sim \mathcal{F}\}.$$

Likewise, we introduce the quotient set

$$\mathbb{F}_R^*(\mathbb{M}_L^*) / \sim := \{[\mathcal{F}]_\sim \mid \mathcal{F} \in \mathbb{F}_R^*(\mathbb{M}_L^*)\}.$$

We now define a new relation \geq^* on $\mathbb{F}_R^*(\mathbb{M}_L^*) / \sim$ by the rule

$$[\mathcal{F}_1]_\sim \geq^* [\mathcal{F}_2]_\sim \quad \text{if and only if} \quad \mathcal{F}_1 \geq \mathcal{F}_2.$$

Clearly, \geq^* inherits completeness and transitivity from \geq . It also happens to be antisymmetric, as we now show. Let $[\mathcal{F}_1]_\sim, [\mathcal{F}_2]_\sim \in \mathbb{F}_R^*(\mathbb{M}_L^*) / \sim$ be such that $[\mathcal{F}_1]_\sim \geq^* [\mathcal{F}_2]_\sim$ and $[\mathcal{F}_2]_\sim \geq^* [\mathcal{F}_1]_\sim$. Then we have $\mathcal{F}_1 \geq \mathcal{F}_2$ and $\mathcal{F}_2 \geq \mathcal{F}_1$, implying $\mathcal{F}_1 \sim \mathcal{F}_2$. Since \sim is transitive, anything that is \sim to \mathcal{F}_1 is also \sim to \mathcal{F}_2 , and vice versa. Thus, the two equivalence classes are the same set, implying that $[\mathcal{F}_1]_\sim = [\mathcal{F}_2]_\sim$. This means \geq^* is antisymmetric.

Next, we define a function

$$\varphi : \mathbb{F}_R^*(\mathbb{M}_L^*) / \sim \longrightarrow \mathbb{Q}.$$

Pay special attention to the codomain: we are using the rationals, not the reals. Since the former is a subset of the latter, it would suffice to identify a function that works on the rationals and then extend the codomain. Now, since $\mathbb{F}_R^*(\mathbb{M}_L^*)$ is countable, so too is its quotient set $\mathbb{F}_R^*(\mathbb{M}_L^*) / \sim$. Thus, we can write it by

$$\mathbb{F}_R^*(\mathbb{M}_L^*) / \sim = \{[\mathcal{F}_1]_\sim, [\mathcal{F}_2]_\sim, \dots\}.$$

The rationals are also countable—this is a rather famous result due to Cantor that your humble will not bother to address here—so we can write them as

$$\mathbb{Q} = \{q_1, q_2, \dots\}.$$

We construct φ inductively like so:

1. set $\varphi([\mathcal{F}_1]_\sim) = q_1$;
2. for each $i = 2, 3, \dots$, we set $\varphi([\mathcal{F}_i]_\sim)$ to be the first q_i (with respect to its index) with the same ordering with respect to $\varphi([\mathcal{F}_1]_\sim), \dots, \varphi([\mathcal{F}_{i-1}]_\sim)$ as $[\mathcal{F}_i]_\sim$ does with respect to $[\mathcal{F}_1]_\sim, \dots, [\mathcal{F}_{i-1}]_\sim$. In other words:
 - (a) in case $\mathcal{F}_i \geq \mathcal{F}_j$ for all $j < i$, we set $\varphi([\mathcal{F}_i]_\sim)$ to be the first rational number (with respect to the index) strictly greater than all $\varphi([\mathcal{F}_j]_\sim)$ for $j < i$;
 - (b) similar reasoning goes in case $\mathcal{F}_i < \mathcal{F}_j$ for all $j < i$;
 - (c) in case \mathcal{F}_i is indifferent to any \mathcal{F}_j for $j < i$, we set $\varphi([\mathcal{F}_i]_\sim)$ to have the same value as $\varphi([\mathcal{F}_j]_\sim)$ for the first $j < i$ that is indifferent to \mathcal{F}_i ;
 - (d) in case \mathcal{F}_i is strictly better than some \mathcal{F}_j for $j < i$ and strictly worse than others (and indifferent to none), we identify where it fits in the order and set $\varphi([\mathcal{F}_i]_\sim)$ to be the average of the two closest rational numbers.

This is well-defined because of the completeness and transitivity of \geq and the order-density of the rationals.

Given the terms of the construction, we have the desired biconditional relationship between \geq and \geq , and we have constructed a function that represents the doctrine. [\[Back to the text.\]](#) ■

References

Aliprantis, Charalambos D. and Kim C. Border. 2006. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Third ed. Berlin: Springer.

Allison, Graham. 2017. *Destined for War: Can America and China Escape Thucydides's Trap?* Boston, MA: Houghton Mifflin Harcourt.

Awodey, Steve. 2010. *Category Theory*. Number 52 in "Oxford Logic Guides" second ed. Oxford, UK: Oxford University Press.

Baliga, Sandeep and Tomas Sjöström. 2004. "Arms Races and Negotiations." *Review of Economic Studies* 71(2):351–369.

Baliga, Sandeep and Tomas Sjöström. 2008. "Strategic Ambiguity and Arms Proliferation." *Journal of Political Economy* 116(6):1023–1057.

Bas, Muhammet A. and Andrew J. Coe. 2016. "A Dynamic Theory of Nuclear Proliferation and Preventive War." *International Organization* 70(4):655–685.

Bas, Muhammet A. and Andrew J. Coe. 2018. "Give Peace a (Second) Chance: A Theory of Nonproliferation Deals." *International Studies Quarterly* 62(3):606–617.

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

Bollobás, Béla. 1998. *Modern Graph Theory*. Number 184 in "Graduate Texts in Mathematics" New York, NY: Springer.

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.

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

Bury, Patrick. 2021. "Conceptualising the Quiet Revolution: The Post-Fordist Revolution in Western Military Logistics." *European Security* 30(1):112–136.

Camerer, Colin F. 2006. "Wanting, Liking, and Learning: Neuroscience and Paternalism." *University of Chicago Law Review* 73(1).

Carroll, Robert J. and Brenton Kenkel. 2019. "Prediction, Proxies, and Power." *American Journal of Political Science* 63(3):577–593.

Chandler, David G. 1966. *The Campaigns of Napoleon*. New York: Macmillan.

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.

Congressional Budget Office. 2021. "The U.S. Military's Force Structure: A Primer." Congressional Budget Office. Technical report.

Cormen, Thomas H., Charles E. Leiserson, Ronald L. Rivest and Clifford Stein. 1990. *Introduction to Algorithms*. Fourth (2022) ed. Cambridge, MA: The MIT Press.

Dahl, Robert A. 1957. "The Concept of Power." *Behavioral Science* 2(3):201–215.

Debreu, Gerard. 1964. "Continuity Properties of Paretian Utility." *International Economic Review* 5(3):285–293.

Debs, Alexandre and Nuno P. Monteiro. 2014. "Known Unknowns: Power Shifts, Uncertainty, and War." *International Organization* 68(1):1–31.

Diestel, Reinhard. 1998. *Graph Theory*. Number 173 in "Graduate Texts in Mathematics" fifth ebook (2016) ed. New York: Springer.

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.

Dowding, Keith, ed. 2011. *Encyclopedia of Power*. Thousand Oaks, CA: Sage.

Euclid. 300 B.C. *The Thirteen Books of Euclid's Elements*. Second (1926) heath-translated ed. New York: Cambridge University Press.

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

Fearon, James D. 2018. "Cooperation, Conflict, and the Costs of Anarchy." *International Organization* 72(3):523–559.

Fong, Brendan and David I. Spivak. 2019. *An Invitation to Applied Category Theory: Seven Sketches in Compositionality*. Cambridge, UK: Cambridge University Press.

Fréchet, Maurice. 1906. "Sur Quelques Points Du Calcul Fonctionnel." *Rendiconti del Circolo Matematico di Palermo* 22(1):1–72.

Fung, Allen. 1996. "Testing the Self-Strengthening: The Chinese Army in the Sino-Japanese War of 1894–1895." *Modern Asian Studies* 30(4):1007–1031.

Garfinkel, Michelle R. 1990. "Arming as a Strategic Investment in a Cooperative Equilibrium." *American Economic Review* 80(1):50–68.

Goldman, Emily O. and Leslie C. Eliason, eds. 2003. *The Diffusion of Military Technology and Ideas*. Stanford, CA: Stanford University Press.

Grossman, Herschell I. 1991. "A General Equilibrium Model of Insurrections." *American Economic Review* 81(4):912–921.

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.

Howard, Michael. 1976. *War in European History*. Updated ed. Oxford, UK: Oxford University Press.

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. 1969. *The Outbreak of the Peloponnesian War*. Ithaca, NY: Cornell University Press.

Kagan, Donald and Gregory F. Viggiano, eds. 2013. *Men of Bronze: Hoplite Warfare in Ancient Greece*. Princeton, NJ: Princeton University Press.

Kelley, John L. 1975. *General Topology*. Number 27 in "Graduate Texts in Mathematics" Berlin: Springer.

Kissinger, Henry. 1973. "Moral Purposes and Policy Choices." Speech presented to the Center for the Study of Democratic Institutions on the occasion of *Pacem in Terris III*, Washington, DC.

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.

Lawvere, F. William. 1973. "Metric Spaces, Generalized Logic, and Closed Categories." *Rendiconti del Seminario Matematico e Fisico di Milano* 43(1):135–166.

Mac Lane, Saunders. 1971. *Categories for the Working Mathematician*. Number 5 in "Graduate Texts in Mathematics" 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.

Mallett, Michael. 1974. *Mercenaries and Their Masters: Warfare in Renaissance Italy*. Republished (2009) ed. Barnsley, UK: Pen & Sword.

Martin, Bernd. 1995. *Japan and Germany in the Modern World*. Providence, RI: Berghahn Books.

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.

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.

Munkres, James R. 2000. *Topology*. Second ed. Upper Saddle River, NJ: Prentice Hall.

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.

Ok, Efe A. 2007. *Real Analysis with Economic Applications*. Princeton, NJ: Princeton University Press.

Ortega y Gasset, José. 1930. *The Revolt of the Masses*. Kerrigan (1985) ed. Notre Dame, IN: University of Notre Dame Press.

Polmar, Norman. 2006. *Aircraft Carriers: A History of Carrier Aviation and Its Influence on World Events*. Dulles, VA: Potomac Books.

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

Resende-Santos, João. 2007. *Neorealism, States, and the Modern Mass Army*. Cambridge: Cambridge University Press.

Riehl, Emily. 2016. *Category Theory in Context*. Mineola, NY: Dover.

Riesen, Kaspar. 2015. *Structural Pattern Recognition with Graph Edit Distance: Approximation Algorithms and Applications*. Advances in Computer Vision and Pattern Recognition Switzerland: Springer International Publishing.

Santayana, George. 1896. *The Sense of Beauty: Being the Outlines of Aesthetic Theory*. Critical (1986) ed. Cambridge, MA: The MIT 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.

Serratosa, Francesc. 2021. "Redefining the Graph Edit Distance." *SN Computer Science* 2(6):438.

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.

Strange, Susan. 1988. *States and Markets*. Reprinted (1993) ed. London: Pinter.

Thucydides. 432 B.C.E. *The History of the Peloponnesian War*. Project gutenberg (2003) ed. Salt Lake City, UT: Project Gutenberg.

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.

West, Douglas B. 1996. *Introduction to Graph Theory*. Second (2001) ed. Upper Saddle River, NJ: Prentice Hall.

Willard, Stephen. 1970. *General Topology*. Unabridged dover (2004) ed. Mineola, NY: Dover.