

Working Papers

A Simple Model of Government Decision-Making and Plurilateral Cooperation on Critical Materials

Spiros Bougheas, Bernard Hoekman, Douglas Nelson

Imprint:

CESifo Working Papers

ISSN 2364-1428 (digital)

Publisher and distributor: Munich Society for the Promotion
of Economic Research - CESifo GmbH

Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740

Email office@cesifo.de
<https://www.cesifo.org>

Editor: Clemens Fuest

An electronic version of the paper may be downloaded free of charge

- from the CESifo website: www.ifo.de/en/cesifo/publications/cesifo-working-papers
- from the SSRN website: www.ssrn.com/index.cfm/en/cesifo/
- from the RePEc website: <https://ideas.repec.org/s/ces/ceswps.html>

A Simple Model of Government Decision-Making and Plurilateral Cooperation on Critical Materials*

Spiros Bougheas (University of Nottingham)

Bernard Hoekman (EUI)

Douglas Nelson (Tulane University)

July 8, 2026

Abstract: We develop an integrated framework to characterize multi-objective policymaking and apply it to the case of critical materials. The model considers that national policy decisions will be influenced by domestic agencies with different priorities and goals and be informed by actions taken by other countries pertaining to critical materials. This creates a potential basis for a group of countries to form a critical materials club to reduce the probability of supply shocks and/or their consequences. We show that in an environment with multiple policy-active governments, criticality assessments embedded in a multi-objective government with domain-specific bureaucracies and multiple instruments can rationalize divergent policy choices under similar risk environments, provide conditions under which international coordination is more or less likely and apply the model to consider the design and political economy of plurilateral arrangements for critical materials.

JEL codes: E61, D71, F15, F53, F68, L52, Q38

Keywords: critical materials, multi-objective policymaking, plurilateral agreements, clubs

* We are very grateful to Henrik Horn, Petros Mavroidis and participants in the June 2026 EUI workshop on “Plurilateral trade agreements and policy spillovers” for valuable suggestions. Support from the European Research Council, grant no. 101142666-PIANO-ERC-2023-ADG, is gratefully acknowledged.

Introduction

Copper, iron, coal and petroleum have powered economies and militaries for centuries, but the modern world also runs on cobalt, lithium, and rare earth elements (REEs). The latter commodities are essential inputs to microchips, batteries and strong magnets which, in turn, are essential inputs to flat screens, solar panels, windmills, automobiles, aircraft, missiles and drones. From this list, it is easy to see that these goods are of fundamental concern to governments with respect to functioning of the economy in general, environmental policy, and national security. Not surprisingly, the world's major economies are concerned with access to these goods and have policies targeted on that access. As part of such policies, most of those economies maintain lists of what are generally called "critical materials".¹

REEs are components in a large number of economically important products, so reduced access could have significant macroeconomic effects (Leal Filho et al. 2023; Taboga 2024; Neary and Highley 1984). On the seller side, the markets for REEs are generally characterized by small volumes/values traded and concentration in both mining and refining. As a result, these are not generally competitive markets. On the buyer side, REEs are relatively small parts of the economic cost of the products in which they are incorporated. Those uses are embedded in globally complex supply chains, usually relatively early (upstream) in the chain, but the goods cannot be produced without them. This means that the potential for choke points and weaponization in the supply chains is substantial. For example, REEs are needed to produce strong magnets, batteries, and screens, all of which are, in turn, essential to many critical final goods (e.g. EVs, wind turbines, military equipment). Thus, standard intuitions and analyses that rest on competitive markets are problematic. Moreover, the fact that critical materials are embedded in environmental and security policies means that the relevant political economies are very different from those characterizing trade in, say, washing machines or wheat.

Given the geopolitical significance of critical materials-related policies it is surprising how little attention the topic has received by mainstream economists. The main framework for analysis used by policymakers is the Risk-Matrix Approach (Bao, Li, and Wu 2022). This is inadequate

¹ The International Energy Agency (IEA) maintains a website tracking national critical minerals policies. Twenty-three countries maintain critical minerals lists that are the focus of policies. See <https://www.iea.org/data-and-statistics/data-tools/critical-minerals-policy-tracker>. While we use critical minerals as a focal point for concreteness, critical materials extend beyond minerals, e.g., active pharmaceutical ingredients and medical supplies.

for addressing the multi-level complexity of the many, potentially conflicting goals at both the national and the international levels. In this paper we embed the risk-matrix approach in a model of multi-objective policymaking and apply it to the case of critical materials. Specifically, we consider a peak decision-maker that takes policy advice from a number of policy domain-specific bureaus. Each of these bureaus is charged with evaluating possible policies, with respect to critical materials, from the perspective of their specific domain.² Thus, the macroeconomic bureau (the economics ministry) focuses on critical materials policy from the perspective of macroeconomic performance, the national security bureau (the defense ministry) focuses on critical materials policy from the perspective of national security, and so on.

The policy decisions that emerge will be informed by actions taken by other countries pertaining to critical materials. Conversely, decisions (policies) of these countries will take into account those of other nations. This creates a potential basis for a group of countries to engage in self-enforcing plurilateral cooperation, i.e., form a critical materials club. Such a club could seek to reduce the probability of supply shocks and/or their consequences through e.g., coordination of investment in new mining projects, developing mid-stream processing capacity, joint stockpiling, etc.³

The paper proceeds as follows. In Section I we briefly describe the risk-matrix approach used by governments to prepare their lists of critical materials. After briefly describing this methodology and its use in the identification of critical materials, we argue it is a far from complete analysis of a government's policy choice problem over these commodities. In Section II we develop our integrated framework representing decision processes over critical goods. In the model the peak decision-maker responds to changes in the domestic and geopolitical environment, while the

² We call these domain-specific entities “bureaucracies”, but it might be more accurate to think of them as national policy communities anchored on a bureaucracy and incorporating relevant legislative committees. These may, themselves, be linked to an epistemic community with considerably broader membership. Haas (2021, 699) defines epistemic communities as “networks ... of knowledge-based experts with an authoritative claim to policy relevant knowledge within their domain”. In any event, we assume that there is a single entity, the bureau, that speaks authoritatively for this community on issues of the relevant domain.

³ We discuss several examples of such clubs later in this paper. The design of plurilateral cooperation in this area has become a subject of interest to many governments. For example, in 2026, USTR requested comments on the design of a plurilateral agreement on trade in critical minerals and associated measures to strengthen the resilience of critical mineral supply chains (<https://comments.ustr.gov/s/docket?docketNumber=USTR-2026-0034>). More generally, see Hoekman and Sabel (2019, 2021) for a discussion of open plurilateral agreements as a vehicle for issue- or domain-specific cooperation among a group of countries. Hoekman, Mavroidis and Nelson (2023) argue that plurilateral agreements are potential instruments to enhance economic resilience and security.

domain-specific bureaus respond to changes in the technocratic and epistemic environments within their domains.⁴ We begin by analyzing the case of a government with a single policy objective, two critical commodities (lithium and REEs), and a single policy instrument. We then extend the framework by introducing a second policy instrument and considering the case of two policy objectives. The generalization from this case to multiple critical goods and multiple policy objectives is straightforward. We conclude Section II by sketching an analysis of an environment with multiple policy-active governments, which permits us to show how criticality assessments, when embedded in a multi-objective government with domain-specific bureaucracies and multiple instruments, can rationalize divergent policy choices under similar risk environments, and provide conditions under which international coordination is more or less likely. To develop a more complete analysis of international cooperation, in Section III we show how the model can be extended to analyze the design and political economy of plurilateral arrangements for critical materials. We conclude in Section IV with an application of the conceptual framework to recent bilateral and plurilateral initiatives on critical materials.

I. Critical Goods: The Risk Matrix Approach

Loosely speaking, “critical goods” are goods without which the economy (or key sectors), civil society, and/or the state will fall well short of attaining core goals. Historically, critical goods often have been the products of mines—from copper and tin to iron and REEs (Raymond 1984). Access to these commodities, for most communities, has been to one degree or another, for one reason or another, risky. We can be more precise about risk and criticality by adopting the risk matrix framework that has often been used by governments and firms in identifying risks associated with specific goods (Cox 2009; European Commission 2014; National Research Council 2008; Bao, Li, and Wu 2022).

Although details vary widely across fields (finance, insurance, engineering, environmental science, etc.) the basic approach of risk analysis identifies two key elements: the probability of a bad outcome; and the consequence of a bad outcome. For the case of critical goods, a bad outcome is generally a significant reduction in access to the good, while the consequence is a negative effect on some core goal resulting from that lack of access. For the great majority of

⁴ We abstract from any strategic interactions between bureaus.

risk analysis, the relevant consequence has been an effect on economic performance but, following recent work, we will also consider national security and environmental degradation.⁵

The probability of disruption of access to critical materials has been an essential component of the definition of “criticality” by many governments. Thus, suppose we consider a set of potentially critical materials (I), denote the elements of the set $i \in I$, and the relevant bad outcome for i by x_i . An example of i might be REEs and the bad outcome might be China restricting exports.⁶ We can denote the probability that x_i occurs by $q(x_i)$. Similarly, we can denote the consequence of the same bad outcome as $C(x_i)$. The literature on risk matrix analysis treats risk, which we denote $\rho(x_i)$, as the product

$$\rho(x_i) = q(x_i) \times C(x_i).$$


The risk matrix is a representation of this logic. Early applications of this method simply involve identifying high, medium and low levels of both q and C and associate criticality with high levels of each. One of the first such applications, by the US National Research Council (2008, pg. 165), identified 13 minerals (see Figure 1). We will focus on the case where q and C are continuous variables. In this case, critical levels of q and C are chosen and minerals whose values of these variables are above the critical values are declared to be critical minerals.⁷

⁵ Much research on risk relates very micro-level analysis (e.g. firm or portfolio level risk). In this paper, the focus is on macroeconomic and macropolitical risk.

⁶ Of course, there are many potential sources of risk for a critical good (e.g. rare earths): demand might increase dramatically due to technological change; sources of supply might be exhausted. In addition, given how thin the markets for many critical goods are, local political/economic disturbances could result in supply restriction (e.g. civil disturbance, local war). Finally, given the small number of suppliers in some markets, monopolization is a real possibility. For our purposes, we will treat risk as a singleton—i.e. either there really is just one source of risk (China restricts exports) or there is a well-specified risk aggregate (Bjørnsen and Aven 2019; Bao, Li, and Wu 2022).

⁷ The European Commission (2014, p. 24) presents just such a graph.


Figure 1: Commodity Risk in a Simple Risk Matrix



Glöser et al. (2015, p. 42) suggest that it is natural to compare criticality across minerals using iso- ρ contours. That is, we can project all combinations of q and C that yield the same level of ρ into q - C space (see Figure 2). Since $\bar{\rho} = q \times C$ defines a rectangular hyperbola in q - C space for all value of $\rho = \bar{\rho}$.⁸ In this case, rather than choosing critical values for q and C separately, criticality is defined relative to a threshold critical level of ρ . Operationalization of the risk matrix involves estimation of $q(x_i)$ and $C(x_i)$ for all relevant i . Variables commonly used to measure q include country concentration; country risk; depletion time; by-product dependency; demand growth; and country concentration in mining corporations (Achzet and Helbig 2013). Similarly, variables commonly used to measure C include the value of products affected; value of utilized material; strategic importance; and substitutability (Helbig et al. 2016). As a practical matter, a selection of measures for each must be made and then aggregated into single q and C measures. These dimensions are, in principle, operational.

⁸ Importantly, this is true only for the case in which ρ is strictly separable in q and C . It is, of course possible to consider other functional form assumptions which would permit a richer set of relations between q and C (see, for example Blackorby, Primont, and Russell 1978).

Figure 2: Risk Matrix with Risk Contours



It should be clear that, except for simple visualization, the 2-dimensional characterization of the risk matrix is inessential. Even if we assume that $q(x_i)$ is the same for all policy domains, the evaluation of consequences of x_i will generally vary across policy domains, as will the identification of which commodities are critical. Suppose we now denote domain-specific consequences of a given bad outcome by $C^d(x_i)$, where $d \in \mathcal{D}$ (the set of all consequential domains). Most governments now treat environmental goals as core government objectives. Thus, in addition to macroeconomic consequences, $C^M(x_i)$, governments must incorporate environmental consequences, $C^E(x_i)$, in their evaluation of criticality. For example, because rare earth elements play an essential role in inputs to green transition (Depraeter and Goutte 2023), many governments will consider rare earth supply risks to both national economic performance and environmental policy. As a result, environmental consequences have been added to criticality evaluation (Graedel et al. 2012; Graedel et al. 2015). There will be matrices like those in Figures 1 and 2 for each policy domain.

The possibility that C^d -specific ρ can be determined independently does not mean that identification of a decision-maker's overall critical commodity set is independent. In particular, this decision relies on the weights the decision-maker attaches to the various C^d 's, a

fundamentally political decision. These weights will vary across countries as a result of differing political environments. Thus, even relatively similar countries will produce different sets of critical commodities. We suppose that the $\rho(x_i)$, and thus criticality, are only an input to the political economy of critical commodities. That is, once we determine which commodities are critical, these $\rho(x_i)$, and the components $q(x_i)$ and $C^d(x_i)$, become inputs to a broader policy choice problem.

II. Framework for Policy Choice Analysis with Risk Matrices

II.A The Case of a Single Objective: The Risk Environment

To develop a framework for policy analysis, we now assume that, instead of a commodity-specific x_i , there is a future state of the world, $x \in \mathcal{X}$, facing country $k \in K$ that is a complete description of the world (economic, environmental, political, geopolitical, etc.), except for policy. We isolate policy because we will eventually consider both a world where country k is policy-active and a world where all countries in K are potentially policy active.

For a given potentially critical good $i \in I$, from the perspective of country k (with a single objective, we can abstract from bureaus), we can evaluate the probability of a good (G) or bad (B) future outcome and the consequence of each outcome exactly as is done in the standard risk matrix.⁹ If we denote the probability of a bad outcome for commodity i at x as $q_B^{ki}(x)$ and the associated consequence of a bad outcome as $C_B^{ki}(x)$, then we can calculate risk as:¹⁰

$$\rho_B^{ki}(x) = q_B^{ki}(x) \times C_B^{ki}(x) \quad (1)$$

⁹ Practically, where the risk matrix literature evaluates consequences with a number of empirical measures, we will assume that consequences are identified relative to some specific target. Thus, for example, the macroeconomic authority might evaluate consequence in terms of unemployment. If there is more than one target, say unemployment and inflation, we would need to consider a multi-objective problem of the sort we present as representative of the peak decision-maker in the next section.

¹⁰ It is important to be clear that ρ , q , and C are functions of the state of the world, $x \in \mathcal{X}$, that can be good or bad for a given critical good. Thus, for example, $\rho_B^{ki}(x)$ identifies the risk of x being a bad state for commodity i . Thus, with 2 commodities (R and L) and 2 realizations (bad and good) for each commodity, there are four possible states. That is, denoting good and bad states for commodities R and L with subscripts G and B , $\mathcal{X} \equiv \{R_GL_G, R_GL_B, R_B L_G, R_B L_B\}$.

As a practical matter, the points representing the ρ_B^{ki} (for say lithium and REEs) in the risk matrix are identified independently from one another. In the world that generates those data, it is unlikely that ρ_B^{kL} and ρ_B^{kR} are actually independent.

We can illustrate the overall choice environment by supposing that analysts can evaluate q^{ki} and C^{ki} for all $i \in I$. For concreteness, suppose that there are two potentially critical goods, lithium (L) and REEs (R) and, for each good, x can be either bad $q_B^{ki}(x)$ or good $q_G^{ki}(x) = 1 - q_B^{ki}(x)$. Of course, it must be the case that $C_B^{ki}(x) < C_G^{ki}(x)$.

Looking forward from the moment of policy adoption, for country k , assuming that the q 's are independent, the probabilities of a given state are reported in Table 1. There are four possible realizations for this economy in the future. We can denote the total consequence for the macroeconomy as $C_M^k(\cdot)$, where the arguments are the commodity-specific outcomes in each of the four states.¹¹

Table 1: Possible Future States of the World

		Lithium		
		Good	Bad	
REEs	Good	$q_G^{kR} \times q_G^{kL}$	$q_G^{kR} \times q_B^{kL}$	$1 - q_B^{kR}(x)$
	Bad	$q_B^{kR} \times q_G^{kL}$	$q_B^{kR} \times q_B^{kL}$	$q_B^{kR}(x)$
		$1 - q_B^{kL}(x)$	$q_B^{kL}(x)$	

If we denote individual consequences by their probabilities, we obtain the outcomes reported in Table 2.¹²

Table 2: Possible Future Consequences

		Lithium	
		Good	Bad
REEs	Good	$C_M^k(x[R_G, L_G])$	$C_M^k(x[R_G, L_B])$
	Bad	$C_M^k(x[R_B, L_G])$	$C_M^k(x[R_B, L_B])$

¹¹ We will later consider the existence of other objectives such as environmental policy and national security policy.

¹² Table 2 abuses good notation for, we hope, clarity. What is relevant to the macroeconomic authority in each realized state is the consequence for each commodity, assuming x can be either good or bad for product i . Thus, we denote the realized consequences in each of the four cases by whether it is good (G) or bad (B) for R and L .

As in the risk matrix analysis, all entries in both tables can be estimated from currently available data. The points in the risk matrix for each commodity can be calculated as $\rho_B^{ki}(x) = q_B^{ki}(x) \times C_B^{ki}(x)$. However, it should be clear that, from the perspective of the macroeconomic authority, this is not sufficient information to determine a complete evaluation of the policy environment prior to policy adoption. If we denote the probability of a given cell occurring as $p_{GG}^k = q_G^{kR} \times q_G^{kL}$, $p_{GB}^k = q_G^{kR} \times q_B^{kL}$, and so on, we can write the expected consequence as:

$$\begin{aligned} \mathbb{E}^M[x] = & p_{GG}^k C_M^k \left(C_G^{kR}(x), C_G^{kL}(x) \right) + p_{BG}^k C_M^k \left(C_B^{kR}(x), C_G^{kL}(x) \right) \\ & + p_{GB}^k C_M^k \left(C_G^{kR}(x), C_B^{kL}(x) \right) + p_{BB}^k C_M^k \left(C_B^{kR}(x), C_B^{kL}(x) \right) \end{aligned} \quad (2)$$

II.B The Case of a Single Objective: Policy Choice

Consider a peak decisionmaker, concerned only with the macroeconomic consequences of policy, that can adopt a set of policies that will apply with certainty to the future state of the world.¹³ That is, we suppose that, after receiving an evaluation of those effects from the macroeconomic domain authority (M), the peak decision-maker selects a policy vector $\mathbf{a}^k$ from a set $\mathcal{A}$ of all possible policies.¹⁴ These policies might be general and apply to all critical materials (a uniform border measure, a uniform tax subsidy, etc.) or target particular critical commodities.¹⁵ The decisionmaker may choose not to use any or all of them. Thus, the realized state of the world facing country k in the future will be $(x, \mathbf{a}^k)$.

The domain-specific authority (M) must support this policy decision by providing an evaluation of all relevant policy mixes. For example, it might evaluate critical materials policies in terms of their effect on growth (or unemployment, or inflation, ...). In evaluating policy vectors, the

¹³ We leave open the possibility that the policymaker adopts a specific policy in $t = 0$ or adopts a commitment to adopt a policy in t_1 in the case of a specific realization. Here, we focus on the case of a policy adopted in $t = 0$ which cannot be changed in $t = 1$.

¹⁴ In fact, we will assume that the domain authorities and, ultimately, the peak decision-maker are constrained to a subset $S \in \mathcal{A}$. This permits us to recognize that the elements of S can vary over time and countries. For example, export taxes are unconstitutional in the US, but not in most other countries. Similarly, the US has entered into a number of international commitments (e.g. GATT/WTO, USMCA, etc.) that restrict its choice of policy instruments.

¹⁵ Note that the elements of $\mathbf{a}$ can range from very specific (e.g., subsidy to the production of a specific mineral), to broader policies (e.g., subsidies to domestic production of batteries, that use many critical commodities as inputs), to general policies (e.g. tariffs on imports of all critical commodities, or quantitative restrictions on exports of all critical materials). They may be continuous (like dollars of subsidy); on/off (like an export ban); or category bins (like high, medium and low; PTA, MFN, non-MFN). It is a modest convenience to order the measures by intrusiveness from low to high.

authority will consider the effect of a given vector on the components of ρ . Thus, for example, creating a stockpile of a critical material will have its main effect via a reduction in q , while support for research on substitutes will have its main effect via a reduction in the consequence of a reduction in access to imported supply. Most policies, including these, will affect both. The essential effect of concern is the way an increase or decrease in risk affects the policy target.

Using full information about the structure of the economy, M can evaluate the effects of any policy vector on, say, growth, and use this information to rank policies. Thus, we assume that this authority can rank policy vectors in terms of its expected macroeconomic consequence $\mathbb{E}^M[x, \mathbf{a}^k]$. We represent this ranking via the relation $\succsim_{kd}$. Thus, $\mathbf{a}^k \succsim_{kd} \mathbf{a}^{k'}$ is taken to mean that the authority for domain d in country k believes that policy vector $\mathbf{a}^k$ will reduce risk, as evaluated in terms of its policy target (growth), more than will $\mathbf{a}^{k'}$.¹⁶ It is a standard result that if $\succsim_{kd}$ is complete and transitive, it can be represented by a real-valued evaluation function.¹⁷ We denote this function $v^{kd}(\mathbf{a}^k)$. That is, it produces a number such that if $\mathbf{a}^k \succsim_{kd} \mathbf{a}^{k'}$ at x , then $v^{kd}(\mathbf{a}^k) \geq v^{kd}(\mathbf{a}^{k'})$. Thus, M seeks to:¹⁸

$$\begin{aligned} & \max_{\mathbf{a} \in \mathcal{A}} v^{kd}(\mathbf{a}^k) \\ & \text{s. t. } x \in \mathcal{X} \end{aligned} \tag{3}$$

Our analysis presumes that a domain authority reports this function to a central decision-maker, thus providing the information base on which it (i.e. d) evaluates alternative policy mixes from the perspective of domain d . That is, if M determines its policy preference based on the effect of a given policy vector on growth, that information is embedded in $v^{kd}(\mathbf{a}^k)$.

Starting with the case where the government possesses a single instrument of general applicability whose value, for illustrative purposes, can be either off (0) or on (1). That is, $a^k = 0$ or $a^k = 1$. If $a^k = 0$, then the probabilities and the consequences are exactly as in equation (2).

The expected consequence is:

¹⁶ Recall that we are focused on risk as defined in equation (1).

¹⁷ See, for example, Mas-Colell, Whinston, and Green (1995, Proposition 1.B.2). In consumer theory, this is usually called a “utility function” and, while “utility” has no psychological meaning beyond the fact that if $\mathbf{a}^k \succsim_{kd} \mathbf{a}^{k'}$, then the numbers given by the utility function must be such that $u(\mathbf{a}^k) \geq u(\mathbf{a}^{k'})$, we prefer to simply call the functions representing the order “evaluation functions” to remind readers that “governments” don’t have psychic states, so these reflect evaluations in terms of policy targets with no implication of utility as a psychological state.

¹⁸ Recall that maximizing v involves minimizing the risk arising from reduced access to critical materials.

$$\begin{aligned}\mathbb{E}^M[x, a^k] &= p_{GG}^k C_M^k \left(C_G^{kR}(x, a^k), C_G^{kL}(x, a^k) \right) + p_{BG}^k C_M^k \left(C_B^{kR}(x, a^k), C_G^{kL}(x, a^k) \right) \\ &\quad + p_{GB}^k C_M^k \left(C_G^{kR}(x, a^k), C_B^{kL}(x, a^k) \right) + p_{BB}^k C_M^k \left(C_B^{kR}(x, a^k), C_B^{kL}(x, a^k) \right)\end{aligned}\quad (4)$$

That is, if $a^k = 0$ the expected outcome is $\mathbb{E}^M[x, a^k] = \mathbb{E}^M[x]$. If $a^k = 1$, there is another set of cases. For each commodity, for each outcome (bad or good), a^k can either increase or reduce the risk. Similarly, for the consequences. There are still four cases, for each of the four outcomes, the consequence will be better or worse with a^k . The macroeconomic authority can directly compare the expected outcomes with and without policy. That is, $v_{a^k=0}^{kM} \geq v_{a^k=1}^{kM}$.

There is no reason, other than clarity of exposition, for the policy to take one of two values. Thus, we can suppose that single policy can take continuous values in, say, the unit interval-- $a_1 \in [0,1]$. The key thing to note is that while it is reasonable to assume continuity of v^{km} on $[0,1]$, there is no reason to assume anything else. In particular, not only will these functions not generally be monotonic, but there is also no particular reason for them to be single peaked. That is, it should be clear that the quantitative magnitudes of these effects will vary across the two commodities. Over some range (even at one point), the policy might be good for both, good for one but not the other, or bad for both. In the case where a given policy, over some range, is good for one sector and bad for another, the result of the analysis will depend on magnitudes of the effects and the way those effects are weighed in v^{km} .

Now suppose that there are two policies, $\mathbf{a}_1^k = (a_1^k, a_2^k)$ each of which can be off (0) or on (1)—for example, an export ban on one of the two goods. In this case, the feasible set of policies $\mathcal{A} = \{(0,0), (1,0), (0,1), (1,1)\}$. In any state of the world (see Table 1), these four policy combinations can be ranked by $\succeq_{kd}$ and there is no reason to suppose that there is any particular order to this ranking for any given state, or any particular relationship between the rankings across states. Thus, there are now 4 with-policy outcomes to be evaluated across the 4 possible states of the world. As the number of policies increase, the dimensionality of $\mathcal{A}$ increases as $2^{\#\mathcal{A}}$.¹⁹ The problem becomes more complex if we permit both policies to vary continuously. Even if we assume that the evaluations are monotonic (increasing or decreasing), or single-

¹⁹ If we are willing to associate a specific cost with each policy that is on, and suppose that we have a budget constraint, this is what, in the literature on combinatorial optimization, is called a knapsack problem (Pisinger and Toth 1998). There is a sizable literature on whether, and under what conditions, these problems are $\mathcal{NP}$ -hard (Krass, Sethi, and Sorger 1994).

peaked, unless we make a very strong assumption about the relationship between the effects of the two policies, we have very little information about the continuous effect of the policies on the evaluation function of the domain authority.

II.C The Case of Multiple Objectives: Policy Choice

Most research on policy choice is done from the perspective of a decision-maker with a single, well-specified objective. This is the path that we followed in the previous section. However, those choices can be very different depending on what that objective is. To this point we have identified the choice of the peak decision-maker with that of the macroeconomic authority in choosing policies to reduce the risk associated with the need for critical materials. However, many of these same commodities are also critical to the environmental and national security agendas of many states. Thus, we now define a set of policy domains, $d \in \mathcal{D}$, in terms of which a peak authority will set policy. In particular, we suppose that policy might be adopted to affect macroeconomic conditions, environmental conditions, and/or national security—that is, $d \in \{M, E, S\}$.²⁰

Assume that each domain has a domain-specific authority (a “bureaucracy”) responsible for evaluating policies reflecting the concerns of that domain. Each of these applies some method, represented by equation (3), embodying that method which reflects domain-specific knowledge and the concerns of the policy-making community in that domain (as outlined above for the macroeconomic authority). These methods will be reflected in the consequences measured by the relevant target variables. Thus, the environmental authority might focus on particulates generated by burning carbon-based fuels and the national security authority might focus on perceived geopolitical risk. Because the peak decision-maker knows the v^{kd} , it knows the basis of the d -specific rankings of policy vectors.

The peak decision-maker will ultimately choose the policy vector for the country. Given the presence of multiple objectives, and the need to select one policy vector, it must make a decision that will generally differ from the preferred policies of the individual bureaucracies. If we suppose that the policy domains represent genuine interests of the state, and the bureaucracies represent the actual goals relative to those domains, the peak decision-maker’s choice is

²⁰ Table 1 still gives the states of the world without intervention.

complex. For example, rare earths are not actually rare but mining them is environmentally quite damaging (Riesman and Weber 2012). As a result, many countries chose to restrict such mining and outsourced the supply to China (Kennedy 2016). The consequences of this were made apparent to the United States when China used export restrictions on rare earths in response to US tariffs. Where the environmental authority will prefer to restrict local mining, the national security authority might propose relaxation in environmental regulations associated with mining. As a result, the components of the $\mathbf{a}^{kd}$ vectors cannot simply be averaged (or combined under some weighting scheme).²¹ Thus, we are in the realm of multi-objective decision theory (Keeney and Raiffa 1993; Yu 1985).

The Peak Decision-Maker's Problem

If we denote the top-level decision-maker's (the state's) objective function as $G(\bullet)$, we can write her problem as:²²

$$\begin{aligned} \max_{\mathbf{a}^k \in \mathcal{A}} G\left(v^{kM}(\mathbf{a}^k; x), v^{kE}(\mathbf{a}^k; x), v^{kS}(\mathbf{a}^k; x)\right) \\ \text{s. t. } \mathbf{a}^k \in \mathcal{A} \\ x \in \mathcal{X} \end{aligned} \tag{5}$$

As a first step, we consider the state's feasible set for the case where both policies are either off (0) or on (1), and that $\mathcal{A} = \{(0,0), (1,0), (0,1), (1,1)\}$. Furthermore, for illustrative purposes, we assume that the macroeconomic authority prefers no intervention (0,0), the environmental authority prefers that both policies be adopted (1,1), and the security authority prefers a policy targeted on commodity R , but not on commodity L (0,1). It is easy to see, in this case, that (1,0) is dominated, but that the other three policies remain feasible for the state. The state can only accommodate one of the domain authorities. Alternatively, suppose that each of the policies takes a value in the unit interval. If the evaluation function of each domain authority is single-peaked preference over the unit square, we know from the classic work of Plott (1967) and McKelvey


²¹ For intuition, suppose that in the feasible set, $\mathcal{A}$, $\mathbf{a}^d$ vectors policies that are either "on" or "off". Unless the best policy vector for each domain is identical to the other two domains, there will be a conflict that must be resolved.

²² As we noted in the introduction, G reflects the constraints on the peak-level decision maker, while the v reflect the constraints on the domain-specific authorities (e.g. states of knowledge, technology, etc.).

(1979) that the Pareto sets associated with this case can be quite large.²³ This result is generic for $d \geq 3$.

Consider Figure 3. The unit square is $\mathcal{A}$. The ideal policy from the perspective of each of the domain authorities is one of the points in this space, around which we have drawn level curves for evaluations that decrease away from those points. The tangencies between the curves and the triangular area identified by them is the Pareto set. Any point in the unit square that is not in the Pareto set, is dominated by points in the Pareto set.²⁴ Thus, the Pareto set is the set of undominated policies that might be chosen by the state. The essential point to note is that there are an infinite number of undominated points and at none of the undominated points is there a tangency between all three indifference curves. In this paper, we are not concerned with the way $\mathbf{a}^k$ is selected from the nondominated solutions of this problem. We presume that there is some political process embedded in G , but that is not our concern here. Rather, the point is that, in choosing policies and in constructing international agreements about such policies, governments have multiple, often inconsistent, objectives.

Figure 3: The Pareto Set



Example: Cobb-Douglas Objective Function

The state will make some choice. The first step in the process is identifying the set of undominated policies, the next is to take into account all of the conjunctural factors (political in

²³ This is not the main point of either of these papers, which consider problems with the method of majority decision-making, but both present a preliminary analysis of the Pareto set on the way to a demonstration that this method need not select points in the Pareto set.

²⁴ We have drawn the level sets of the $v(x)$ functions as radially symmetric for graphical convenience. As Plott (1967) and McKelvey (1979) show, this is inessential.

particular) in selecting among the undominated policies. If we suppose that G is strictly quasi-concave over the undominated set, this will select a unique policy vector. For illustrative purposes, suppose that this is a Cobb-Douglas function of the three objectives:

$$G \equiv \left(v^{kM}(\mathbf{a}^k; x)\right)^\alpha \left(v^{kE}(\mathbf{a}^k; x)\right)^\beta \left(v^{kS}(\mathbf{a}^k; x)\right)^\gamma, \alpha + \beta + \gamma = 1,$$

where the parameters (α, β, γ) are themselves functions of those conjunctural factors. It is in the nature of conjunctural factors that they cannot be clearly identified *ex ante*, and yet they play a fundamental role in determining a given equilibrium choice and changes in that choice. For example, if β is greater for a government that prioritizes combating climate change, a shift in government to parties that prioritize security will produce a shift in the $\mathbf{a}$ vector to a set of policies with a smaller emphasis on environmental policy.

For the case of two policy domains (e.g. macroeconomic and environmental) we can illustrate the policy trade-off graphically (Figure 4). We can write the objective function as:

$$G \equiv \left(v^{kM}(x, \mathcal{A})\right)^\alpha \left(v^{kE}(x, \mathcal{A})\right)^\beta, \quad \alpha + \beta = 1$$


$$\mathcal{A} = \{(0,0), (1,0), (0,1), (1,1)\}$$

and let $\hat{v} = \log v$.

The vertical axis in Figure 4 measures $\hat{v}^{kM}(x, \mathcal{A})$ and $\hat{v}^{kE}(x, \mathcal{A})$. With two policies and binary values there are four possible values for each function. For $\alpha = 1$ the values of $\hat{v}^{kM}(x, \mathcal{A})$ are also the values for G and for $\beta = 1$ the values of $\hat{v}^{kE}(x, \mathcal{A})$ are also the values for G . For intermediate weights the values are given by points on the solid lines connecting the corresponding value functions. For example, for $\alpha = \beta = 0.5$ the values of G for each of the four policy options are given by the midpoints of those lines.

In this example we have: $v^E(1,1) > v^M(0,0) > v^E(1,0) > v^M(1,0) > v^M(0,1) > v^E(0,0) > v^M(1,1) > v^E(0,1)$. Notice that the policy (0,1) is clearly dominated by the policy (1,0) and will never be chosen. For $\alpha = 1$ the optimal policy choice is (0,0) while for $\beta = 1$ the optimal policy choice is (1,1). For $\alpha = \beta = 0.5$ the optimal policy choice is (1,0). The bold segments correspond to the optimal policy choices for each value of $\beta = 1 - \alpha$.

Figure 4: Policy trade-offs



In terms of the language we will deploy in the next section, $G(\bullet)$ summarizes the *objectives* of the government (it is naturally the objective function). The $\mathbf{a}^k$ -vector ultimately chosen, i.e. the policies adopted by the government, are expected to achieve specific *goals* (i.e. identifiable effects that serve the objectives).²⁵ Thus, a goal is a first-order result of a policy. For empirical work, a central part of the inference structure assumes that countries facing common risk environments and confronting common costs from restricted access to critical materials will adopt relatively common policies ($\mathbf{a}^k$ -vectors). We expect policies to vary with both the risk and cost environments. For example, two countries might have very similar economic structures with respect to the use of critical commodities (i.e. their $C(\bullet)$ functions will be very similar) but one might possess natural sources of supply and thus would have a low ρ for the given good (that country would probably not declare that good “critical” and its $\mathbf{a}^k$ -vector would not reflect concern about supply). However, as we develop in the next section, even when countries have similar risk and cost environments, and thus similar evaluations of the first-order effects on policy domains, they might well prefer different policy mixes as a result of differing objectives.

²⁵ We might think of objectives as the “ends” of policy and goals as the “means”.

II.D Adding international considerations

When all countries are potentially policy-active, the choice problem of country k must be specified relative to the policies of other governments. To incorporate this, we denote by $\neg k \in K$ the set of all countries other than k , and by $\mathbb{a}^{\neg k}$ the matrix collecting the policy vectors of those countries. Just as domestic policy $\mathbf{a}^k$ can affect either the probability of a given bad outcome or the consequence of that outcome, foreign policies $\mathbb{a}^{\neg k}$ will also enter both the q - and the C^d -terms of country k 's risk matrix. China's use of export restrictions on rare earth elements in response to U.S. tariffs is an obvious example: a shift in Chinese export policy changed the perceived probability of disruption for other states and, by extension, the evaluation of rare-earths risk in their macroeconomic and security domains.²⁶ In the framework developed above, this can induce a change in the national-security authority's value function $v^{kS}(\cdot)$ and in the peak decision-maker's choice of $\mathbf{a}^k$.

In a world with multiple, policy-active governments, the peak decision-makers in country k will take the existing policy environment into account when identifying the overall optimal policy. Using the information contained in the maximum value functions for the domain authorities, we can write a given country's overall policy problem for critical materials as:²⁷

$$\begin{aligned} \max_{\mathbf{a}^k \in S^k} & G(v^{kM}(\mathbf{a}^k; x, \mathbb{a}^{\neg k}), v^{kE}(\mathbf{a}^k; x, \mathbb{a}^{\neg k}), v^{kS}(\mathbf{a}^k; x, \mathbb{a}^{\neg k}); \mathbb{a}^{\neg k}) \\ \text{s. t. } & \mathbf{a}^k \in S^k \\ & \mathbb{a}^{\neg k} \in \prod_{j \neq k} S^j \\ & x \in \mathcal{X} \end{aligned} \tag{6}$$

The central authorities of other countries face a similar problem. Clearly these problems are interdependent and finding the overall solution can be very complex. Under the strong supposition that each central authority knows the objective functions (this would include both the G and v) of all central authorities, the solution, if it exists, would be a Nash Equilibrium: each country's policy vector $\mathbf{a}^k$ will be an optimal response to the policy vectors of all other countries,

²⁶ These measures also affected the EU as China imposed a licensing regime on all exports of REEs. The EU obtains all of its heavy REEs, 85% of its light REEs and 98% of its rare-earth magnets from China (Szczepański, 2025).

²⁷ As we noted in footnote 11, we now need to recognize that the feasible set of policies may vary across countries. Thus, we now assume that every country k has access to a k -specific feasible set of policies $S^k \in \mathcal{A}$.

$\mathfrak{a}^{-k}$. Formally, the solution to problem (6) yields an optimal response (reaction) function for country k , $\hat{\mathbf{a}}^k = \varphi(\mathfrak{a}^{-k}) = \arg\max_{\mathbf{a}^k \in S^k} G(\cdot)$. A Nash equilibrium is then defined as set of K policy vectors, one for each country, such as $\hat{\mathbf{a}}^j = \varphi(\hat{\mathbf{a}}^{-j})$ for every $j \in K$.

The existence of Nash-equilibrium requires that any unilateral deviation is unprofitable. However, for the present purpose, the most relevant international policies are those that involve explicit cooperation among a subset of countries. To capture this, we allow countries to enter into non-binding, self-enforcing agreements and we use Coalition-Proof Nash (Bernheim et al. 1987) as the equilibrium concept. When a group of countries conclude such an agreement, we say that they form a “critical materials plurilateral” (CMP) club.

For concreteness, let J be a proper subset of the set of countries K .²⁸ Suppose that the peak decision makers of these countries have decided to choose jointly some of the $\mathbf{a}^k$, $k \in J$, policy parameters. We denote the matrix of policy parameters chosen by club members (a vector of parameters for each member) as $\mathfrak{a}^k(J)$, by $\mathfrak{a}^{-k}(J)$ the corresponding matrix for all members other than country k , and by $\mathfrak{a}(-J)$ the corresponding matrix of all other countries. Clearly if the parameters in $\mathfrak{a}^k(J)$ were the same as those that club members would have chosen under a conventional Nash equilibrium the formation of the plurilateral agreement would be redundant. The new set of strategies (including the responses of the peak decision-makers of the complement set of countries) will be a Coalition-Proof Nash equilibrium if the following condition is satisfied: there is no proper subset of players taking the actions of the complement set of players as fixed that can deviate so that all its members are better off and all considered deviations have also to be self-enforcing. This last requirement differentiates the Coalition-Proof Nash and the Strong Nash (Aumann 1959) concepts. In our example, just by finding a subset of the set of countries among those in the set J that formed the plurilateral that can profitably deviate is sufficient to show that the plurilateral cannot be part of a Strong Nash equilibrium. However, it can still be a part of a Coalition-Proof Nash equilibrium if the proposed deviation itself is not self-enforcing.

²⁸ This subset may be as small as two, i.e., a bilateral agreement. We consider the more general case that includes bilateral agreements. In principle, plurilateral agreements spanning more than two countries are likely to be more effective given that critical materials supply chains often will span multiple countries.

III. Critical-Materials Plurilaterals through the Lens of the Model

In this section we apply the framework developed above to the design and political economy of plurilateral arrangements for critical materials. Plurilateral agreements can be understood as institutions that provide common instruments affecting both the probability and the consequences of supply disruptions, negotiated by governments that are themselves solving multi-objective decision problems.

Conceptually, a CMP is a club that provides instruments which operate on both the probability and consequence terms in the risk matrix. On the probability side, a CMP may coordinate investment in new mining and processing capacity, operate joint strategic stockpiles, and establish information-sharing and early-warning mechanisms. For members, these measures can reduce the probability of bad outcomes $q_B^{ki}(x)$ for some set of critical materials, conditional on a given profile $(a^k(J), a(\neg J))$. On the consequence side, a CMP can promote cooperative research on substitution and recycling, harmonize environmental and labor standards in extraction and processing, and support resilient downstream production in sectors such as batteries and permanent magnets. These instruments act primarily on the consequence terms $C_B^{kd}(x)$, potentially lowering the macroeconomic, environmental, and security costs of a disruption for participating countries.

III.A Objectives and domains at the plurilateral level

At the national level, we have assumed that the state's decision problem can be represented as the choice of a policy vector a^k that maximizes an aggregating function $G(\cdot)$ defined over domain-specific evaluation functions v^{kM}, v^{kE}, v^{kS} . The macroeconomic authority evaluates policies in terms of their implications for growth, employment, or inflation; the environmental authority focuses on emissions and ecological damage; and the security authority is concerned with military readiness and geopolitical vulnerability. Each domain uses a risk-matrix logic—probabilities of disruption and domain-specific consequences—but the peak decision-maker must choose a single policy vector. At the plurilateral level, these underlying structures are not eliminated. Instead, each member government arrives with its own domain-specific value functions and its own evaluation of those domains.

Crucially, the same CMP will not be evaluated in the same way by all domestic authorities. From the macroeconomic perspective, joining a CMP may be attractive because deeper allied markets, shared stockpiles, and more diversified supply chains reduce expected output and employment losses in bad states of the world. At the same time, membership can entail constraints on the use of trade and industrial policies – for example, disciplines on export restrictions or local-content requirements—that the macroeconomic bureau might regard as valuable in some states of the world. The resulting $v^{kM}(\cdot)$ need not be monotonic in the depth of CMP commitments: shallow membership (with limited information-sharing and weak disciplines) may dominate both non-membership and deep, highly constraining cooperation. The environmental authority may view the same CMP quite differently. If the plurilateral embeds high common standards, participation lowers the environmental cost of expanding supply and supports the green transition, increasing v^{kE} . If, instead, the CMP relaxes standards in the interest of rapid capacity expansion, the environmental authority may prefer non-membership combined with stricter domestic regulation, even at the cost of higher disruption risk. Finally, the security authority will typically value a CMP that reduces dependence on a single non-allied supplier, supports allied processing capacity, and provides coordinated responses to coercive export policies, but may view with suspicion commitments that limit the ability to weaponize access to critical materials against non-members.

The tension between incentives to join and not to join therefore arises from two sources. First, there are genuine policy trade-offs at the level of each domain: joining a CMP moves both $q_B^{ki}(x)$ and $C_B^{kd}(x)$ in ways that may help in some states of the world and hinder in others, and the sign and magnitude of these effects differ across domains. Second, the peak decision-maker aggregates these domain-specific evaluations using weights that reflect domestic politics and conjunctural factors. A government that places greater weight on environmental and security objectives may find that membership in the CMP increases its objective $G(\cdot)$ for a given international environment, while a government with stronger emphasis on traditional industrial policy autonomy may prefer non-membership even in the presence of a seemingly attractive CMP. Because the benefits from the CMP (for example, in terms of diversification and shared capacity) depend on the membership decisions of others, the problem is further complicated by club-like externalities: as the size of the plurilateral expands, the gains to joining for a marginal country typically rise, but so may the perceived costs of accepting common disciplines. In the

language of the model, even if countries share broadly similar risk matrices for particular materials, differences in domain-specific value functions and in the weights applied by their peak decision-makers imply heterogeneous equilibrium choices over CMP membership.

A club does not maximize a single global “criticality” function; rather, it aggregates governments that differ systematically in how they trade off macroeconomic performance, environmental protection, and security. The negotiation over the club’s rules and instruments is therefore a negotiation over how much weight to assign to each domain in the common package. A security-oriented government will push for arrangements that expand allied capacity and restrict dependence on geopolitical rivals; an environmentally oriented government will seek to embed strong environmental and labor standards; and a macro-oriented government will prioritize instruments that minimize expected costs of disruptions for the domestic economy, even if that implies weaker standards or less ambitious diversification. The resulting plurilateral design reflects a compromise over these divergent domain-specific preferences, not the solution to a single maximization problem.

III.B Plurilateral instruments as operations on q and C

The risk-matrix framework provides a useful way of organizing the instruments that a critical-materials plurilateral might deploy. At the level of individual countries, we have distinguished between policies that primarily affect the probability of a bad outcome $q_B^{ki}(x)$ and policies that primarily affect the consequence terms $C_B^{kd}(x)$ in the different domains. The same distinction applies at the club level.

On the “probability” side, a plurilateral can:

- Coordinate investment in new mining and processing capacity across members, thereby diversifying supply away from a narrow set of producers and reducing the probability of severe disruption in any given state of the world.
- Establish shared strategic stockpiles and mutual support arrangements that function as insurance against idiosyncratic shocks and third-country export restrictions.
- Develop common early-warning systems and information-sharing mechanisms that allow members to respond more quickly to emerging risks in supply chains.

- Establish disciplines on the use of certain trade and industrial policies among members (for example, time-consistent rules on export restrictions in crises) that reduce the risk that members themselves become sources of disruption.

On the “consequence” side, a plurilateral can:

- Promote cooperative R&D on substitution, recycling, and material efficiency to lower the technological and economic costs of adjusting to a disruption.
- Embed common environmental and labor standards in mining and processing, thereby affecting the environmental consequences $C_B^{kE}(x)$ associated with marginal supply expansions.
- Provide coordinated support for strategically important downstream activities (such as permanent magnets, batteries, and semiconductor components) that are central to macroeconomic and security objectives, thereby limiting the macroeconomic and strategic consequences $C_B^{kM}(x)$ and $C_B^{kS}(x)$ of shocks.

From the perspective of an individual country, joining a plurilateral that supplies such instruments is equivalent to gaining access to an augmented menu of policies that act on both q_B^{ki} and C_B^{kd} , subject to the constraints embedded in the club’s rules. The key point is that the same common package of club-level instruments will be evaluated differently by each member’s domain-specific authorities. The macroeconomic authority may value deeper export disciplines and shared stockpiles highly; the environmental authority may prioritize strong standards and green R&D; the security authority may insist on robust allied capacity and the ability to discriminate against rivals. There is no reason to expect that any given plurilateral design will lie at the ideal point of any member’s domestic bureaucracy.

III.C Membership, depth, and heterogeneity

Our model suggests two distinct but related decisions at the country level: whether to join a plurilateral and, conditional on joining, how much “depth” in terms of commitments and disciplines the government is willing to support. Both are shaped by the structure of the domestic decision problem and by the profile of other countries’ policies and membership choices.

To fix ideas, suppose that a given plurilateral is characterized by a “depth vector” $\delta = (\delta_q, \delta_c)$, where δ_q measures the extent to which membership reduces disruption probabilities for key materials and δ_c summarizes the extent to which membership changes the relevant consequence terms in each domain. A deeper club (e.g., with stronger commitments to coordinated investment, stockpiling, standards, and disciplines on export restrictions) will typically be associated with larger absolute values of δ_q and δ_c . The sign and magnitude of these changes are club-specific: deeper export disciplines may reduce q_B^{ki} for members by making the supply environment more predictable, but they may also increase the perceived cost of policy constraints for domestic authorities that value autonomy.

For any given (δ_q, δ_c) and international policy profile $(\mathbb{a}^k(J), \mathbb{a}^k(\neg J))$, a country’s domain-specific authorities will update their value functions v^{kM}, v^{kE}, v^{kS} and report their preferred policy vector, including a recommendation on membership. The peak decision-maker then compares the value of the objective function $G(\cdot)$ under membership and non-membership and under alternative depths of commitments. Because the benefits of the plurilateral typically increase with the number and characteristics of its members (the classic club-good externality), the incremental gain from joining will depend on the existing membership set J . As more large, technologically advanced economies join and implement concrete cooperative projects, the potential reductions in $q_B^{ki}(x)$ and $C_B^{kd}(x)$ for additional members increase, but so may the distributional and political costs associated with accepting the club’s rules.

Under plausible assumptions about the heterogeneity of domain-specific value functions and of the weights applied by peak decision-makers, the model predicts non-trivial, and often incomplete, participation in critical-materials plurilaterals. Even countries that share broadly similar risk matrices for particular materials may take different positions on membership and depth because they weigh environmental risks differently, face different domestic political constraints, or assign different importance to security relative to macroeconomic performance. This mechanism helps explain why, in practice, one would expect to observe a variety of overlapping and partially inclusive initiatives in the critical-minerals space rather than a single, comprehensive global agreement.

III.D Comparative statics: shocks to objectives and club design

The framework developed in section II allows us to interpret changes in critical materials policy, including participation in plurilateral arrangements, as responses to shocks that either alter the mapping from policies to domain-specific outcomes or reweigh the importance of different domains. In the Cobb–Douglas example, we assumed that the peak decision-maker’s objective could be written as

$$G \equiv (v^{kM}(\mathbf{a}^k; x))^\alpha (v^{kE}(\mathbf{a}^k; x))^\beta (v^{kS}(\mathbf{a}^k; x))^\gamma, \alpha + \beta + \gamma = 1,$$

where v^{kM}, v^{kE}, v^{kS} are the value functions reported by the macroeconomic, environmental, and security authorities, respectively, and (α, β, γ) capture the relative weight that the peak authority attaches to each domain. With two domains and a discrete set of policies $\mathcal{A} = \{(0,0), (1,0), (0,1), (1,1)\}$, Figure 4 illustrated how the optimal policy choice changes as the relative weight on environmental versus macroeconomic concerns change. For $\alpha = 1$ the macroeconomic authority’s preferred policy is chosen; for $\beta = 1$ the environmental authority’s preferred policy is chosen; and for intermediate weights the state selects a compromise policy.

Critical materials plurilaterals can be analyzed in precisely the same way. In this context, the relevant policy vector $\mathbf{a}^k$ includes both domestic instruments and the membership decision m^k , and the discrete policies whose values are plotted in Figure 4 can be interpreted as alternative combinations of domestic and plurilateral commitments. One corner might correspond to non-membership with limited domestic intervention; another to an aggressive domestic industrial policy without participation in a club; a third to shallow membership in a plurilateral with modest disciplines; and a fourth to deep membership in a plurilateral with strong common standards and coordinated investment. For each of these combinations, the domain authorities compute $v^{kM}(\mathbf{a}^k; x, \mathbb{a}^k(J), \mathbb{a}^{-k}(J))$, $v^{kE}(\mathbf{a}^k; x, \mathbb{a}^k(J), \mathbb{a}^{-k}(J))$, and $v^{kS}(\mathbf{a}^k; x, \mathbb{a}^k(J), \mathbb{a}^{-k}(J))$, and the peak decision-maker compares the resulting values of $G(\cdot)$ as the weights (α, β, γ) vary. The “bold segments” in the Figure 4 represent the subset of policy combinations that are optimal for some configuration of (α, β, γ) . As political or conjunctural factors change the weights, the state may switch discretely from one corner solution to another, including changes in membership and in the depth of plurilateral commitments.

Shocks can also operate through the domain-specific value functions themselves. As rare earth elements and related materials become more central to growth sectors and to the energy transition, the macroeconomic consequences $C_B^{kM}(x)$ and environmental consequences $C_B^{kE}(x)$ of a given disruption increase, raising the evaluated gains from policies that reduce either the probability of bad states or their consequences. In our notation, this can be represented as an upward shift in $v^{kM}(\mathbf{a}^k; x, \mathbb{a}^{-k}(J), \mathbb{a}(\neg J))$ and $v^{kE}(\mathbf{a}^k; x, \mathbb{a}^{-k}(J), \mathbb{a}(\neg J))$ at policy points that expand access—such as domestic production support, participation in a plurilateral with coordinated investment, or both. Similarly, an increase in geopolitical tension or the demonstrated use of export controls by a dominant supplier raises the security consequences $C_B^{kS}(x)$, thereby increasing $v^{kS}(\cdot)$ at policy vectors that diversify supply toward allies, develop allied processing capacity, or embed robust crisis-coordination mechanisms. In terms of Figure 4, these shocks change the vertical coordinates of the domain-specific evaluation points associated with each policy combination and hence reshape the upper envelope that determines the state’s optimal choice for a given (α, β, γ) .

Taken together, shocks to (α, β, γ) and to v^{kd} provide a simple comparative-statics lens on the emergence and evolution of critical materials plurilaterals. A government that initially chooses non-membership in a shallow cooperative arrangement may, after a shock that increases the weight on environmental and security objectives or raises the perceived costs of disruption in those domains, find that its optimal policy lies at a different corner, one involving deeper membership in a club that coordinates investment, stockpiling, and standards. Conversely, a government that experiences a change in domestic politics that lowers the weight on environmental concerns may move from a deep, standards-oriented plurilateral equilibrium to a configuration that retains some security cooperation but relaxes environmental disciplines. The next subsection uses this lens to interpret recent changes in critical materials policy.

IV. Applied Illustrations: Examples of critical minerals clubs

IV.A From MSP to FORGE

The Minerals Security Partnership (MSP), launched in 2022, is a natural starting point for an applied illustration. As of 2025 it had 15 members (Australia, Canada, Estonia, Finland, France, Germany, India, Italy, Japan, Norway, Korea, Sweden, United Kingdom, United States and the

EU). The MSP was designed to be a policy coordination mechanism for participating states and a platform to facilitate high-level critical minerals-related project identification. It was complemented in 2024 with an MSP Forum to support deliberations between MSP members and producer/supplier countries.²⁹ The Forum on Resource Geostrategic Engagement (FORGE), launched in February 2026, built on the MSP, providing a platform to go beyond identification of potential projects and state-level policy dialogue to also cover investment commitments and project implementation. FORGE aims to provide a structured framework for project development, risk mitigation, and financing of projects by businesses working alongside public entities (Ahn, 2026). It is intended to continue to serve as a platform for participating states to coordinate critical minerals policies but go beyond the MSP framework to deepen cooperation on projects to enhance critical mineral supply chain security. Its operational details have yet to be determined, but reportedly FORGE may include setting price floors for minerals procurement (US Department of State, 2026), e.g., contracts for difference.

The MSP as a predecessor club. The MSP fixed the configuration of domain weights and club design against which subsequent developments can be assessed. The MSP can be interpreted as a critical-minerals club among broadly like-minded economies aimed at promoting secure and sustainable supplies of materials such as lithium, cobalt, nickel, graphite, and REEs. In the terms developed above, the MSP operationalized a particular combination of instruments affecting both the probability and consequence components of members' risk matrices.

On the probability side, MSP members emphasized coordinated support for mining and processing investment outside a narrow set of existing producers, together with information-sharing and joint project-risk assessment among export-credit and development-finance institutions – measures that, in our framework, reduce $q_B^{ki}(x)$ for a subset of materials. On the consequence side, the MSP linked security-of-supply objectives explicitly to environmental and social standards, raising v^{kE} for participants while supporting allied processing capacity in ways intended to raise v^{kM} and v^{kS} as well. The resulting depth vector $\delta = (\delta_q, \delta_c)$ was, by design, weighted toward environmental and governance disciplines as the price of membership. This configuration suited a particular alignment of domain weights. Members were concentrated

²⁹ These included Argentina, DRC, Dominican Republic, Ecuador, Greenland, Kazakhstan, Mexico, Namibia, Peru, the Philippines, Serbia, Türkiye, Ukraine, Uzbekistan, and Zambia. https://policy.trade.ec.europa.eu/help-exporters-and-importers/accessing-markets/raw-materials/minerals-security-partnership-forum_en

among high income democracies for whom the weights (α, β, γ) in $G(\cdot)$ placed substantial mass on environmental and security objectives, and for whom domestic politics rewarded combining green and security narratives.

FORGE as a successor club. In February 2026 the Trump administration announced FORGE as the explicit successor to the MSP, alongside a parallel domestic stockpiling vehicle (Project Vault) and a wave of new bilateral minerals frameworks. The shift to FORGE for cooperation on critical minerals illustrates how a club’s design and its equilibrium activity level are separate things, and that the latter can lag the former even absent any change in the underlying risk matrices. FORGE was launched at an inaugural Critical Minerals Ministerial meeting of more than fifty countries, with chairing responsibilities rotating to the Republic of Korea through mid-2026 and the United States driving the forum’s project work (Baskaran and Schwartz 2026). Unlike the MSP’s emphasis on standards-linked coordination, FORGE has been described as a plurilateral coalition oriented around a preferential trade-and-investment zone (Miller and Morandi 2026). FORGE is expected to encompass coordinated reference prices or price floors intended to address disincentives affecting long-cycle mining investment in participating countries (Baskaran and Schwartz 2026).

FORGE maps onto the same q_B^{ki} and $C_B^{kd}(x)$ decomposition as the MSP, but with a different balance of instruments. On the probability side, FORGE’s ambition to set coordinated price floors and provide members with preferential market access functions much like the MSP’s investment-coordination measures, aiming to reduce $q_B^{ki}(x)$ by making capacity expansion more predictable for members. Accompanying bilateral frameworks — more than twenty signed within several months of the ministerial — extend this logic outside the core plurilateral agreement, consistent with the model’s prediction that overlapping, partially inclusive arrangements will be the norm rather than the exception.

On the consequence side, however, FORGE’s center of gravity sits closer to v^{kM} and v^{kS} than to v^{kE} . Reference pricing and preferential access primarily address macroeconomic exposure to price volatility and security exposure to coercive market manipulation by dominant suppliers. The environmental and governance disciplines that were central to the MSP’s design are not a comparable focus of FORGE’s public framing. In terms of the depth vector, FORGE can be read

as retaining, or even deepening, commitments on δ_q , while pulling back on the environmental-standards component of δ_c relative to its predecessor.

The transition from MSP to FORGE is a natural illustration of the comparative-statics logic developed in the previous section, i.e., a shock to (α, β, γ) . Holding the underlying physical risk matrices for REEs, lithium, and related materials roughly fixed, what changed between 2022 and 2026 was the weighting that the relevant peak decision-makers assigned to environmental relative to macroeconomic and security domains — alongside, plausibly, an increase in the perceived consequence terms $C_B^{kM}(x)$ and $C_B^{kS}(x)$ themselves, as China’s 2025 export-licensing restrictions sharpened the macroeconomic and security stakes of dependence for the U.S., the EU and other OECD member countries.

In the language used previously, a government that experiences a change in domestic politics lowering the weight on environmental concerns may move from a deep, standards-oriented plurilateral equilibrium to a configuration that retains security cooperation but relaxes environmental disciplines. The shift from the MSP to FORGE is close to a textbook case of exactly this corner-to-corner move. It is not a story about the underlying criticality of rare earths or cobalt changing — $q_B^{ki}(x)$ and the physical consequence terms are not what moved. It is a story about a change in (α, β, γ) , consistent with a new U.S. administration, changing priorities in the EU and other US allies, and a reweighting of domain priorities, that relocates the optimal $\mathbf{a}^k$ to a different point on the same Pareto frontier described in Section III. Both the MSP and FORGE sit on the same envelope of feasible club configurations for the same set of underlying risk matrices. What changed in going from MSP to FORGE is the preferred point on that envelope for a government with updated weights.

IV.B Pax Silica and overlapping clubs

The model’s prediction of partial, overlapping membership is reinforced by the near-simultaneous emergence of Pax Silica, a US State Department-coordinated initiative launched in December 2025 with a narrower membership (including Japan, the Republic of Korea, Australia, the United Kingdom, and several Gulf states). Pax Silica’s mandate spans a wider “silicon stack”

than FORGE — from mineral extraction through semiconductors and AI infrastructure — and was framed by U.S. officials as a positive-sum coalition rather than an exclusionary bloc.

For our purposes, Pax Silica is best treated not as a rival CMP but as a club defined over a different, partially overlapping set of domains: it shares FORGE’s emphasis on v^{kM} and v^{kS} , but extends the relevant consequence terms beyond critical-materials supply into downstream technology security. Its narrower, more capability-selected membership — weighted toward countries with semiconductor manufacturing, chip-design, or sovereign capital rather than toward the broad ministerial coalition FORGE assembled — is consistent with the model’s club-good externality logic: as the relevant good becomes more technologically specific, the set of countries for whom membership clears the participation constraint shrinks accordingly. That the United States is simultaneously party to FORGE, its more limited in scope predecessor, the MSP, Pax Silica, a growing lattice of bilateral frameworks and a member of G7 initiatives (e.g., the 2025 G7 Critical Minerals Action Plan) is not institutional incoherence so much as the predicted outcome of multi-objective, multi-club bargaining among heterogeneous governments.

The same observation holds for the various arrangements that have been put in place or are being pursued by the EU and other OECD countries. In the case of the EU these include Clean Trade and Investment Partnerships (CTIPs), framework agreements through which the EU seeks to manage strategic dependencies and bolster the security of critical global value chains by improving access to raw materials, clean energy and clean tech (European Commission, 2025). The first CTIP was concluded in 2025 with South Africa and included a Euro 4.7 billion investment commitment.³⁰ The CTIPs are led by the EU’s Directorate General (DG) for Trade and Economic Security but also involve other parts of the EU Commission as well, including the DG for International Partnerships and the DG for Climate, Net Zero and Clean Growth. They therefore illustrate the pertinence of the model developed above, as each of these “bureaus” will

³⁰ As an element of implementing the 2024 Critical Raw Materials Act, the EU has also concluded partnerships on sustainable raw materials with Argentina, Australia, Canada, Chile, the Democratic Republic of Congo, Greenland, Kazakhstan, Namibia, Norway, Rwanda, Serbia, Ukraine, Uzbekistan and Zambia (Jütten, 2025). The Act identifies a list of 34 raw materials deemed critical, sets a limit of 65% of the EU’s annual needs of specified strategic raw materials at any relevant stage of processing that can be sourced from a single non-EU country and imposes de-risking performance benchmarks: by 2030 at least 10% of the EU’s annual needs for extraction must come from the EU, 40% for processing and 25% for recycling. See https://commission.europa.eu/topics/competitiveness/green-deal-industrial-plan/european-critical-raw-materials-act_en

have idiosyncratic priorities. The fact that CTIPs are legally non-binding memoranda of understandings (MoUs) may be motivated in part by the fact this feature removes the need for European Parliament deliberation and approval (Jütten, 2025).

While bilateral agreements can address the objectives of both a supplier of critical materials and the country seeking to diversify its sourcing by including, e.g., investment financing for specific projects or commitments relating to upgrading/greater processing activity in the source country, there is a recognition that they may be second best instruments. This is because critical materials supply chains generally span more than two jurisdictions, necessitating many bilateral agreements that are likely to be less efficient and effective than cooperation among the major players in different value chains. Bilateral agreements can become steppingstones for broader engagement on a plurilateral basis. For example, the EU and the US have agreed on an Action Plan for Critical Minerals Supply Chain Resilience, with the aim of paving the way towards a possible plurilateral trade initiative with global partners by exploring how trade and other policy instruments can be used to support efforts to diversify critical mineral sourcing.³¹

V. Conclusion

Risk-matrix assessments of criticality are a useful starting point, but they are only inputs into a broader political-economy problem faced by multi-objective governments. When these governments act jointly, the interaction of domain-specific value functions, domestic weighting of objectives, and club-like externalities over membership and depth produces a rich variety of institutional outcomes. Understanding this variety requires moving beyond single-objective models of policy choice and focusing instead on how governments trade off multiple, often conflicting, goals when confronting shared risks.

Taken together, the MSP-to-FORGE transition and the parallel emergence of Pax Silica provide a clear illustration. Risk-matrix assessments of criticality did not move in any way that would justify recharacterizing rare earths, cobalt, or lithium between 2022 and 2026. What moved was

³¹ The Action Plan mentions several possible instruments, including border-adjusted price floors, standards-based markets, price gap-subsidies and offtake agreements, development of common standards for mining, processing and recycling; investment promotion; joint research and innovation; stockpiling strategies; and mechanisms for rapid response to supply disruptions. See https://ec.europa.eu/commission/presscorner/detail/en/ip_26_862. See also <https://www.csis.org/analysis/critical-minerals-ministerial-introduces-new-international-cooperation-strategy>.

the political weighting of domains within at least one key decision-maker, and the resulting reallocation across a landscape of club configurations whose existence the model anticipates. Understanding why critical-materials cooperation has fragmented into FORGE, Pax Silica, and a dense layer of bilateral deals — rather than consolidating into a single comprehensive arrangement — requires exactly the move away from single-objective criticality scoring that we have argued for.

This example also highlights the membership and depth issues discussed in the previous subsection. Current membership of plurilateral critical minerals initiatives is concentrated among advanced democracies with broadly similar objectives, especially a shared commitment to democratic capitalism, environmental protection, and alliance-based security cooperation. For these governments, the weights in the objective function $G(\cdot)$ attached to environmental and security domains are relatively high, and the domestic political environment is generally supportive of policies that combine greening and security narratives. This configuration makes it easier to agree on a club design in which strong environmental and governance standards are integral to the security-of-supply strategy. At the same time, even within this group, there is variation in industrial structure, existing mining capacity, and domestic political constraints, leading to ongoing negotiation about how far to go in embedding binding disciplines on export policies, local-content rules, and the allocation of new processing capacity across members.

From the perspective of non-members, a critical minerals club may look quite different. Governments that place greater weight on traditional industrial-policy autonomy or that are less willing to accept high environmental standards as a condition for access to finance may view the club's depth vector (δ_q, δ_c) as unattractive. They may prefer looser forms of cooperation, bilateral arrangements, or even strategies that deliberately exploit their role as non-member suppliers to extract rents or political concessions. Our framework suggests that such heterogeneous responses are not anomalies but rather the natural outcome of multi-objective decision-making under risk: the same plurilateral that increases the objective function $G(\cdot)$ for one set of countries may reduce it for others, even if the underlying physical properties of the materials and the basic structure of global supply chains are common.

Countries evaluating whether to participate in a plurilateral club will need to consider not only their own strategic objectives and goals but also the incentives confronting (potential) partner

countries. A necessary condition for success is to design critical materials clubs so as to both reduce the probability and potential adverse consequences of supply shocks (including susceptibility to weaponization of trade). How this might be done for specific minerals or specific economic activities – e.g., resource extraction, midstream processing, recycling, etc. – through the use of policy instruments is a core design challenge for the club. This calls for an understanding of the factors that will influence investment and production allocation decisions and willingness to accept the objectives and goals of countries hosting extraction and processing activities. For example, the latter may seek to move up the value chain. A key need will be for a club to identify and apply measures that align the incentives of upstream and downstream firms in specific supply chains,³² as well as the interests of these economic actors with those of the participating governments on both the supply and the demand side.

References

- Achzet, Benjamin, and Christoph Helbig. 2013. How to evaluate raw material supply risks—an overview, *Resources Policy*, 38: 435-47.
- Ahn, Dukgeun. 2026. Korea's Trade Policy Strategies through Plurilateral Initiatives: Middle-Power Adaptation under Geoeconomic Fragmentation. PIANO project paper, EUI.
- Aumann, Robert. 1959. Acceptable points in General Cooperative n-Person Games, in Albert Tucker and Robert Luce (eds.), *Contributions to the Theory of Games* Princeton University Press: Princeton): 287-324.
- Bao, Chunbing, Jianping Li, and Dengsheng Wu. 2022. *Risk Matrix: Rating Scheme Design and Risk Aggregation* (Springer Nature: Singapore).
- Baskaran, Gracelin, and Meredith Schwartz. 2026. Critical Minerals Ministerial Introduces New International Cooperation Strategy. In *CSIS Critical Questions*. Washington, DC: Center for Strategic & International Studies.
- Bernheim, Douglas, Bezalel Peleg and Michael Whinston. 1987. Coalition-Proof Nash Equilibria I: Concepts, *Journal of Economic Theory*, 42: 1-12.
- Bjørnsen, Kjartan, and Terje Aven. 2019. Risk aggregation: What does it really mean?, *Reliability Engineering & System Safety*, 191: 106524.
- Blackorby, Charles, Daniel Primont, and R. Robert Russell. 1978. *Duality, separability, and functional structure: theory and economic applications* (North-Holland: New York).
- Cox, Louis A. 2009. *Risk analysis of complex and uncertain systems* (Springer: New York).
- Depraeter, Lisa, and Stephane Goutte. 2023. The role and challenges of rare earths in the energy transition, *Resources Policy*, 86: 104137.
- European Commission. 2014. *Report on critical raw materials for the EU*. Enterprise and Industry Directorate General.

³² Upstream miners are likely to demand price and offtake guarantees, whereas downstream firms will be concerned about the associated consequences for their production costs. The latter may seek subsidies and/or trade protection to compensate them for such higher costs.

- European Commission. 2025. The Clean Industrial Deal: A joint roadmap for competitiveness and decarbonisation. COM(2025) 85 final.
- Glöser, Simon, Luis Espinoza, Carsten Gandenberger, and Martin Faulstich. 2015. Raw material criticality in the context of classical risk assessment, *Resources Policy*, 44: 35-46.
- Graedel, T. E., Rachel Barr, Chelsea Chandler, Thomas Chase, Joanne Choi, Lee Christoffersen, Elizabeth Friedlander, Claire Henly, Christine Jun, Nedal T. Nassar, Daniel Schechner, Simon Warren, Man-yu Yang, and Charles Zhu. 2012. Methodology of Metal Criticality Determination, *Environmental Science & Technology*, 46: 1063-70.
- Graedel, T. , E. Harper, N. Nassar, P. Nuss and B. Reck. 2015. Criticality of metals and metalloids, *Proceedings of the National Academy of Sciences*, 112: 4257-62.
- Haas, Peter M. 2021. Epistemic Communities, in Lavanya Rajamani and Jacqueline Peel (eds.), *The Oxford Handbook of International Environmental Law* (Oxford University Press: New York): 698-705.
- Helbig, Christoph, Lars Wietschel, Andrea Thorenz, and Axel Tuma. 2016. How to evaluate raw material vulnerability - An overview' *Resources Policy*, 48: 13-24.
- Hoekman, Bernard and Charles Sabel. 2019. Open plurilateral agreements, international regulatory cooperation and the WTO. *Global Policy*, 10(3), 297-312.
- Hoekman, Bernard and Charles Sabel. 2021. Plurilateral cooperation as an alternative to trade agreements: Innovating one domain at a time. *Global Policy*, 12, 49-60.
- Hoekman, Bernard, Petros C. Mavroidis and Douglas Nelson. 2023. Geopolitical competition, globalisation and WTO reform. *The World Economy*, 46(5), 1163-88.
- Jütten, Marc. 2025. Clean trade and investment partnerships A new instrument in the EU's trade policy toolbox. European Parliamentary Research Service (PE 769.576).
- Keeney, Ralph and Howard Raiffa. 1993. *Decisions with multiple objectives: preferences and value tradeoffs*. Cambridge University Press.
- Kennedy, James C. 2016. Rare Earth Production, Regulatory USA/International Constraints and Chinese Dominance: The Economic Viability Is Bounded by Geochemistry and Value Chain Integration. in Ismar Borges De Lima and Walter Leal Filho (eds.), *Rare Earths Industry*. Elsevier: Boston, 37-55.
- Krass, D., S. P. Sethi, and G. Sorger. 1994. Some Complexity Issues In A Class Of Knapsack Problems: What Makes A Knapsack Problem “Hard”?, *INFOR: Information Systems and Operational Research*, 32: 149-61.
- Leal Filho, Walter, Richard Kotter, Pinar Gökçin Özuyar, Ismaila Rimi Abubakar, João Henrique Paulino Pires Eustachio, and Newton R. Matandirotya. 2023. Understanding Rare Earth Elements as Critical Raw Materials, *Sustainability*, 15: 1919.
- Mas-Colell, Andreu, Michael Whinston, and Jerry Green. 1995. *Microeconomic Theory* (Oxford University Press: New York).
- McKelvey, Richard D. 1979. General Conditions for Global Intransitivities in Formal Voting Models, *Econometrica*, 47: 1085-112.
- Miller, Hugh, and Pau Morandi. 2026. Collective security in a fragmented world: avoiding the scramble for critical minerals through coordinated stockpiling. *CETEX Policy Report* London School of Economics, June.
- National Research Council. 2008. *Minerals, critical minerals, and the U.S. economy* (National Academies Press: Washington, D.C.).
- Neary, C. R., and D. E. Highley. 1984. The Economic Importance of the Rare Earth Elements, in P. Henderson (ed.), *Developments in Geochemistry* (Elsevier): 423-66.

- Pisinger, David, and Paolo Toth. 1998. Knapsack Problems. in Ding-Zhu Du and Panos M. Pardalos (eds.), *Handbook of Combinatorial Optimization: Volume 1–3* (Springer US: Boston, MA): 299-428.
- Plott, Charles. 1967. A Notion of Equilibrium and its Possibility under Majority Rule, *American Economic Review*, 57: 787-806.
- Raymond, Robert. 1984. *Out of the fiery furnace: the impact of metals on the history of mankind* (The Pennsylvania State University Press: University Park, PA).
- Riesman, David J., and Robert J. Weber. 2012. *Rare Earth Elements: A Review of Production, Processing, Recycling, and Associated Environmental Issues* (U.S. Environmental Protection Agency: Washington, D.C.).
- Szczepański, Marcin. 2025. China's rare-earth export restrictions. European Parliamentary Research Service (PE 779.220).
- Taboga, Marco. 2024. The potential macroeconomic relevance of critical materials: some preliminary evidence. *Bank of Italy Occasional Paper* #897.
- US Department of State. 2026. Fact Sheet: 2026 Critical Minerals Ministerial. February 4.
- Yu, Po-Lung. 1985. *Multiple-criteria decision making: concepts, techniques, and extensions* (Plenum Press: New York).