

Designing robust global value chains under trade-policy uncertainty

A regret-based network design framework

SEN2331

Silvester den Boer

Delft University of Technology

Designing robust global value chains under trade-policy uncertainty

A regret-based network design framework

by

Silvester den Boer

Chair:	M.Y. Maknoon
1st Supervisor:	P.S.A. Stokkink
2nd Supervisor:	R. Stöllinger
Advisor :	S. Momen
External Supervisor:	B. Rijnders (KPMG)
Project Duration:	February, 2026 - October, 2026
Faculty:	Faculty of Technology, Policy and Management, Delft

Cover:	An AI generated image
Style:	TU Delft Report Style, with modifications by Daan Zwaneveld

Preface

With this thesis I complete my MSc in Complex Systems Engineering and Management (CoSEM) at Delft University of Technology. I came to this topic through my interest in the workings of global supply chains and in network dynamics. Tariffs have been around for centuries, and recent geopolitical tensions have brought them back into mainstream attention. The uncertainty around them is a fairly new subject for supply chain network design. That made the work interesting, and it also meant I ran into many challenges. Throughout, I was most intrigued by how decisions that governments take, not always rationally, turn into trade barriers and ripple through value chains.

Trade measures change within months, while the production networks they tax take years to build, so a firm has to commit to its network before it knows which measures it will face. I found that a few well-placed alternative locations can limit its exposure to those measures, and that some exposure is worth keeping. I am happy with the result, and I believe the foundation of the model offers a new modelling perspective on a problem that is itself new. The topic holds much more potential for research, and I see this work as a good starting point. Writing this thesis also helped me grow personally. It taught me to be strict with myself and to be certain of my case before putting it on paper, and it showed me that today's tools can slow you down as much as they speed you up.

First of all, I thank the chair of my committee, Yousef Maknoon, with whom I selected this topic. I enjoyed working with him, and he challenged me to get to the bottom of the problem. I thank my first supervisor, Patrick Stokkink, for being a helping hand in the beginning on the core of robust optimisation. I also thank my second supervisor, Roman Stöllinger, for offering an extra view from the trade economics side of the story and for giving me more insight into the functioning of the WTO. Lastly, I especially thank Sara Momen, who joined at a later stage but contributed a great deal, for helping me with the mathematical foundation of the model and for taking the time to go through every detail with me until the very end. More than once I needed flexibility from them, and I am grateful they gave it.

I wrote this thesis in collaboration with the Procurement & Supply Chain team of KPMG Netherlands, which let me test how academic research translates into practice. I thank the whole team for their interest, for being a sounding board for my ideas, for a good look into consultancy, and for the laughs and good times at the office. Special thanks go to my external supervisor, Bob Rijnders, for his genuine interest and help, for sometimes simply being a listening ear and for never holding back his opinion. I also thank Erik van Zeijl for his sharp questions and his dedication to turning this research into insights a business can use.

With this thesis I also close seven years of my academic life, first in Groningen and then in Delft. In those years I lived through the highest highs and the lowest lows, and I would not have made it without the constant support of my parents. My final thanks go to them.

Cannot wait for all that is to come.

*Silvester den Boer
Delft, October 2026*

AI statement

This thesis was developed with ChatGPT 5.6 (OpenAI), Claude Opus 5 and Claude Code (Anthropic), and NotebookLM (Google). The statement follows the disclosure guidance of Cleland et al. (2026) and the advice on AI-assisted coding of Bridgeford et al. (2026). All data used in this thesis are public, and no confidential or commercially restricted information was given to an AI system.

Research and literature ChatGPT and Claude were used as discussion partners to test assumptions, explain optimisation concepts, critique arguments and suggest alternative formulations and experiments. Every choice among these suggestions was the author's, as were the research question, the regime framework, the formulation, the robustness rule and the experimental design. NotebookLM was used to find passages in the collected literature and policy documents, and every citation was checked against its source.

Code and experiments Most of the code in this thesis was written by Claude Code. The mathematical formulation, the experiments and the expected behaviour were specified first, Claude Code wrote and revised the implementation, and the numbers of each outcome were then checked by hand. When they matched, the code behind that outcome was taken to work, and when they did not, the implementation was rejected and specified again. The code is therefore verified through its outcomes rather than line by line.

This way of working made a scale of experimentation possible that could not have been reached by hand. The main grid alone covers 84 instances and more than 2,000 solves, the targeted experiments vary supplier origins and supplier dependence, and the solar case adds another 252 solves. Claude Code also went through the output at that scale, checking the termination status of every solve behind a reported result and evaluating every committed design under every scenario. Several experiments, such as varying the length of the production chain, were run and later left out because they did not change the mechanisms.

Reflection The ease of running experiments had a cost. With this much output, it was at times hard to keep an overview, to decide which results mattered and to keep the text in line with the latest runs. Several rounds of revision went into bringing the chapters back to the numbers and to the few mechanisms they support. The main lesson of working with these tools is that they make results easy to produce, and the real work becomes deciding which of them to trust and which to report.

Writing and figures ChatGPT and Claude were used extensively to improve structure, clarity and style, and parts of the text were drafted or substantially rewritten with their help. Several TikZ and pgfplots figures were generated from data supplied by the author. All text and figures were reviewed and revised by the author.

Responsibility Full responsibility for the research question, the model, the interpretation of the results and the conclusions of this thesis rests with the author.

Summary

Trade measures change within months, while the production networks they tax take years to build, so a firm commits its network before it knows which measures it will face. Duty falls on the full customs value at every border a good crosses, so where each production step sits decides how much a measure can tax.

Trade economics prices duty on the full value of goods at each border they cross, but takes the production network as given. Network design chooses locations, but treats tariffs as a known cost. This thesis asks how a firm should design a multi-stage global value chain so that it remains cost-efficient under normal trading conditions and limits its exposure to trade policy it cannot foresee.

Its main contribution is a two-stage robust optimisation model that needs no probabilities for political decisions. The firm first commits the locations where each production step can run, and once a measure is known it adjusts production and shipments within those locations. Six regimes represent future policy. Each is a documented class of measure, such as trade defence, safeguards or export controls. It has a rule for which shipments it reaches and a cap, the increment it adds in full to their tariff. The model accepts a design only if, under every regime, its worst-case cost stays within a chosen percentage of the best the firm could have secured had it known that regime. Among the accepted designs, it commits the one that is cheapest in normal times. This adapts the p -robustness of Snyder and Daskin to trade policy, and tightening the percentage traces what each level of protection costs. To the author's knowledge, no earlier model combines production across successive stages, uncertainty specific to each trade measure and a tariff base that follows the design.

Applied to three constructed chain architectures and to a solar photovoltaic chain serving the United States, the model yields three findings. First, a firm can protect itself without forecasting tariff rates, because what a measure reaches matters more than its rate. Defending against bilateral escalation costs up to 12 percent more than the cheapest network under normal tariffs. Defending against anti-circumvention, at the same cap, costs nothing extra. Second, a few alternative locations deliver most of the protection. A design's worst shock cost is the most that any regime adds to its cost. A second location per step achieves 91 to 94 percent of the cut in worst shock cost that allowing every location achieves. The most protected designs cost about 4 to 10 percent more in normal times, 6.6 percent in solar, and each unit of that extra cost removes about ten units of worst shock cost. Third, protection stops where the remaining exposure is cheaper to keep than to remove. The most protected solar design cuts the worst shock cost from 189.3 to 5.5 MUSD per GW. It keeps wafer and cell production in China, because leaving China would raise its extra cost in normal times from 6.6 to 9.2 percent.

A firm should hold its alternatives where its bill of materials forces goods across a border, at merges that join separately made components, at the closing step and where supplied inputs arrive. These steps account for 155 of the 164 changes the model makes as protection tightens. Because value jumps at each merge, moving these steps towards the market leaves the crossing on cheaper goods upstream. This lowers the value a tariff can reach and protects against regimes on movements into a market. Conditions outside the production network, where the firm sells and where it buys, decide how much protection it needs.

The findings are mechanisms, and the numbers show how large they are in these chains. The numbers are less certain, because the constructed chains rest on declared cost shares, the solar case on public cost data, and the caps, commitment costs and minimum output per facility are not systematically varied. The solar case demonstrates that the mechanisms hold in a real world instance. The next steps are to vary these inputs and to apply the framework to one firm with its own data.

Contents

Preface	i
AI statement	ii
Summary	iii
1 Introduction	1
1.1 Research gap and questions	1
1.2 Approach	2
1.3 Contributions and findings	3
2 Literature review	5
2.1 Fragmented production and the gross-value tariff base	5
2.1.1 Value creation along global value chains	5
2.2 Trade rules and policy shocks	6
2.2.1 Trade rules and the measures they permit	6
2.2.2 Low and predictable tariffs	7
2.2.3 Trade measures before 2018	7
2.2.4 The post-2018 escalation	8
2.2.5 What trade measures cost firms	11
2.3 Network design under trade-policy uncertainty	11
2.3.1 Tariffs in network design models	12
2.4 Two literatures on tariffs	12
2.5 Research gap	13
3 Problem and research questions	15
3.1 Main question	16
3.2 Sub-questions	16
3.3 Scope	16
4 Trade-policy regimes	18
4.1 From measures to regimes	18
4.2 The six regimes	19
4.2.1 Exposure rules	20
4.2.2 Alignment	21
4.2.3 Regime caps	21
4.3 Bundles	23
5 Methodology	24
5.1 Conceptual model	24
5.2 Mathematical formulation	24
5.2.1 Sets and network structure	24
5.2.2 Parameters	24
5.2.3 Decision variables	26
5.2.4 Objective function	26
5.2.5 Constraints	27
6 Robust optimisation under trade-policy uncertainty	29
6.1 Overview	29
6.2 The uncertainty set	30
6.3 The model	31
6.3.1 The reference	31

6.3.2	The robust design problem	32
6.4	Solution procedure and experimental measures	33
6.4.1	Robustness frontier and floor	33
7	Computational experiments	35
7.1	Experimental design	35
7.1.1	Reference architectures	35
7.1.2	Geographic setting and demand	36
7.1.3	Economic parameterisation	37
7.1.4	Location rights, minimum throughput and capacity	40
7.1.5	Regime instantiation	40
7.1.6	Experimental grid	41
7.1.7	Implementation	41
7.1.8	Reading the results	41
7.2	The robustness frontier	42
7.2.1	Reference networks	42
7.2.2	Design changes along the frontier	44
7.2.3	Binding regimes along the frontier	45
7.2.4	Limits of maximum-regret protection	47
7.2.5	Break-even caps	49
7.3	Regime effects	49
7.3.1	Operational response to each regime	49
7.3.2	Premium and protection along the frontier	51
7.3.3	Responses by regime type	52
7.4	Conditions shaping robustness	53
7.4.1	Supplier origins and input dependence	53
7.4.2	Production-location flexibility	56
7.4.3	Demand geography	57
7.5	Managerial implications	59
8	Case study: the solar photovoltaic value chain	61
8.1	Case framework	61
8.1.1	Case setting	61
8.1.2	Production structure	62
8.1.3	Economic parameterisation	63
8.1.4	Regime instantiation	65
8.1.5	Experiments	66
8.2	Results	67
8.2.1	Reference networks	67
8.2.2	Design changes along the frontier	68
8.2.3	Break-even caps	70
8.2.4	Solar value chain strategies	70
8.2.5	The solar floor design under each regime	72
8.2.6	Leaving China and the reference firms	73
8.3	Case synthesis	74
9	Discussion and conclusion	76
9.1	Answers to the research questions	76
9.2	Contribution and relevance	80
9.3	Limitations	81
9.4	Further research	82
	Bibliography	83
A	Computational detail	89
A.1	Instance calibration	89
A.2	Full design frontiers	90
A.3	Regime flows by architecture	93
A.4	The robustness floors in full	96

A.5	Supplier dependence in full	97
A.6	Cost under shock along the frontier	98
B	Solar PV instance calibration	101
B.1	Network	101
B.2	Gross values	101
B.3	Cost calibration	101
B.4	Baseline tariff environment	103
B.5	Source exhibits	104
C	Implementation and reproducibility	106
C.1	Solution procedure	106
C.2	Solver record	106

1

Introduction

On 2 April 2025 the United States announced tariffs that raised the combined rate on imports from China to 54 percent, set 20 percent for the European Union and 10 percent for most remaining trading partners, and within eight weeks revisions, suspensions for some partners and challenges in court had followed (Ignatenko et al., 2025). The announcement extended a confrontation that started in 2018, when additional tariffs came to cover more than \$300 billion of United States imports (Fajgelbaum et al., 2020). Each measure was discretionary, aimed at chosen products and often at chosen partners, and moved far faster than the networks it taxed (Ossa and Redding, 2026).

That asymmetry creates the strategic problem. A production network takes years to establish, so firms operate networks committed for one policy environment while facing another. The question is therefore which network to commit before the firm knows the policy it will face, since any response to a measure works within that network.

Fragmented manufacturing amplifies the problem. Intermediate goods cross several borders before reaching the customer (Hummels et al., 2001), and an ad valorem tariff falls on the customs value presented at each crossing rather than on the value added since the previous one (World Trade Organization, 1994a), so later crossings expose progressively larger values to the same rate (Yi, 2003). Where the firm locates a step decides which commodity crosses which border and at what value, so for a firm, tariff exposure is a consequence of network design rather than an external parameter.

1.1. Research gap and questions

Existing research captures one side at a time. Trade economics prices duty on the gross value of goods in process and documents how trade policy behaves, but treats the production network as given (Yi, 2003; Rouzet and Miroudot, 2013). Supply chain network design treats location as a decision, but prices tariffs as a cost on flows, without representing which measure produced them or what it can reach (Melo et al., 2009; Dong and Kouvelis, 2020). Neither combines a tariff base that follows the design with policy the firm cannot know when it commits. Therefore the main research question of this thesis is:

How should a firm design a multi-stage global value chain so that it remains cost-efficient under normal trading conditions and limits its exposure to trade policy it cannot foresee?

This main question is split into four sub-questions, which Chapter 3 states in full. The first asks how the configuration of a fragmented value chain determines its exposure to trade measures, and which part of that exposure the network design can remove. The second asks how uncertainty about future trade policy can be represented for robust network design when no probabilities can be assigned to it. The third asks how that uncertainty changes the cost-minimising network and what determines the result. The fourth asks what a robust network costs under normal trading conditions and under the measures it may face, and where further protection stops paying.

1.2. Approach

The method in brief The firm must choose where to produce before it knows which trade measures it will face, and it can adjust production and shipments only within the locations it has committed. The method compares candidate networks in three steps. It first finds the network that is cheapest under the tariffs of normal times. It then takes each kind of trade measure governments have used, such as anti-dumping duties or export restrictions, and finds the network the firm would have built had it known that kind of measure was coming. Finally, it accepts a network only if, whichever kind of measure comes, its worst-case cost stays within a chosen percentage of that best network's. Among the accepted networks, it commits the one that is cheapest in normal times, and lowering the percentage shows what each extra degree of protection costs.

The model This thesis answers the question with a two-stage robust optimisation model. In the first stage, the firm commits a design by qualifying production steps at candidate locations, through owned capacity or external producers. In the second stage, once a measure is known, production and flows adjust within that design. Each qualified facility carries a per-period commitment cost and, in the applications, must produce at least 5 percent of demand whatever measure comes.

Normal trading conditions Under normal trading conditions each crossing pays its baseline tariff, the bilateral effectively applied rate for that commodity and country pair. The design with the lowest total cost under these rates is the deterministic network. It provides the benchmark for what protection costs in normal times.

Uncertainty about trade policy The source of uncertainty is the trade policy the firm will face after committing its design. A government decides which goods and countries a measure targets and how high to set its rate, so the firm faces a double unknown, which measure will come and how severe it will be. No record gives the frequency of such discretionary measures, so the thesis represents this uncertainty without probabilities.

Regimes and caps The thesis constructs six regimes from documented classes of measure, each defined by the movements it can reach, where a movement is one commodity shipped between two locations. Bilateral escalation reaches all goods moving between two countries. Four regimes act on goods entering a given market. Friendshoring reaches all goods from non-allied origins, and trade defence reaches the finished good from one non-allied origin. Market protection reaches the finished good from every foreign origin, and anti-circumvention reaches the finished good or its direct input from a neutral hub. Export control reaches a supplied input leaving its origin. A seventh regime, the composite, combines trade defence and anti-circumvention against the same market.

Each regime works as a rule. Applied to the candidate network, which holds every location the firm could use, it generates every measure of its class that is possible within that network. Each such measure is one scenario, and a scenario names the movements its measure reaches. Trade defence, for example, generates one scenario for every non-allied origin that could ship the finished good into a market. Bilateral escalation, by contrast, generates one scenario for every pair of countries in the network. A regime's cap is the fixed tariff increment it adds in full to the baseline tariff of every movement it reaches (Equation (6.3)). A cap of 0.54 adds 54 percentage points. The caps, from 0.20 to 1.00, rest on recorded measures and published estimates (Chapter 4). Because each cap is set before optimisation, the model enumerates the first unknown and fixes the second. Within the model, the uncertainty therefore lies only in which regime comes and which of its scenarios occurs.

Regret and the robustness frontier Under a regime, a design's worst-case cost is its total cost in that regime's costliest scenario, after production and flows adjust. For each regime, the model first computes the reference cost, the lowest worst-case cost any design within the candidate network can attain. The design that attains it, the reference network, is the one the firm would choose if it knew the regime but not its scenario beforehand. A design's regret is how far its worst-case cost exceeds the reference cost, as a share of the reference cost. Regret thus counts commitment and operating costs as well as duty.

The firm chooses a regret tolerance, θ , the largest regret it is prepared to accept, and the model selects the design that is cheapest in normal times while keeping its regret within that tolerance under every regime. Whatever regime comes, the committed design's worst case then costs at most a share θ more than the worst case of the reference network, the design the firm would have chosen had it known that regime. This is the p -robust rule of Snyder and Daskin (2006), adapted to regimes and to cost in normal times. A design's premium is its extra normal-times cost relative to the deterministic network. Tightening the tolerance traces a frontier between premium and largest regret across regimes, from the deterministic network to the floor design. The floor design's largest regret, the robustness floor, is the lowest any design can attain. The frontier therefore shows what each attainable level of protection costs.

Scope The model compares alternative commitments over a single period, with fixed demand and input requirements. Measures enter as tariffs or their ad valorem equivalents, and outright route closures lie outside its scope. The comparison includes per-period commitment and operating costs, but does not model investment timing, financing or recovery over time.

Applications The thesis applies the framework to three constructed chain architectures, Serial, Modular and Complex, and to a crystalline silicon photovoltaic chain serving the United States. The constructed chains can produce in four blocs, a high-cost market and, relative to it, an allied bloc, a neutral hub and a low-cost adversarial bloc. Architecture and demand geography give 21 instances, three of which, one per architecture, serve only the high-cost market and are reported in most detail. Further runs vary supplier origins and location options. Demand geography and supplier origins lie outside the production network, so these runs also test how conditions the firm inherits, where it sells and where its suppliers are, shape how robust its network can be. Varying one property at a time isolates the mechanisms (Chapter 7).

The photovoltaic case examines those mechanisms on documented locations and costs (Chapter 8). A firm buys polysilicon and makes wafers, cells and modules for the United States, at locations in the United States, allied Korea, neutral ASEAN hubs and adversarial China where Trina Solar and Hanwha Qcells document capability, with 2022 IEA costs and baseline tariffs from WITS. Two controls restrict supply to China and remove Chinese production, and the committed designs are set beside the footprints of both firms. Both applications trace the frontier and evaluate committed designs under every regime. They support conclusions about the modelled mechanisms. The industry case is a demonstration rather than a validation against a firm's observed decisions.

1.3. Contributions and findings

The thesis makes four contributions, which Chapter 9 positions against the literature.

1. A rule for choosing a network without probabilities. It accepts only designs whose regret stays within the tolerance under every regime and commits the one among them that is cheapest in normal times. The rule adapts the p -robustness of Snyder and Daskin (2006), which compares a design with the best design for each scenario and minimises expected cost. Here the comparison is with the best design for each regime, because the firm cannot know which origin a measure will name, and normal-times cost replaces expected cost, because no probabilities exist.
2. A way to represent trade-policy uncertainty that follows trade law. The closest earlier robust model lets tariff rates vary within a range (Ahmadi Dighehsara et al., 2025). Here each regime is a class of measure that governments actually use, defined by the goods and origins its rules allow it to target, so every scenario is a measure the firm could face and the set can be checked against the record.
3. Three findings on how a firm can protect itself. It can do so without forecasting tariff rates, a few alternative locations deliver most of the protection, and protection stops paying where the remaining exposure is cheaper to keep than to remove. The findings below state each with its evidence.
4. A framework that joins trade economics and network design, the two literatures that so far treat tariffs separately. From the first it takes the duty on the full value of goods at every crossing, and

from the second the firm's choice of where to place each step, so that the tariff base follows the design. Any multi-stage value chain within its scope enters as data, and the framework returns what each level of protection costs for that chain.

The thesis argues that a firm can limit its exposure to trade policy it cannot foresee by holding a few well-placed alternative locations, and that some exposure is worth keeping. Three findings support this.

First, a firm can protect itself without forecasting tariff rates, because which of its shipments a measure reaches matters more than the rate it charges. Bilateral escalation against a non-allied country and anti-circumvention both add 54 percentage points to the tariff. Yet defending against bilateral escalation costs up to 12 percent of the deterministic network's cost, while defending against anti-circumvention costs nothing, because the deterministic networks produce nothing in the neutral hub it reaches. Defending against a regime also pays quickly. If the firm knew which regime would come, defending against one that reaches the finished good would cost less than keeping the deterministic network even if that regime charged only a fifth of its cap.

Second, a few alternative locations deliver most of the protection. A design's worst shock cost is the most that any regime adds to its cost. In the constructed chains, allowing each step a second location achieves 91 to 94 percent of the cut in worst shock cost that allowing every location achieves. Conditions outside the network decide how much of it is needed. Where the firm sells sets how much exposure enters, and where its suppliers are accounts for much of what remains.

Third, protection stops where the remaining exposure is cheaper to keep than to remove. In the solar chain the floor design carries a premium of 6.6 percent and cuts the worst shock cost from the deterministic network's 189.3 to 5.5 MUSD per GW, yet keeps wafer and cell production in China, since leaving China would add 2.6 percentage points to the premium.

Chapter 2 reviews the literature and Chapter 3 states the design problem and the four sub-questions. Chapter 4 constructs the regimes and their caps, Chapters 5 and 6 develop the deterministic and the robust model, Chapters 7 and 8 apply the framework to the constructed chains and the photovoltaic case, and Chapter 9 answers the questions and states the limitations.

2

Literature review

This chapter reviews the two literatures the thesis joins. Trade economics and trade law establish what duty falls on and how trade measures behave, and network design supplies a model for committing production before uncertainty resolves. The chapter compares the two and ends with the gap they leave.

2.1. Fragmented production and the gross-value tariff base

A tariff is the rate a government charges on goods entering its territory, and a duty is the amount the importer pays. An ad valorem tariff is a percentage of the good's customs value and a specific tariff a fixed amount per unit. Each importer sets its rates per product in a schedule built on the Harmonized System, the nomenclature of the World Customs Organization that gives each product group a six-digit code and that more than 200 countries use as the basis of their customs tariffs (World Customs Organization, n.d.). The WTO valuation rules base customs value first on the transaction value, the price actually paid or payable for the good when sold for export (World Trade Organization, 1994a).

The thesis speaks of trade-policy uncertainty because a firm cannot foresee which measure a government will take, which origins and products it will name and how much it will charge. It prices a measure by the cost that measure adds to goods crossing a border, so it works with ad valorem tariffs and brings in any other measure through its ad valorem equivalent, the ad valorem tariff that would raise the cost of importing the good by the same amount (Anderson and van Wincoop, 2004). A specific duty, a quota or licence, the withdrawal of a preference and a condition on subsidy eligibility can each enter this way, whatever their legal form. A measure that closes a route outright, such as a ban or detention at the border, and a charge on a base other than customs value lie outside the thesis. The exclusion is a modelling choice and does not rank those measures below the ones modelled.

Hummels et al. (2001) document the growth of vertical specialisation, in which exports embody imported intermediates and goods in process cross more than one border. An ad valorem tariff falls on the full customs value presented at each crossing, which includes the value added at every earlier step, so a step may add little value while its output carries a large base. Yi (2003) shows that repeated trade costs along sequential chains make trade respond disproportionately to tariff changes. Exposure therefore depends on the tariff rate, on how often the product crosses a border and on the value it carries at each crossing. A government sets the rate, while a firm that designs its production network sets the other two by where it places each step.

2.1.1. Value creation along global value chains

The smiling curve holds that activities at the start and end of a value chain, such as design and branding, capture more value than fabrication and assembly in the middle (Mudambi, 2008). In its original form the curve describes business functions rather than production stages, but Rungi and Del Prete (2017) find the same profile within production alone: tracing value added from primary goods through intermediates to final goods, they find it concentrated at the beginning and end of the production process. Gross value

accumulates along the chain, so this profile sets the customs value presented at each border crossing, and with it the tariff base.

Customs value is the transaction price (World Trade Organization, 1994a), so it also carries the margin the seller captures, and margins differ substantially between positions in the same chain (Shin et al., 2012). What a step adds to a product's value therefore differs from what it spends on production and materials, and customs value does not rise in step with production cost. This is the one property of the curve that the constructed chains of Chapter 7 use.

2.2. Trade rules and policy shocks

Trade rules decide when a country may tax goods from one origin more heavily than goods from another, on which products and for how long. Governments use the measures these rules permit, and since 2018 they have increasingly acted outside them. Which goods and origins a measure covers decides what a firm can escape by moving production, and the rate it charges decides what the firm pays where it cannot.

2.2.1. Trade rules and the measures they permit

The World Trade Organization (WTO), established by the Marrakesh Agreement in 1995, is the organisation that oversees the multilateral trade agreements, with the aim of reducing barriers to trade and ending discrimination between trading partners (World Trade Organization, 1994b). The rules on tariffs for goods are in the General Agreement on Tariffs and Trade (GATT), which rests on two commitments. Each member lists a bound rate for each product in its schedule, a ceiling negotiated in the trade rounds, and under the most-favoured-nation (MFN) clause it must give the goods of all members the same treatment. The MFN clause matters most for the rates actually applied. A member may charge less than its bound rate, but any lower rate it applies to one member it must offer to all (World Trade Organization, n.d.b). Regional trade agreements are the main exception. The European Union's customs union and its free trade agreements, for example, give partners lower or zero rates that other members do not receive, and members may also grant developing countries preferential access. The rate an importer charges on a product from a given partner is therefore its bilateral effectively applied rate, the lower of its applied MFN rate and any preferential rate it grants that partner (World Bank, n.d.), and the trade measures below are charged on top of it.

Three WTO trade measures may depart from these commitments under set conditions. Anti-dumping and countervailing duties protect domestic producers from imports priced too low for them to compete with. An anti-dumping duty, permitted by the GATT and the Anti-Dumping Agreement, answers low prices set by the exporting firm, which sells abroad below its normal value, usually the price it charges at home or, where that cannot be used, its cost of production plus a normal profit. A countervailing duty, regulated by the Agreement on Subsidies and Countervailing Measures, answers low prices made possible by the exporting government, which subsidises its firms, and offsets that subsidy. Both require an investigation showing injury to the domestic industry and normally expire after five years.

A safeguard, permitted by the GATT and the Agreement on Safeguards, answers imports that rise too fast rather than imports priced unfairly. It temporarily restricts imports of a product when a surge causes or threatens serious injury to the domestic industry, to give that industry time to adjust. National authorities impose these measures, and a member that considers another member's measure inconsistent with the agreements can take it to WTO dispute settlement (World Trade Organization, n.d.a).

The three differ in reach. Anti-dumping and countervailing duties follow a product from named countries, so production moved to another origin escapes them. A safeguard applies to the product from every supplier, except developing countries that each supply at most 3 percent of imports and together at most 9 percent (Read, 2005). Blanchard et al. (2016) find that bilateral tariffs and WTO trade measures are lighter against partners whose final goods contain the imposing country's value added, because the duty then falls partly on the imposing country's own input suppliers. The regimes of trade defence and market protection follow this difference in reach (Chapter 4).

Governments also act outside the three WTO trade measures, on different legal footing. Withdrawing a preference returns a partner to the MFN rate, which the agreements permit. Export restrictions fall under separate rules, and a WTO ruling found China's rare-earth export quotas and duties inconsistent

with them (Mancheri, 2015). A government can close the escape that anti-dumping and countervailing duties leave open by extending the duty order to goods completed in a third country. The agreements contain no rule on such extensions (Zhu, 2024). Tariffs that governments impose under domestic law, such as the United States tariffs on national-security and unfair-trade grounds, became prominent after 2018 (Bown, 2021).

Two recent EU measures influence where production locates but fall outside the scope of this thesis, because neither is a border charge on customs value. The Carbon Border Adjustment Mechanism charges embodied emissions (European Commission, 2024a), and the Critical Raw Materials Act sets benchmarks for where the EU's supply comes from (European Commission, 2023).

2.2.2. Low and predictable tariffs

The post-war trading system, built through successive rounds of the General Agreement on Tariffs and Trade and consolidated under the WTO, produced a sustained decline in trade barriers, which took average effective tariffs in the United States below two percent by the mid-2010s (Ossa and Redding, 2026). Figure 2.1 shows the United States series, customs duty revenue as a share of merchandise imports, chosen because it runs from 1790 and because the post-2018 escalation began with United States measures. Falling tariffs and expanding global value chains reinforced one another, since fragmented production pays only while border-crossing costs stay low (Yi, 2003).

Figure 2: U.S. Average Effective Tariff Rate Since 1790




Figure 2.1: US average effective tariff rate since 1790. Source: Ossa and Redding (2026), Figure 2, with data from the Yale Budget Lab, Historical Statistics of the United States, the Monthly Treasury Statement and the Bureau of Economic Analysis.

Bound rates made these low tariffs predictable as well. A member can change a bound rate only after negotiating with its trading partners, which may mean compensating them for lost trade (World Trade Organization, n.d.b). Handley and Limão (2015) show that commitments which remove the risk of reversal to higher tariffs raise export investment even when applied tariffs are already low, and they cite evidence that the WTO's bound rates have this effect.

Firms could therefore treat tariffs as fixed parameters, and Ossa and Redding (2026) describe them as a largely settled issue until their return as an active policy tool after 2018. That stability was conditional, since bound rates deliver predictability only while the commitments behind them remain credible.

2.2.3. Trade measures before 2018

Bown and Crowley (2013) trace the use of the three trade measures by five industrialised economies from 1988 to 2010. Worldwide, WTO members opened 6,353 investigations under the three trade measures between 1995 and 2017 and imposed 4,033 measures, nine in ten of them anti-dumping measures. India opened the most anti-dumping investigations, 888, ahead of the United States with 661 and the European Union with 502 (World Trade Organization, 2025). Table 2.1 records five measures from before 2018, with the legal ground of each, how it ended and how firms and markets responded.

Table 2.1: Five trade measures before 2018.

Measure	Legal ground	Ruling or end	Observed response
US steel, 2002	Safeguard	Withdrawn after an adverse WTO ruling (21 months)	Higher costs for steel-consuming industries (Read, 2005)
US tyres from China, 2009	China-specific safeguard under China's accession terms	Upheld by the WTO, expired on schedule (3 years)	Imports shifted to other East Asian suppliers (Charnovitz and Hoekman, 2013)
US solar from China, 2012	Anti-dumping and countervailing duties	In force, extended to Southeast Asian completion through circumvention findings (U.S. Department of Commerce, International Trade Administration, 2023)	Assembly relocated to Southeast Asia (Zhu, 2024)
EU solar from China, 2013	Anti-dumping and countervailing duties	Expired September 2018 (European Commission, 2018)	Price undertaking covering 75 percent of Chinese imports (Goron, 2018)
China rare earths, 2010	Export quotas and taxes	Adverse WTO ruling 2014, abolished 2015 (Mancheri, 2015)	Downstream export effects differed by sector (Chen et al., 2021)

The five are a selection, and their only purpose is to show that governments used such measures before 2018 and that firms responded to them, by relocating production, changing sourcing or absorbing higher input costs. A selection establishes no regularity, and large industrialised economies or China imposed all five on partners able to retaliate or to challenge them at the WTO. Adverse WTO rulings or fixed terms ended three of the five, the European Union let its solar duties expire, and only the United States solar orders persisted, extended through circumvention findings.

2.2.4. The post-2018 escalation

The trade war between the United States and China serves here as the main example of a wider shift after 2018, towards larger and less predictable tariff increases imposed outside the WTO trade measures. Between January 2018 and January 2021, the average United States tariff on Chinese goods rose from 3.1 to 19.3 percent, and the Chinese rate on United States goods from 8.0 to 20.7 percent (Bown, 2021). By January 2021, additional tariffs covered 66.4 percent of United States imports from China and 58.3 percent of Chinese imports from the United States. Unlike the product-specific measures before 2018, the escalation reached most bilateral trade (Figure 2.2).

At product level, Fajgelbaum et al. (2020) report that United States import tariffs rose from 2.6 to 16.6 percent across 12,043 products covering \$303 billion of trade. Trading partners retaliated, raising average rates on United States exports from 7.3 to 20.4 percent across 8,073 products covering \$127 billion, and Figure 2.3 traces the United States increases and the retaliation that followed through 2018.

Two properties of the shock matter for network design. It was discrete, since 99.8 percent of targeted import varieties received an additional tariff of either 10 or 25 percentage points, and it fell mostly upstream, with intermediate and capital goods making up 85 percent of the value of targeted imports (Fajgelbaum et al., 2020). The model represents measures accordingly, as increases that arrive in one step and can reach intermediate goods as well as final ones.

The post-2018 tariffs also persisted. The tariffs between the United States and China remained largely in place after the Phase One agreement, with some extended to further products (Bown, 2021). The announcements of April 2025 carried the approach beyond that conflict with large country-specific increases (Ignatenko et al., 2025), many imposed under emergency economic powers after a declared national emergency (Ossa and Redding, 2026). The three WTO trade measures stayed in use alongside them, as in the European Union's 2024 countervailing duties of up to 35.3 percent on battery





Figure 2.2: US and Chinese tariff rates towards each other and the rest of the world, January 2018 to January 2021 (Panel A), and the share of bilateral trade subject to additional tariffs (Panel B). Source: Bown (2021).



Notes: Figure shows the unweighted average tariff rate of targeted import and export varieties for each tariff wave before and after they are targeted. Import tariffs constructed from U.S. International Trade Commission (USITC) documents, and retaliatory tariffs constructed using official documents from foreign finance and trade ministries.

Figure 2.3: Trade-war timeline: statutory tariff rates on targeted US imports (Panel A) and retaliatory tariffs on US exports (Panel B), January to December 2018. Source: Fajgelbaum et al. (2020).

electric vehicles from China (European Commission, 2024b), and governments attached origin and ownership conditions to subsidies, as in the United States restrictions on clean-energy credits for prohibited foreign entities (Internal Revenue Service, 2026).

Since 2018, measures have been large, discrete, specific to origins and products, and open-ended, so a firm now commits its network against the measures a government may reach for and the cost they add to the crossings they name.

2.2.5. What trade measures cost firms

Exporters may absorb part of a tariff by lowering their prices, and how much depends on the measure. Nizovtsev and Skiba (2020) find that exporters absorb part of ordinary tariffs but pass anti-dumping duties on more fully. The 2018 tariffs passed through almost completely into US import prices, so importing firms and their customers bore them (Amiti et al., 2019; Fajgelbaum et al., 2020). Fajgelbaum et al. (2020) estimate gross losses to consumers and importing firms of \$51 billion, an aggregate real income loss of \$7.2 billion net of tariff revenue and gains to protected producers. Those gains came with higher domestic prices and markups (Amiti et al., 2019), so firms buying tariffed inputs faced more expensive domestic alternatives as well.

Responding costs the firm as well. Miller et al. (2025) distinguish adjustment costs, such as relocating production and equipment, transaction costs, such as qualifying suppliers and renegotiating contracts, and the opportunity costs of responding early or late. The model prices only the per-period cost of holding a location qualified, so one-off relocation and qualification costs and the timing of a response lie outside it.

2.3. Network design under trade-policy uncertainty

Supply chain network design determines which facilities or production capabilities a firm establishes and how it operates the resulting network (Melo et al., 2009). The firm makes those commitments before it knows trade policy, so the model needs a representation of that uncertainty.

Stochastic optimisation needs probabilities that trade policy does not supply. This thesis follows instead the robustness perspective of Rosenhead et al. (1972), evaluating commitments across specified possible futures, so its designs depend on which policy realisations the uncertainty set includes.

Robust optimisation replaces probability distributions with an explicit uncertainty set, the states a solution must withstand. In the two-stage adjustable framework of Ben-Tal et al. (2004), operational decisions adapt once the uncertainty resolves. The set can be a bounded region, as in the budgeted uncertainty of Bertsimas and Sim (2004), or a finite list of scenarios, as in this thesis (Zeng and Zhao, 2013).

A second strand of the literature evaluates decisions by their regret rather than their worst-case cost. Savage (1951) measures a decision by its loss against the best action for the state that occurs, so the benchmark in each scenario is the optimum under perfect information (Madansky, 1960). Regret therefore judges a decision against what was attainable, and a cost no decision could have avoided counts against none of them.

Snyder and Daskin (2006) turn regret into a constraint. A solution's relative regret in a scenario is its cost in excess of that scenario's optimum, as a share of that optimum, and their stochastic p -robustness requires it to stay below a set tolerance in every scenario while the objective minimises expected cost. The tolerance sets the trade-off. With none imposed the model minimises expected cost alone, tightening it lowers the maximum regret at a higher expected cost, and the tightest tolerance that still admits a solution yields the minimax-regret solution. In their test instances large reductions in maximum regret cost little in expected terms. Their models are two-stage, with facility locations fixed before the scenario and customer assignments following it.

Bruckler et al. (2024) synthesise the supply chain resilience literature around three capacities. Absorptive capacity limits performance losses during a disruption, adaptive capacity adjusts operations to changing conditions, and restorative capacity recovers performance afterwards. They associate robustness with the first and resilience with all three. In a two-stage design model the committed network

provides absorptive capacity and the re-optimised flows provide adaptive capacity, so a cost measured after re-optimisation reflects both. Restorative capacity lies outside a single-period model.

The supply chain reliability literature models physical failures and asks how a network keeps operating when assets become unavailable (Snyder and Daskin, 2005). A measure that closes a route raises the same question. The measures this thesis represents leave every asset and route in place and change what a route costs, so a network stays able to operate but may cease to be worth operating in its original configuration.

2.3.1. Tariffs in network design models

Supply chain network design models already include tariffs, in decisions on facility location, sourcing, capacity and product mix (Ahmadi Digebsara et al., 2025). Most treat the tariff as a rate on a fixed base and ask how that rate changes the network. In Zhou et al. (2017) the base depends on the design, because a carbon border adjustment taxes embodied emissions at a known rate. Cui et al. (2025) plan capacity along a five-stage photovoltaic chain for 12 regions to 2030, where a known European ban on Chinese imports sets those flows to zero. Ahmadi Digebsara et al. (2025) make tariffs and import quotas uncertain in a two-stage robust design, with a single production stage and tariffs as uncertain cost coefficients on bilateral flows. Dong and Kouvelis (2020) develop stylised models in which tariffs change the configuration and operational flexibility of a global network. They also set out four directions for tariff research in network design, covering tariff uncertainty, product interdependence, competition and the strategic responses of partners in decentralised supply chains.

2.4. Two literatures on tariffs

Economic studies and network design models both ask how tariffs change the cost of production spread across borders. Miller et al. (2025) find far less supply chain research on the question than economic research, and no framework that connects the two. The two differ in what they take as given. Economic studies take the production structure as given, observed in input-output data or set by market equilibrium, and measure outcomes for markets, industries and countries (Yi, 2003; Rouzet and Miroudot, 2013; Blanchard et al., 2026). Network design models take tariff rates as given and optimise one firm's facilities, capacities and flows (Melo et al., 2009; Dong and Kouvelis, 2020). Handley and Limão (2015) attach probabilities to future tariffs, while network design models represent uncertainty through scenarios or uncertainty sets (Snyder and Daskin, 2006; Ahmadi Digebsara et al., 2025).

Each is strong where the other is thin. Economic studies measure how much of a tariff reaches import prices (Amity et al., 2019; Fajgelbaum et al., 2020) and keep prices and wages consistent in general equilibrium (Yi, 2003), but with no firm-level choice of locations they cannot tell a firm where to produce. Network design models are prescriptive and can also represent production steps and the timing of decisions (Cui et al., 2025), but they price tariffs as a cost on flows (Ahmadi Digebsara et al., 2025). Dong and Kouvelis (2020) add that their own models fix finished-goods prices and leave out input substitution and changes in product mix.

This thesis follows network design, because its decision has the two stages that tradition models. A manufacturer commits production capabilities before it knows trade policy, and production and flows adapt once it does, the here-and-now and wait-and-see decisions of Shapiro et al. (2009). Where network design speaks of a supply chain, the thesis speaks of a *value chain*, because the duty a design pays follows the value each crossing carries.

From economics and trade law it takes four elements. The first is the tariff base: a tariff falls on the full value of the good in process at each crossing (Yi, 2003), and the model prices duty on that value. The second is reach: measures apply to defined sets of origins and products (World Trade Organization, n.d.a), and Chapter 4 defines each class of measure by what it reaches. The third is one-sided risk: only the possibility of higher tariffs affects a sunk commitment (Handley and Limão, 2015), so measures in the model only add cost and carry no probabilities. The fourth is incidence: the 2018 duties passed almost fully into import prices, so duty counts as a cost to the importing firm, an upper bound where exporters absorb part of a tariff (Nizovtsev and Skiba, 2020). It leaves out the general-equilibrium adjustment of prices and wages and the pricing and competition responses some network models carry, so its results describe one firm's choices at given prices and demand.

Table 2.2 positions the studies that meet at least one of the two requirements this review has set out: a tariff base that follows where a firm locates its steps, and policy unknown when it commits them. The column headings give a common basis for comparison and do not follow the terminology of the original studies.

2.5. Research gap

No study reviewed combines multi-stage design, uncertain trade policy with measure-specific reach, and a tariff base that follows the value accumulated at each crossing. Zhou et al. (2017) make a tariff base an outcome of the design through embodied emissions, and Cui et al. (2025) optimise a multi-stage photovoltaic chain under known restrictions, so each holds part of the combination. The combination matters because relocating a step changes which commodity crosses which border, and with it the customs value on which duty falls, so exposure becomes an outcome of the design. Governments already set trade policy in response to global value chains (Blanchard et al., 2016), and configuration is how a firm answers it.

This thesis builds that combination into one optimisation framework and takes up three of the four directions Dong and Kouvelis (2020) set out for tariff research in network design. From the first it takes tariff uncertainty across multiple stages of a supply chain, in a single period, so it models neither the timing nor the duration of a measure. From the second, product interdependence, it takes the structure of the bill of materials, through the chain architectures of Chapter 7, with fixed input requirements, so a firm can change where an input comes from but cannot replace it with another. From the third, competition, it takes a firm's vulnerability to trade policy in facility network design, and not the competitive response. The fourth, the strategic responses of partners in decentralised supply chains, lies outside its scope.

Table 2.2: Positioning of this study against the trade and network design literature.

Study	Strategic network design	Production representation	Trade policy representation	Policy uncertain at design
<i>Economics</i>				
Yi (2003)	No firm-level facility choice	Sequential stages with endogenous country allocation	Tariffs on gross goods in process	No
Rouzet and Miroudot (2013)	No	Observed multi-stage input-output structure	Cumulative duties embodied across production stages	No
Handley and Limão (2015)	No (export entry only)	Single-stage exporting firms	Probabilistic tariff process bounded by a worst-case rate	Yes (probabilistic)
Blanchard et al. (2026)	No	Global input-output production network	Tariffs on final goods and intermediate inputs	No
<i>Network design</i>				
Zhou et al. (2017)	Yes	Three-tier supply chain	Carbon tariffs on embodied emissions	No
Dong and Kouvelis (2020)	Yes, in reviewed and illustrative models	Facility networks and product allocation	Tariff effects on configuration and operational flexibility	Yes, in some models
Cui et al. (2025)	Yes, industry-wide capacity expansion to 2030	Five-stage production chain	Aggregate trade-cost indices, and a ban as prohibited flows	No
Ahmadi Dighsara et al. (2025)	Yes	Single-stage network	Uncertain bilateral tariffs and import quotas	Yes
This study	Yes	Multi-stage production with step qualification	Duty on each commodity's customs value, with reach set per class of measure	Yes, finite scenarios with relative-regret limits

3

Problem and research questions

A firm that places production steps in several countries pays duty at every border its goods cross, on the customs value they carry there (Figure 3.1). It must commit those locations before governments decide which trade measures to impose and at what rate. The firm therefore faces a *double unknown*, which measure will materialise and how severe it will be, and must design its network against a cost that depends on the design itself.




Figure 3.1: Tariff exposure under three configurations with the same tariffs, input and market. Moving production steps changes where borders are crossed and at what customs value. Proportions are illustrative.

Concentrating production in the lowest-cost locations minimises operating cost under stable trading conditions, at the price of higher exposure and less room to respond. A network with one location per step has nowhere to shift volume, so it pays the duty on every unit a measure reaches. Qualifying a

second location buys that option, at a cost paid in normal times whether or not the option is ever used. The firm must therefore decide how much efficiency it is willing to give up today for flexibility it may not need.

3.1. Main question

How should a firm design a multi-stage global value chain so that it remains cost-efficient under normal trading conditions and limits its exposure to trade policy it cannot foresee?

The decision maker is a firm designing its global production network. It decides which production steps to qualify at each candidate location, whether it performs them itself or through contract manufacturers, since the commitment concerns production capability and not ownership. The choice is hard to reverse, and once a measure materialises the firm can adjust production and flows only within the network it has committed. The problem is therefore a two-stage decision, shown in Figure 3.2.




Figure 3.2: The firm's two-stage decision.

Answering the main question therefore means establishing what a given increase in normal-times cost buys in exposure the firm no longer bears, and where that purchase stops paying.

3.2. Sub-questions

Four sub-questions break the main question down.

- **RQ1. What the design determines.** How does the configuration of a fragmented value chain determine its exposure to trade measures, and which of that exposure can production-network design remove?
- **RQ2. What the design is decided against.** How can uncertainty about future trade policy be represented for robust production network design when no probabilities can be assigned to it?
- **RQ3. How the design changes.** How does uncertainty about trade policy change the cost-minimising production network, and what determines the resulting network configuration?
- **RQ4. What the design costs, and its limit.** What does a robust network cost in normal trading conditions and under the trade measures it may face, and where does further protection stop paying?

Chapters 7 and 8 answer RQ1 by quantifying how exposure changes with the placement of production steps, using the deterministic model of Chapter 5. Chapter 4 answers RQ2 by turning the first part of the double unknown into a finite set of classes and fixing the second at each regime's cap, and Chapter 6 embeds both in a two-stage optimisation model. Chapters 7 and 8 also answer RQ3, by tracing how the committed network changes across chain architectures, the conditions a firm inherits and a documented industry, and RQ4, by reporting what a committed network costs under baseline tariffs and under each measure, and where further protection stops paying. Chapter 9 states the four answers directly.

3.3. Scope

This thesis builds a framework for the strategic design of a firm's global value chain, from purchased inputs through production to goods delivered in the country of sale. A chain within this scope enters the framework as data: its bill of materials and the customs value of each commodity, the candidate

locations of each production step, its suppliers and markets with their demand, the country of each node and whether each pair of countries is allied, neutral or adversarial, the baseline tariffs and the commodities each class of measures covers, and its costs in one consistent unit, such as shares of the final good's value. The representation of trade measures, the robust criterion and the solution procedure stay fixed, so two chains differ only in their data. Chapter 7 applies the framework to constructed chains built to isolate mechanisms, and Chapter 8 to a solar photovoltaic chain calibrated from public sources.

Trade policy enters as the ad valorem tariffs and equivalents defined in Section 2.1. A network answers such a measure by moving volume to a crossing the measure does not reach.

The thesis holds two things fixed so that designs differ in configuration alone. Demand is a parameter with a constant total, and its geographic distribution varies across the experimental instances, because a measure reaches only the share of output sold in the market it targets. The model also covers a single period and evaluates an initial commitment, so inventory, recovery over time and the up-front investment lie outside the study. A qualified step instead carries a cost per period, depreciation on owned capital or the charge for reserved capacity, so the model compares designs on what they cost to hold and operate and leaves out the timing and financing of the outlay. The two-stage structure still captures the order of events, since the firm commits the design before the measure and adapts the flows after it.

Exchange rates, lead times, yields and the physical disruption of facilities also bear on which network a firm should build. The thesis isolates the mechanism specific to trade policy, a duty whose number of charges and base both follow from the design, and leaves its interaction with those uncertainties to further research.

4

Trade-policy regimes

A *regime* is a class of trade-policy measures together with the rule that decides which movements the class reaches. Its cap is the fixed tariff increment it adds in full to the baseline tariff of every movement it reaches (Equation (6.3)). The term is the thesis's own and has no legal meaning. A movement is one commodity carried on one link of the network. Applied to a network, a regime creates every possible measure of its kind between the network's locations, and each measure becomes one scenario, priced as an ad valorem charge on the movements it reaches (Section 2.1).

This chapter defines six regimes, the movements each reaches and the cap each adds, and the composite of trade defence and anti-circumvention, which the thesis counts as a seventh regime. With Chapter 6 it answers RQ2. Each regime rests on the trade rules, which constrain what a government may decide, and on the decisions governments have taken. The rules fix what its measures can reach, and the recorded decisions fix whom they name and how much they charge. Chapter 6 then turns the scenarios of each regime into the uncertainty set of the design problem.

4.1. From measures to regimes

Three properties of trade-policy measures bear on their representation. First, every tariff change follows a policy decision. A government decides whether to invoke a measure, names the movements it covers and sets its rate, and enacts the decision in a legal act such as a proclamation, a determination or a regulation. These decisions change tariffs in large discrete steps (Fajgelbaum et al., 2020). Trade remedy investigations number in the thousands (Section 2.2.3), but no record gives the frequency with which a given measure reaches a given crossing of a given network, so the regimes carry no probabilities. Second, measures are one-sided. They raise the tariff on affected movements above the baseline tariff, the bilateral effectively applied rate (Section 2.2.1), never below it, and for a sunk commitment only that increase matters (Handley and Limão, 2015). Friendshoring also lowers the cost of qualifying origins, through a preference or a subsidy, but what it changes for a firm is the gap to the origins that do not qualify, so the model represents it as an increase on those origins. Third, measures target movements according to legal criteria, so one measure reaches every movement that satisfies those criteria.

These properties rule out the two standard representations. A probabilistic formulation requires a distribution over these decisions that the record does not supply. A budgeted interval set of the kind proposed by Bertsimas and Sim (2004), and applied to tariffs by Ahmadi Dighehsara et al. (2025), conflicts with the record in three ways. Each coefficient may take any value in its interval, whereas observed measures cluster at a few discrete rates. Coefficients vary independently, admitting scenarios in which movements the same legal criterion covers are affected differently. And the budget limits how many coefficients deviate at once, whereas the legal basis of a measure fixes how far it reaches.

The representation treats the two parts of the double unknown of Chapter 3 separately. Which measure materialises becomes a finite set of regimes that follows the institutional structure, since each class of measures is invoked or not and, once invoked, applies to a legally defined set of movements. How

severe it becomes one cap per regime, two for bilateral escalation, which every scenario charges in full (Section 4.2.3).

4.2. The six regimes

Read against the WTO rules of Section 2.2.1, the measures governments have taken fall into six regimes, grouped by the movements they reach. Three build on the WTO trade measures. The other three follow what governments have done outside them, imposing tariffs under domestic law (bilateral escalation), attaching origin conditions to preferences and subsidies (friendshoring) and restricting exports of critical inputs (export control). Table 4.1 gives the legal ground of each and the episodes that document it.

Table 4.1: Regimes, their legal ground and documented record.

Regime	Mechanism	Legal ground	Documented record
r_1	Bilateral escalation	Domestic statutes and emergency powers, outside the WTO trade measures	US–China 2018 (Fajgelbaum et al., 2020; Bown, 2021); reciprocal measures 2025 (Ignatenko et al., 2025)
r_2	Trade defence	WTO trade measure, anti-dumping and countervailing duties (Anti-Dumping Agreement, Agreement on Subsidies and Countervailing Measures)	EU–China solar 2013 (Goron, 2018); EU–China battery electric vehicles 2024 (European Commission, 2024b)
r_3	Market protection	WTO trade measure, safeguard (Agreement on Safeguards)	US steel 2002 (Read, 2005); US solar and washing machines 2018 (Bown, 2021)
r_4	Friendshoring	Withdrawal of preferences, a return to MFN, and origin or ownership conditions on subsidies under national law	US preference lapses (Handley and Limão, 2015); prohibited-foreign-entity rules for US clean-energy credits (Internal Revenue Service, 2026); EU Net-Zero Industry Act resilience criteria 2024 (European Union, 2024)
r_5	Anti-circumvention	No agreed WTO rule: extension of trade defence under national law	US findings on solar completed in Southeast Asia (U.S. Department of Commerce, International Trade Administration, 2023)
r_6	Export control	Licensing, quotas and duties on exports of critical upstream inputs, where WTO rules generally prohibit export quotas (GATT Art. XI) but not export duties	China rare earths 2010–2014 (Chen et al., 2021); China gallium and germanium 2023 (Javorcik et al., 2024); US semiconductor controls 2019–2020 (Bown, 2021)

Bilateral escalation (r_1) raises duties on trade between two jurisdictions under domestic statutory authority, without a prior injury finding. It targets a country pair rather than a producer or a product group, and both directions of the pair are exposed once retaliation follows. The record also shows that escalation need not be uniform across partners, which motivates the two caps of Section 4.2.3.

Trade defence (r_2) covers anti-dumping and countervailing duties. They differ in what triggers them, a firm's export pricing or a government's subsidy, but both follow an investigation into imports of a defined product from a defined origin and reach the same movements. Splitting them would add a regime with identical scenarios, so the model treats them as one.

Market protection (r_3), the safeguard, responds to a finding of serious injury from increased imports and applies to the affected product category irrespective of origin, leaving aside the exemptions for small developing-country suppliers. It is the only import-side regime that reaches every foreign origin at once.

Friendshoring (r_4) conditions a benefit on the origin or control of the supplier rather than imposing a charge. The market remains open, and a non-qualifying origin loses a preference, a subsidy or an

advantage in public auctions, which raises what it costs to serve that market by the benefit forgone. The regime appears as prohibited-foreign-entity restrictions conditioning United States clean-energy credits on the ownership or control of the producer (Internal Revenue Service, 2026), as the resilience criterion of the European Union's Net-Zero Industry Act, which public auctions and procurement must apply where a single third country supplies more than half of a net-zero technology (European Union, 2024), and as the withdrawal of tariff preferences, which returns a beneficiary to the MFN rate (Handley and Limão, 2015).

Anti-circumvention (r_5) extends a duty order in force against one origin to goods completed in a third country from that origin's inputs and sold into the same protected market. The WTO agreements contain no rule on such extensions, so each government sets its own test (Zhu, 2024). In the United States, for example, the test asks whether the processing in the third country is minor or insignificant (U.S. Department of Commerce, International Trade Administration, 2023), which confines an extension to the last steps of production. It can therefore reach two kinds of goods leaving a hub, the finished good and the inputs to its final step, which together form the *protected group*.

Both kinds are needed because anti-circumvention acts together with the trade defence it extends. Trade defence taxes the finished goods coming from the origin, and a firm can route around it through a hub in two ways, by finishing the good in the hub or by making the inputs there and finishing the good in the market. The extension follows both routes, as the US solar orders cover cells whether or not assembled into modules (U.S. Department of Commerce, International Trade Administration, 2023), so a design that splits its supply between the origin and a hub pays on both. Each kind generates its own scenario, because a market can extend the order to the hub's finished goods or to its inputs alone. Section 4.3 puts the two in force together as the composite.

The model keys the extension to the hub where the step takes place and does not trace content. Real extensions can also test what a good contains, as the 2023 solar determination did (U.S. Department of Commerce, International Trade Administration, 2023), so a modelled duty can fall on a good such a test would clear or miss one it would catch.

Export control (r_6) restricts the outbound movement of designated upstream inputs, through licensing, quantitative limits, export duties, or a combination of the three. It differs from the other five in that the affected movements leave the imposing jurisdiction rather than entering it, so exposure attaches to origin rather than destination. China restricted exports of rare earth minerals and left the finished products containing them unrestricted, so the measure reached the upstream input alone, until a WTO ruling led China to remove its quotas and duties (Mancheri, 2015). Such restrictions are common, and their use on industrial raw materials grew more than fivefold between 2009 and 2020, mostly as export taxes, which WTO rules generally allow while they generally prohibit export quotas (Kowalski and Legendre, 2023). The regime therefore covers a common practice that the agreements discipline only in part.

Chen et al. (2021) show that the downstream consequences of such restrictions vary systematically with the importance of the restricted input in production cost and with the elasticity of substitution across downstream sectors. The regime therefore places a cost wedge only on movements of the restricted input. Downstream cost then rises with the input's value in the chain, through the gross-value base of Section 2.1, and with where the consuming step sits, so the model needs no added penalty. Substitution away from the input lies outside it, since its input requirements are fixed (Section 2.5).

4.2.1. Exposure rules

Each regime generates its scenarios by one fixed rule, stated in Table 4.2. A rule names the anchor the measure keys on and the movements it reaches from that anchor, and it refers to no particular network. A network supplies what the rules read: the country of each node, the commodities it moves and the alignment of each pair of countries (Section 4.2.2). An application fixes three further inputs, the commodity scope of each rule, drawn from the measure's own product definition, the anchors it admits and the membership of the protected group. The same six rules therefore apply to any network, and applications differ in what the rules select, never in how the scenarios are formed, so their results remain comparable.

Table 4.2: Exposure rules of the six regimes.

Regime	Anchor	Movements exposed	Alignment	Scenarios generated
r_1 Bilateral escalation	country pair	every movement between the two countries, in either direction, carrying a commodity in the escalation schedule	lower cap for allied pairs	one per admitted pair
r_2 Trade defence	origin and protected market	movements from that origin into that market carrying a commodity in the subject-merchandise definition	origins not allied	one per admitted origin and market
r_3 Market protection	protected market	movements of the affected product category into that market, from every foreign origin	all origins alike	one per protected market
r_4 Friendshoring	protected market	movements into that market from non-eligible origins, carrying a commodity subject to the eligibility rule	origins not allied	one per protected market
r_5 Anti-circumvention	hub and protected market	movements of one protected-group commodity from that hub into that market	neutral hubs only	one per commodity, hub and market
r_6 Export control	restricting origin	outbound movements of one designated input from that origin, to every destination	all destinations alike	one per designated input

4.2.2. Alignment

Several regimes distinguish between allied, neutral and adversarial trading relationships. The distinction has a legal footing in one place. The agreements let members depart from MFN through free trade agreements and preferences (Section 2.2.1), so a partner holding preferential access is treated differently by law, and friendshoring works by withdrawing exactly that treatment. Elsewhere the categories are modelling constructs. Trade defence, anti-circumvention and tariffs imposed under domestic authority are all legally capable of reaching an ally, and none contains an alliance criterion.

Behaviour supports the constructs. Blanchard et al. (2016) find less protection applied against partners whose production is linked to the imposing country's own, across the measures behind three of the regimes here. That evidence is about value-added content rather than political or security alignment, so the categories are this model's proxy for integration.

Table 4.2 gives the alignment condition of each regime, and each follows the record. Friendshoring reaches non-allied origins by design, since it conditions a benefit on origin or ownership. Trade defence and anti-circumvention have reached adversarial origins and the neutral hubs that served them (Table 4.1), in line with the lighter protection of integrated partners above. Bilateral escalation reached allies too, at a lower rate (Section 4.2.3). Market protection reaches every origin by law, and export control every destination from the restricting origin, so neither depends on alignment. Categories stay fixed across scenarios, so the model never applies trade defence to an allied origin or anti-circumvention through an allied hub.

4.2.3. Regime caps

Each regime carries a cap, $\bar{\Delta}$, and bilateral escalation carries two, a lower one for allied pairs and a higher one for all others. A cap is the fixed tariff increment a regime adds in full to the baseline tariff, the bilateral effectively applied rate, of every movement it reaches (Equation (6.3)). It is a stress-test rate, which forecasts no future rate and is distinct from the bound rate a WTO member commits to. The regimes and their caps form the uncertainty set, which specifies which movements a shock can reach and by how much, while each application sets the rates the shocks are added to (Section 5.2.2). Each scenario applies the full cap, which is safe for cost, since a milder rate can only lower what a committed design pays.

For bilateral escalation, trade defence and market protection, the cap is read from published rates. Anti-

circumvention inherits the rate of the order it extends, which the model truncates. Friendshoring and export control publish no duty on the movements they affect, so their caps are ad valorem equivalents, one taken from the literature and one converted from an observed price wedge. Where several regimes expose the same movement, the model applies the largest cap, although duties can accumulate in practice, as the anti-dumping and countervailing rates in the 2023 solar determinations do.

Each regime is anchored to a documented episode. The robust design evaluates a committed network against these caps without determining their values.

Table 4.3: Regime caps, the episodes they are anchored to and how each value was obtained.

Regime	Cap $\bar{\Delta}$	Anchoring episode	How the value was obtained	Status
r_1	0.54	Reciprocal measures, April 2025	Announced reciprocal rate on Chinese goods plus the fentanyl tariff already in force	Published rate
	0.25	US steel tariff on national-security grounds, 2018	Rate the steel tariff applied	Published rate
r_2	0.36	EU anti-subsidy proceeding on battery electric vehicles, 2024	Highest rate in the proceeding, 35.3 percent, rounded up	Published rate
r_3	0.30	US steel safeguard 2002, US solar safeguard 2018	Highest rate in the schedule across both episodes	Published rate
r_4	0.20	General-equilibrium friendshoring scenarios	Tariff or trade-cost increase assumed in that literature	Borrowed parameter
r_5	0.54	US circumvention findings, 2023	Truncated from orders well above 100 percent to the highest finished-good rate in the set, to keep the regimes comparable	Modelled truncation
r_6	1.00	Chinese rare-earth export quota, 2010–2014	Converted from export-to-domestic price ratios, bottom of the observed range	Measurement, converted here

Bilateral escalation Bilateral escalation uses two caps because escalation differs across relationships. The higher cap takes 0.54, the 34 percent reciprocal tariff announced for Chinese goods in April 2025 plus the 20 percent fentanyl tariff already in force (Ignatenko et al., 2025). The lower cap, for allied pairs, takes 0.25, the 25 percent tariff the United States applied to steel on national-security grounds from 2018, which reached allied suppliers including Canada, the European Union and Japan (Bown, 2021). It takes this applied rate rather than the 20 percent announced for the European Union in April 2025, which was cut to 10 percent within a week. Neither measure formally differentiated partners by alliance status, and the model maps the lower observed escalation to allied pairs and the higher to all others.

Trade defence Trade defence takes a cap of 0.36, the highest rate determined in the European Union's 2024 anti-subsidy proceeding on battery electric vehicles from China, 35.3 percent for SAIC, rounded up (European Commission, 2024b).

Market protection Market protection takes a cap of 0.30, corresponding both to the highest rate in the 2002 United States steel safeguard (Read, 2005) and to the opening rate of the 2018 solar safeguard (Bown, 2021). The opening rate is the highest in a safeguard schedule, since the safeguard rules require measures lasting more than a year to be progressively liberalised (Charnovitz and Hoekman, 2013).

Friendshoring No authority publishes a rate for friendshoring, because its effects arise from the withdrawal of preferential treatment or of eligibility for subsidies and public auctions rather than from a customs duty, and no study reviewed here estimates an ad valorem equivalent for it. The cap takes

0.20, the 20 percent that Javorcik et al. (2024) assume for trade between geopolitical blocs, as a tariff or as an iceberg trade cost, a scenario value they set to match the scale of the 2018 US–China tariffs.

Anti-circumvention Anti-circumvention extends an existing duty order to third-country completion and therefore inherits whatever rate that order already carries. In the 2023 United States solar determinations, exporters without company-specific rates faced an anti-dumping duty of 238.95 percent together with an all-others countervailing duty of 15.24 percent (U.S. Department of Commerce, International Trade Administration, 2023).

Using those rates directly would place this regime far above the others and make the regimes difficult to compare. The thesis therefore truncates the cap to 0.54, the highest cap any other regime applies to a finished good, which understates the rates observed in practice.

Export control Export control is capped at 1.00. Chen et al. (2021) report that the economically binding measure in the Chinese rare-earth episode was not the published export duty of 15 or 25 percent but the export quota, cut by 40 percent while the production quota fell by 19 percent. The quota created a price wedge between domestic and export markets. Comparing prices for identical rare-earth materials at the same stage of production, they find export prices between two and six times domestic prices after 2010. Paying twice the domestic price is economically equivalent to a 100 percent ad valorem tariff. The cap therefore takes the bottom of the measured range, since the objective is a severe but comparable rate on an exposed movement.

Status of the calibration The caps do not all have the same evidential status, which Table 4.3 records, and two carry the most judgement, the borrowed friendshoring figure and the truncated anti-circumvention rate.

The calibration bears on the second unknown only, the severity of the regimes, and leaves the movements they expose unchanged. Which regimes a network can clear follows from the exposure rules. Which regime binds depends on the caps, and the thesis does not vary them.

4.3. Bundles

A regime applied to a network generates several scenarios, one for each country pair, origin, hub or input its rule names (Table 4.2). On a network that imports into the United States from China and Vietnam, for example, trade defence generates two scenarios, one against each origin, and which of them hits a design harder depends on where the design produces. The scenarios of a regime form its *bundle*, and a design is judged by its worst scenario in that bundle, wherever the measure happens to land. The eight bundles are fixed in advance and are the same for every network, while the scenarios they hold are every possibility the rules admit among the network's locations and commodities, whatever design the firm commits.

Because a design is also tested in normal circumstances, when no regime is in force, bundle b_0 , the baseline, holds a single scenario in which every movement carries its baseline tariff. A design's cost there is its cost in normal times. Bundles b_1 to b_6 each hold the scenarios of one regime, so that a design faces every regime on its own.

Measures can accumulate in practice, but the thesis combines regimes only where one exists to extend another. Anti-circumvention extends a trade-defence order and so presupposes one, as in the 2023 solar findings (U.S. Department of Commerce, International Trade Administration, 2023), and Section 4.2 shows how the two act together on the routes around that order. Bundle $b_7 = \{r_2, r_5\}$, the *composite*, therefore holds scenarios that pair one trade-defence measure with one anti-circumvention measure protecting the same market. It carries the caps of its two parts, 0.36 on the movements from the origin and 0.54 on those through the hub, and the thesis counts it as a seventh regime. Anti-circumvention also keeps a bundle of its own, b_5 , in which only the route through the hub is taxed, so the results can show whether a design is hurt by that route alone or only once trade defence taxes the origin as well.

Bundles b_1 to b_7 are the adverse bundles, one for each of the seven regimes. Which regime comes, and which of its scenarios, is the first part of the double unknown, and the regime caps fix the second. Chapter 6 turns the bundles into the uncertainty set and evaluates designs against each regime.

5

Methodology

5.1. Conceptual model

This chapter formulates the deterministic design model, through which RQ1 traces how duty accumulates along a fragmented network. Any chain within the scope of Section 3.3 enters it as data. The model represents the chain as a directed network in which a firm chooses where production capabilities are established and how commodities move between them. It finds the least-cost network that meets a given demand through decisions on qualification, production and flows.

The network contains three node types. Supplier nodes introduce externally sourced commodities, production nodes are candidate locations at which individual production steps may be qualified and operated, and customer nodes represent final demand. Warehouses are excluded because the analysis ends when the final product reaches the destination market.

Production follows a bill of materials over several steps. Each production step consumes specified inputs and produces specified outputs, so production can be distributed across locations, which changes where borders are crossed and at what customs value (Figure 3.1).

This deterministic formulation provides the baseline design problem. Chapter 6 retains the same network structure and makes the tariff vector scenario-dependent, while allowing production and flows to adapt after the scenario is known.

5.2. Mathematical formulation

5.2.1. Sets and network structure

The network is directed, with node set V and arc set E , where arcs represent feasible shipping links. Table 5.1 lists the sets.

While E describes the physical topology, X restricts flows to commodity movements consistent with the bill of materials and node roles. It is enumerated directly from the network structure, so no variable is created for an inadmissible movement. Three movement types are admitted: externally sourced commodities from suppliers to production locations, intermediates between production locations, and final goods from production locations to customers.

Self-loops at production locations, where $i = j$, carry intermediates between consecutive co-located steps. They cross no border and carry no shipping charge, since an internal handoff is not a shipment, whereas movements between distinct nodes are charged whether or not a border is crossed. Co-locating consecutive steps therefore removes both the tariff exposure and the shipping cost between them.

5.2.2. Parameters

The model uses parameters describing production technology, commodity values, tariffs, costs, capacities and demand, summarised in Table 5.2.

Table 5.1: Sets used in the deterministic model.

Symbol	Description
V	Set of nodes
V^S	Set of supplier nodes, $V^S \subseteq V$
V^P	Set of candidate production locations, $V^P \subseteq V$
V^C	Set of customer nodes, $V^C \subseteq V$
E	Set of directed arcs between nodes
C	Set of countries, each a trade-policy jurisdiction
$c(i)$	Country of node i , $c(i) \in C$
G	Set of commodities, including inputs, intermediates and final goods
G^F	Set of final-product commodities, $G^F \subseteq G$
$\mathcal{P}$	Set of production steps, indexed by $p \in \mathcal{P}$
X	Set of feasible arc–commodity triples, $X \subseteq E \times G$

Table 5.2: Model parameters.

Symbol	Description
<i>Production technology</i>	
α_{pg}^{in}	Input quantity of commodity g required per unit of production step p
α_{pg}^{out}	Output quantity of commodity g produced per unit of production step p
<i>Commodity value and tariffs</i>	
v_g	Customs value of one unit of commodity g
τ_{ijg}^0	Baseline tariff, the ad valorem rate on commodity g on arc (i, j) in normal times
<i>Variable costs</i>	
c_{ig}^{supply}	Per-unit cost of procuring commodity g from supplier node i
c_{ijg}	Per-unit shipping cost of commodity g on arc (i, j)
c_{ip}^{prod}	Per-unit cost of running production step p at location i
<i>Commitment costs</i>	
f_{ip}^P	Per-period cost of holding production step p qualified at location i
<i>Capacity and demand</i>	
K_{ip}	Maximum capacity of production step p at location i
ρ	Minimum throughput of a qualified production step, an absolute quantity
D_{ig}	Demand for final commodity g at customer node i

The coefficients α_{pg}^{in} and α_{pg}^{out} define the bill of materials by specifying the commodities each production step consumes and produces. Technological admissibility is carried by capacity. Where $K_{ip} = 0$, step p cannot be performed at location i , so no qualification or production variable is created for that pair, and every sum and constraint over $i \in V^P$ and $p \in \mathcal{P}$ below runs over pairs with $K_{ip} > 0$. External supply is unbounded.

Production cost c_{ip}^{prod} is specific to a production step and location and captures the variable cost of performing that activity there, including differences in labour, energy, locally purchased inputs and production efficiency. The commitment cost f_{ip}^P is likewise defined at the step–location level and represents the per-period cost of maintaining that production capability. For dedicated capacity it carries depreciation on the capital the step requires together with the fixed cost of keeping it available, and for reserved external capacity it is the charge for the reservation. The up-front investment is not modelled as an outlay and enters only through that depreciation. Both cost terms are therefore stated per period, as in network design models that annualise every cost component (Shapiro et al., 2009) over the single-period horizon common in facility location (Melo et al., 2009).

Customs value rests on the transaction value of the goods (World Trade Organization, 1994a), so it differs from the cost of producing the commodity. The model simplifies it to one value v_g per commodity, the same whatever location produced it, so locating a step in a dearer location can raise production cost without changing v_g . The baseline tariff τ_{ijg}^0 is the rate in force before any regime applies, the bilaterally effectively applied rate that country $c(j)$ charges on commodity g from country $c(i)$. It is the lower of that country's MFN rate and any preferential rate it grants $c(i)$ (Section 2.2.1). Every measure the regimes of Chapter 4 represent stays out of it and enters only through them, and each application sets its value.

5.2.3. Decision variables

The model distinguishes production activity from commodity movement, because qualification, capacity and minimum throughput apply to a production step at a location, whereas flows record commodities moving between nodes. Table 5.3 lists the three decision variables.

Table 5.3: Decision variables.

Symbol	Description
x_{ijg}	Flow of commodity g from node i to node j , defined for $(i, j, g) \in X$
y_{ip}	Production quantity of step p at location $i \in V^P$
a_{ip}	Binary variable equal to 1 if step p is qualified at location $i \in V^P$, 0 otherwise

Production cannot be read off the flows. Under the minimum throughput of (5.6), a qualified step may produce more than the network ships onward, with the surplus freely disposed of, so recovering production from the flows would need a surplus variable in its place. Carrying y_{ip} lets capacity, the throughput floor and production cost attach directly to the activity they restrict.

The binary variable a_{ip} records the strategic placement decision. Qualifying a step establishes production capability at that location for the period and incurs f_{ip}^P irrespective of its subsequent utilisation. It costs the same whether or not other steps already sit at that location, since no fixed cost attaches to a site as such. Placing steps together therefore pays off only through the shipping and duty it saves between them.

5.2.4. Objective function

The deterministic design problem minimises the total cost of meeting demand:

$$\begin{aligned}
\min \quad & \underbrace{\sum_{i \in V^P} \sum_{p \in \mathcal{P}} f_{ip}^P a_{ip}}_{\text{commitment}} + \underbrace{\sum_{\substack{(i,j,g) \in X: \\ i \in V^S}} c_{ig}^{\text{supply}} x_{ijg}}_{\text{procurement}} \\
& + \underbrace{\sum_{i \in V^P} \sum_{p \in \mathcal{P}} c_{ip}^{\text{prod}} y_{ip}}_{\text{production}} + \underbrace{\sum_{(i,j,g) \in X} c_{ijg} x_{ijg}}_{\text{shipping}} + \underbrace{\sum_{\substack{(i,j,g) \in X: \\ c(i) \neq c(j)}} \tau_{ijg}^0 v_g x_{ijg}}_{\text{duty}}
\end{aligned} \tag{5.1}$$

5.2.5. Constraints

Production Production follows the bill of materials, enforced separately on the input and output sides. A production location cannot consume more of a commodity than it receives,

$$\sum_{p \in \mathcal{P}} \alpha_{pg}^{\text{in}} y_{ip} \leq \sum_{j: (j,i,g) \in X} x_{jig} \quad \forall i \in V^P, \forall g \in G, \tag{5.2}$$

and cannot ship more of a commodity than it produces,

$$\sum_{j: (i,j,g) \in X} x_{ijg} \leq \sum_{p \in \mathcal{P}} \alpha_{pg}^{\text{out}} y_{ip} \quad \forall i \in V^P, \forall g \in G. \tag{5.3}$$

These two constraints conserve flow at every production location, and the demand constraint below does so at every customer, so together they are the flow-conservation constraints of the model. They hold as inequalities, so a location may receive more of an input than it consumes and ship less of an output than it produces. Cost minimisation rules out both unless something forces them, and the minimum throughput below can force the second, since a step committed before the policy decision is known runs whether or not its output is wanted.

Demand Demand for each final commodity is met as an equality, so the firm cannot reduce its tariff exposure by withdrawing from a market:

$$\sum_{j: (j,i,g) \in X} x_{jig} = D_{ig} \quad \forall i \in V^C, \forall g \in G^F. \tag{5.4}$$

Qualification and capacity Production of a step is bounded above by the capacity available where that step has been qualified,

$$y_{ip} \leq K_{ip} a_{ip} \quad \forall i \in V^P, \forall p \in \mathcal{P}, \tag{5.5}$$

and below by the minimum throughput,

$$y_{ip} \geq \rho a_{ip} \quad \forall i \in V^P, \forall p \in \mathcal{P}. \tag{5.6}$$

The upper bound ties production to qualification, so a step produces only where its capability is committed, and it binds directly where capacity is scarce. Production cost is charged on realised output, so capacity is priced at the margin through c_{ip}^{prod} rather than chosen (Liang and Wilhelm, 2013), and the design decides only where capability is qualified, at cost f_{ip}^P .

The lower bound requires a qualified step to sustain a minimum level of production, so that it is a genuine production route rather than a token option. Together the two bounds make production either zero or at least ρ , the minimum order quantity of backup sourcing under capacity reservation, which

Hou et al. (2017) likewise state in absolute units. The floor ρ is a share of total demand in units of the final good.

Every step in both applications converts one unit of input into one unit of output, so one floor serves every step, where a chain with other coefficients would need one per step. A floor proportional to K_{ip} would give identical processes different minimum scales only because their candidate capacities differ, which the absolute floor avoids.

Domains

$$x_{ijg} \geq 0 \quad \forall (i, j, g) \in X, \quad y_{ip} \geq 0, \quad a_{ip} \in \{0, 1\} \quad \forall i \in V^P, \forall p \in \mathcal{P}. \quad (5.7)$$

The deterministic model minimises (5.1) subject to (5.2) to (5.7).

6

Robust optimisation under trade-policy uncertainty

6.1. Overview

This chapter turns the two-stage decision of Chapter 3 into a robust model and, with Chapter 4, answers RQ2. The model tests each candidate design against every measure the regimes can generate. It accepts a design only if, under each regime, its worst-case cost exceeds that of the best design for the regime by at most a chosen share. Among the accepted designs, it commits the one that is cheapest in normal times.

Formally, the first stage commits a design, the qualification decisions a , common to every scenario. The second stage re-optimises production and flows (y^ω, x^ω) within that design once a scenario ω is known, following the two-stage robust network-design structure of Atamtürk and Zhang (2007). The uncertainty lies only in which scenario occurs. A scenario decides which movements a regime names, and each named movement pays its baseline tariff plus that regime's cap, a fixed rate set in Chapter 4. The scenario therefore fixes the tariff vector τ^ω , while the network and its constraints stay the same, so production and flows are recourse decisions. The deterministic model of Chapter 5 solves the same design problem under baseline tariffs alone, the tariffs that apply in normal times. Its optimum C_0^* , attained by the deterministic network, is the benchmark, and what a robust design costs above it in normal times is that design's normal-times premium. The rest of this section states the criterion in words and argues for it. Section 6.2 then builds the finite set of scenarios, Section 6.3 writes the criterion out as a model and Section 6.4 sets out how it is solved and measured.

The criterion

The model is min-max p -robust over a finite set of scenarios. The regimes generate the scenarios, and the scenarios of each regime form its bundle b (Section 4.3). For each regime, the reference cost C_b^* is the lowest cost any design can guarantee against every scenario of that regime.

A committed design's regret under a regime is how far its worst-case cost there sits above the reference cost, as a share of the reference cost. The firm sets a regret tolerance θ , the largest regret it is prepared to accept. A design is admissible at θ when its regret stays within θ under every regime, and among the admissible designs the model returns the one that is cheapest under baseline tariffs. The tightest tolerance that still admits a design is the robustness floor. The design returned there, the *floor design*, is the one whose largest regret across the regimes is smallest, the minimax-regret design (Snyder and Daskin, 2006). Section 6.3 writes out both halves of the name, the min-max reference in Section 6.3.1 and the p -robust bound in Section 6.3.2.

Why this criterion

The record supplies no probabilities (Section 4.1), so the model has no expected cost to minimise. It asks instead whether a design performs acceptably in every scenario, the usual approach of robust

optimisation when probabilities are unknown (Snyder and Daskin, 2006).

Performance is measured by regret rather than by cost. A limit on worst-case cost, minimax cost, would let the harshest measure decide everything, even when no design could avoid its duty. Regret compares a design with the best design for the same regime and asks how much more the chosen commitment costs. Duty that every design pays does not count, only duty that a better design avoids, so the level of a cap matters less than it would under minimax cost. Regret also mirrors how decision makers are judged, after the measure has come, against what a better choice would have cost under it (Snyder and Daskin, 2006).

The reference cost is set per regime rather than per scenario, because which scenario of a regime hits hardest depends on the design (Section 4.3). A design that moves its assembly from China to Vietnam to escape trade defence against China, for example, becomes exposed to trade defence against Vietnam. Snyder and Daskin (2006) set the reference per scenario, which compares designs under one fixed scenario, as if the firm knew in advance which origin trade defence would name. Here every design, the reference network included, is judged by its own worst scenario of the regime, and the reference cost is the lowest of these worst cases. Regret then measures the cost of not knowing which regime will come.

The tolerance applies to each regime separately. Each regime is a severe shock, and the firm wants to be covered whichever one comes, so doing well against one regime cannot make up for failing another, as an average across regimes would allow. The separate checks matter because protecting against one regime can open a design to another. Moving the final step from an adversarial origin to a neutral hub, for example, escapes bilateral escalation against that origin but brings the design within reach of anti-circumvention.

Among the designs that pass, the model picks the one that is cheapest in normal times, the cost the firm pays for as long as no regime arrives. The criterion thus makes two changes to the stochastic p -robustness of Snyder and Daskin (2006). Normal-times cost replaces expected cost, and the reference cost per regime replaces the best design per scenario.

In the resilience terms of Bruckler et al. (2024), set out in Section 2.3, the committed design absorbs a measure and the re-optimised production and flows adapt to it. Regret is measured after both, so it captures the two together without telling them apart. Recovery over time lies outside the single-period model.

6.2. The uncertainty set

Each regime r yields a finite library $\mathcal{A}^r$ of measures. A measure is written as the set of movements $A \subseteq X$ it reaches together, such as trade defence on final goods from one origin into one market. The rules of Table 4.2, with the commodity scope of Table 6.1, generate every library.

Table 6.1: Bundles in this thesis.

Bundle	Regimes in force	Commodities reached	Alignment
b_0	none, the baseline scenario ω_0 only	none	none
b_1	r_1 Bilateral escalation	all commodities	lower cap for allied pairs
b_2	r_2 Trade defence	final goods	origins not allied
b_3	r_3 Market protection	final goods	all origins alike
b_4	r_4 Friendshoring	all commodities	origins not allied
b_5	r_5 Anti-circumvention	final goods and direct inputs to the final step, one scenario each	neutral hubs only
b_6	r_6 Export control	the restricted upstream input	all destinations alike
b_7	r_2 and r_5 , the composite	as in b_2 and b_5	as in b_2 and b_5

Market protection in market c_m , for example, yields the measure

$$A = \{(i, j, g) \in X : c(i) \neq c_m, c(j) = c_m, g \in G^F\}, \quad (6.1)$$

every foreign-origin movement of a final good into that market, all reached at once, with a cap of $\bar{\Delta}(A) = 0.30$.

A scenario ω consists of the measures in force, one for a single regime and two for the composite b_7 , and the baseline scenario ω_0 contains none. The uncertainty set of bundle b is

$$\mathcal{U}(b) = \begin{cases} \{\omega_0\} & \text{for the baseline } b_0, \\ \mathcal{A}^r & \text{for a single regime, } b = \{r\}, \\ \text{pairs in } \mathcal{A}^{r_2} \times \mathcal{A}^{r_5} \text{ protecting the same market} & \text{for the composite } b_7, \end{cases} \quad (6.2)$$

and the full uncertainty set is $\mathcal{U} = \bigcup_{b \in \{b_0, \dots, b_7\}} \mathcal{U}(b)$. Each pair in $\mathcal{U}(b_7)$ is a *joint scenario*. The reference cost and the tolerance apply per regime, while the second stage responds to the scenario, so the realised tariffs and the second-stage decisions carry the index ω .

In scenario ω , a movement pays its baseline tariff plus the cap of a measure in ω that names it,

$$\tau_{ijg}^\omega = \begin{cases} \tau_{ijg}^0 + \bar{\Delta}(A) & \text{if } (i, j, g) \in A \text{ for some } A \in \omega, \\ \tau_{ijg}^0 & \text{otherwise.} \end{cases} \quad (6.3)$$

Each cap is fixed by the regime that generates the scenario, so the uncertainty lies only in which scenario occurs. In the composite, trade defence and anti-circumvention can both name the final good moving from the neutral hub into the market. If that happens the movement pays its baseline tariff plus the larger of the two caps (Section 4.2.3).

6.3. The model

The model reuses the network, variables and constraints of Chapter 5, with production and movements indexed by scenario. Table 6.2 lists the quantities used to evaluate a committed design.

Table 6.2: Quantities used to evaluate a committed design.

Symbol	Meaning
$Q^\omega(a)$	recourse cost: cheapest operating cost of committed design a under scenario ω
$Q^b(a)$	worst-scenario operating cost: largest recourse cost over the scenarios of bundle b
C_b^*	reference cost of the regime with bundle b : the lowest worst-case total cost over all designs
C_0^*	reference cost of the baseline bundle b_0 : the cost of the deterministic network
$R^b(a)$	relative regret: how far the worst-case total cost of a under the regime with bundle b sits above C_b^*
θ	regret tolerance: the largest regret the firm accepts, the bound on $R^b(a)$ under every regime
$\pi(a)$	normal-times premium: how far the cost of a under baseline tariffs sits above C_0^*

6.3.1. The reference

The objective of Chapter 5, Equation (5.1), is decomposed into a fixed commitment cost that does not vary across scenarios and an operating cost that does,

$$C^{\text{fixed}}(a) = \sum_{i \in V^P} \sum_{p \in \mathcal{P}} f_{ip}^P a_{ip}, \quad (6.4)$$

$$C^{\text{oper}}(x^\omega, y^\omega; \tau^\omega) = \sum_{\substack{(i,j,g) \in X: \\ i \in V^S}} c_{ig}^{\text{supply}} x_{ijg}^\omega + \sum_{i \in V^P} \sum_{p \in \mathcal{P}} c_{ip}^{\text{prod}} y_{ip}^\omega + \sum_{(i,j,g) \in X} c_{ijg} x_{ijg}^\omega + \sum_{\substack{(i,j,g) \in X: \\ c(i) \neq c(j)}} \tau_{ijg}^\omega v_g x_{ijg}^\omega, \quad (6.5)$$

where f_{ip}^P is the per-period commitment cost of Section 5.2.2, so the two cost terms are compared on the same time basis.

For a fixed design, the cheapest response to a scenario and the worst such response within a bundle are

$$Q^\omega(a) = \min_{x^\omega, y^\omega \geq 0} C^{\text{oper}}(x^\omega, y^\omega; \tau^\omega) \quad \text{s.t. constraints (5.2) to (5.6),} \quad (6.6)$$

$$Q^b(a) = \max_{\omega \in \mathcal{U}(b)} Q^\omega(a). \quad (6.7)$$

The reference cost of the regime with bundle b is the lowest worst-case cost any design attains, for the reasons set out in Section 6.1. Substituting (6.7) and (6.6) shows that it is a min-max over the bundle,

$$\begin{aligned} C_b^* &= \min_a \left[C^{\text{fixed}}(a) + Q^b(a) \right], \quad b \in \{b_0, \dots, b_7\}, \\ &= \underbrace{\min_a}_{\text{design}} \left[C^{\text{fixed}}(a) + \underbrace{\max_{\omega \in \mathcal{U}(b)}}_{\text{worst scenario}} \underbrace{\min_{x^\omega, y^\omega \geq 0} C^{\text{oper}}(x^\omega, y^\omega; \tau^\omega)}_{\text{best response}} \right]. \end{aligned} \quad (6.8)$$

The outer minimum chooses the design before any scenario is known. For each design, the maximum takes the scenario of the bundle that costs it most, and the inner minimum gives the cheapest production and flows the design allows under that scenario, within the constraints of Chapter 5. C_b^* is therefore the lowest total cost any design can guarantee whichever scenario of the bundle comes, and the design attaining it is the regime's *reference network*. This is the min-max part of the criterion. The p -robust part follows in Section 6.3.2.

Because b_0 contains only ω_0 , its reference cost is the optimum of the deterministic model of Chapter 5, written C_0^* throughout, and the design attaining it is the *deterministic network*.

6.3.2. The robust design problem

The robust problem commits one design a . It returns the cheapest design under baseline tariffs whose worst-case cost under every regime stays within a fraction θ of that regime's reference cost, the p -robustness condition of Snyder and Daskin (2006) with θ in place of p ,

$$\min_a C^{\text{fixed}}(a) + Q^{\omega_0}(a) \quad \text{s.t.} \quad C^{\text{fixed}}(a) + Q^b(a) \leq (1 + \theta) C_b^*, \quad b \in \{b_1, \dots, b_7\}. \quad (6.9)$$

Because the uncertainty set is finite, each recourse cost can be written out with its own production and flow variables (Zeng and Zhao, 2013). With the references fixed, (6.9) then becomes one mixed-integer program, with one set of design decisions and a copy of the production and flow variables for every scenario:

$$\min_{a, \{x^\omega, y^\omega\}} C^{\text{fixed}}(a) + C^{\text{oper}}(x^{\omega_0}, y^{\omega_0}; \tau^{\omega_0}) \quad (6.10)$$

$$\text{s.t.} \quad C^{\text{fixed}}(a) + C^{\text{oper}}(x^\omega, y^\omega; \tau^\omega) \leq (1 + \theta) C_b^* \quad \forall b \in \{b_1, \dots, b_7\}, \omega \in \mathcal{U}(b) \quad (6.11)$$

$$\text{constraints (5.2) to (5.6) on } (x^\omega, y^\omega), \quad x^\omega, y^\omega \geq 0 \quad \forall \omega \in \mathcal{U} \quad (6.12)$$

$$a_{ip} \in \{0, 1\} \quad \forall i \in V^P, p \in \mathcal{P}. \quad (6.13)$$

The objective prices the design under baseline tariffs alone. Trade policy enters only through constraint (6.11), where each scenario's tariff vector τ^ω sets the operating cost that must stay within $(1 + \theta) C_b^*$.

For a committed design, the relative regret under the regime with bundle b is

$$R^b(a) = \frac{C^{\text{fixed}}(a) + Q^b(a)}{C_b^*} - 1. \quad (6.14)$$

Constraint (6.11) is the requirement $R^b(a) \leq \theta$ written out, so θ bounds the regret under every regime, and a design that meets the requirement under all of them is *admissible* at θ . Since the objective gains nothing from closing the slack, an admissible design usually stays below θ under most regimes.

The tolerance does not apply to the baseline bundle. Its cost is what the model minimises, and how far that cost sits above the cost of the deterministic network is the price of protection, the *normal-times premium*

$$\pi(a) = \frac{C^{\text{fixed}}(a) + Q^{\omega_0}(a)}{C_0^*} - 1, \quad (6.15)$$

which is also the regret of a against b_0 . Figure 6.1 summarises the two evaluations of a committed design.




Figure 6.1: Evaluation of a committed design a . Under each regime its worst-case cost is compared with the regime's reference cost C_b^* , and the regret R^b must not exceed θ . Under baseline tariffs its cost is compared with C_0^* , the cost of the deterministic network, which gives the normal-times premium π .

A design is *feasible* when it meets demand within the constraints of Chapter 5. Tariffs enter only the cost, so feasibility never depends on them. Admissibility does: once a higher tariff lifts a design's worst case more than θ above the reference, the model must commit a different design, relocating, duplicating or consolidating production. A binding capacity would limit how much flow can move away from a taxed movement, so it could make a design inadmissible at a lower tariff but never infeasible. Both applications set capacity high enough never to bind.

6.4. Solution procedure and experimental measures

Solving the model takes two steps, which Algorithm 1 in Appendix C sets out in full. The first computes the eight reference costs in (6.8) once, one for the baseline and one for each regime, and reuses them in every robust solve. For a regime, the reference cost is one mixed-integer program with a copy of the production and flow variables for every scenario and one variable for the worst-case cost, which must be at least the operating cost of each scenario. For the baseline it is the deterministic model.

The second solves the robust problem at a given θ and returns one committed design. Its objective prices only the baseline, so its plans for the other scenarios need only stay within the tolerance. An evaluation pass therefore fixes a and solves (6.6) for every scenario, giving $Q^\omega(a)$, and from these $Q^b(a)$ by (6.7) and $R^b(a)$ by (6.14). These cheapest plans cost no more than the robust problem's own, so every reported regret stays within θ . The regime under which a design's regret is largest is its *binding regime*.

6.4.1. Robustness frontier and floor

Solving the robust problem for progressively tighter values of θ traces the *robustness frontier*. Each tolerance returns the admissible design with the lowest baseline cost, whose normal-times premium is written $\pi(\theta)$. The frontier is the trade-off between cost and maximum regret that Snyder and Daskin (2006) report, with the normal-times premium in place of expected cost. Tightening θ shrinks the admissible set, so $\pi(\theta)$ never falls. Because qualification decisions are discrete, it rises in steps: it stays flat while the same design remains admissible and jumps when that design drops out.

At its loose end the frontier reaches the deterministic network, whose largest regret across the regimes is its *deterministic exposure*. That value is the tolerance at which the deterministic network first becomes admissible, and beyond it the premium is zero.

The *robustness floor*, denoted $\theta_{\min}$, is the smallest tolerance at which an admissible design exists,

$$\theta_{\min} = \min_a \max_{b \in \{b_1, \dots, b_7\}} R^b(a), \quad (6.16)$$

so each design is judged by its largest regret across the regimes, and the robustness floor is the lowest such value any design attains, the minimax-regret value that Snyder and Daskin (2006) reach by tightening the tolerance until no design remains. The design that attains it is the floor design. Bisection locates it, since a design admissible at one tolerance is admissible at every looser one.

Under one condition the robustness floor cannot fall below the premium of the floor design. If some design escapes a regime at no extra cost, that regime's reference cost equals C_0^* . A design's regret under that regime is then its premium plus whatever extra the regime's worst scenario costs it, or the premium alone when the design escapes that regime too. The robustness floor therefore cannot fall below the premium of the floor design.

The robustness floor is meaningful only for a fixed design space, since changing the candidate locations or design constraints changes both the admissible designs and the reference costs they are judged against.

Table 6.3 collects the terms Chapters 7 and 8 use to report committed designs.

Table 6.3: Terms used to report committed designs. Table 6.2 gives the quantities they rest on.

Term	Symbol	Meaning
Reference network	attains C_b^*	The design with the lowest worst-case cost under a regime, its reference cost
Deterministic network	attains C_0^*	The design with the least cost under baseline tariffs, with no tolerance imposed
Deterministic exposure	$\max_b R^b$	Largest regret of the deterministic network across the regimes
Binding regime	$\arg \max_b R^b(a)$	Regime under which a design's regret is largest, several where regimes tie
Robustness frontier	$\pi(\theta)$	Designs returned as θ is tightened, each with its normal-times premium
Robustness floor	$\theta_{\min}$	Smallest tolerance at which an admissible design exists
Floor design	attains $\theta_{\min}$	The design whose largest regret equals the robustness floor

7

Computational experiments

This chapter applies the model to constructed value chains, in which one property can change while the rest stay fixed, and answers RQ1, RQ3 and RQ4 for them. Section 7.1 sets up the instances. Section 7.2 tightens the tolerance and follows how the committed design changes and what each change costs. Section 7.3 holds three designs fixed and asks what each pays under every regime and which exposure it escapes. Section 7.4 varies supplier origins, location rights and demand geography to find what decides how much protection a firm needs and where it stops paying. Section 7.5 turns the findings into priorities for a firm.

7.1. Experimental design

The instances are constructed so that one property can vary while the rest stay fixed, which traces an outcome to the feature that changed. A single observed chain cannot do this, because its topology, geography and cost structure arrive together. Three elements rest on evidence. The architectures follow documented production topologies, the production-cost multipliers lie within measured regional cost differences, and the regimes are those of Chapter 4. The gross values follow the property of Section 2.1.1 that customs value does not rise in step with production cost. The cost shares are declared assumptions, so the magnitudes reported here belong to this calibration. Chapter 8 applies the framework to a documented industry.

The main experimental grid varies three properties. *Architecture* determines how production streams converge and therefore where border crossings arise and how adjustment can propagate through the chain. *Demand orientation* determines how output is distributed across markets and therefore how much of the network a market-facing regime can reach. *Production-location flexibility*, represented by the location-rights parameter q_i , limits how many locations each production step may be qualified in and therefore how much geographic recourse the firm can commit in advance. Targeted experiments later vary supplier origins and supplier dependence to test mechanisms identified in this main grid.

7.1.1. Reference architectures

Three reference architectures span the cases: a single stream, multiple streams converging once, and multiple streams converging in stages. In the network typology of Feldmann and Olhager (2019), Serial is linear and Modular convergent, the snake and the spider of Baldwin and Venables (2013), and Complex adds a second, intermediate convergence.

Serial is a single stream, as in flat steel. Modular has three independent streams meeting at a single closing step, as in consumer electronics. Complex adds an intermediate merge, as in automotive, where two streams combine with a purchased strategic component at mid-chain before joining the third stream at the closing step. The component is supplied from a single fixed origin at 15 percent of final-good value, analogous to a vehicle platform relying on a semiconductor or battery component from one qualified supplier location. Its origin cannot be changed, isolating the effect of dependence on a single-origin input. Figure 7.1 shows the topologies, and the model decides where each step is located

and whether it is duplicated.




Figure 7.1: The three reference architectures. Circles are produced commodities and diamonds are supplied inputs, each shaded by the gross value it embodies as a percentage of the final good. In Complex the upper branch carries the combined value of two streams and the component, so it is darker than the lower stream at the same position along the chain.

Serial has five production steps and one supplied raw material, Modular 13 steps and three raw materials, and Complex 11 steps, three raw materials and the strategic component. The number of production steps, S , divides the production and commitment budgets in Section 7.1.3.

7.1.2. Geographic setting and demand

Production may be located in four blocs: C^H a high-cost industrial market, C^M an allied mid-cost producer, C^T a neutral processing and transit hub, and C^L a low-cost production base, with North America, Central and Eastern Europe or advanced East Asia, Southeast Asia and China as illustrative counterparts. Every step may be located in any bloc, at one candidate location per bloc and that bloc's production-cost multiplier, from 1.00 in C^L to 1.50 in C^H (Table 7.3). Supplied inputs have fixed origins. Serial's raw material comes from C^L , Modular's three from C^L , C^T and C^M , one per stream, and Complex's from C^L and C^T for the two merging streams and from C^M for the third, the bloc that also supplies the strategic component.

The four blocs simplify real trade geography so that each carries one role the regimes of Chapter 4 act on. Seen from C^H , the principal market, C^L is the adversarial origin, C^M the ally that trade defence spares, and C^T the neutral hub through which circumvention runs. The blocs' costs differ because the trade-off requires it: at equal cost, producing inside the served market would avoid most exposure at no extra production cost. For the same reason the calibration makes the adversarial bloc the cheapest

and the principal market the dearest, since a cheapest origin that measures spared would remove the trade-off between cost and exposure. The ally is spared because tariffs fall more lightly on partners whose production is integrated with the imposing market (Blanchard et al., 2016). Measured labour costs place the ratio of high-cost to low-cost production cost between 1.07 and 1.61 (Appendix A.1), so the gap between the dearest and the cheapest bloc is evidence-based, and the two intermediate values are a calibration choice.

Alignment is a relation between blocs rather than a property of one, since the regimes of Chapter 4 treat allied origins differently from adversarial or neutral ones. Table 7.1 defines the relations. Every bloc except C^T can be a served market. C^T is neutral to every bloc and can hold production.

Table 7.1: Relations between blocs, read by row from each served market.

	C^H	C^M	C^T	C^L
C^H	own	allied	neutral	adversarial
C^M	allied	own	neutral	adversarial
C^T	neutral	neutral	own	neutral
C^L	adversarial	adversarial	neutral	own

Total demand is fixed at 100 units and only its distribution over C^H , C^M and C^L varies. Every served market carries its own trade defence, market protection, friendshoring and anti-circumvention, read from its row of Table 7.1, while bilateral escalation and export control do not depend on where demand sits. Demand geography therefore determines which market-anchored regimes can reach the network and what share of output they reach. Table 7.2 defines the seven orientations.

Table 7.2: Demand orientations.

Orientation	Served markets	Shares
Concentrated	C^H , C^M or C^L , each in turn	100
Balanced dual	C^H and C^M	50/50
Asymmetric dual	C^H and C^M , either as the primary	80/20
Dispersed	C^H , C^M and C^L	40/30/30

7.1.3. Economic parameterisation

Table 7.3 lists the economic inputs of the constructed instances, each with its value and source. A facility is one step qualified at one location, so a plant that holds two steps counts as two facilities.

Table 7.3: Economic inputs of the constructed instances. Costs are in units of a final-good value of 100.

Parameter	Notation	Value	Source
Final-good value	v_g	100	Normalisation
Total demand	D	100, geography varied	Reference scale
Raw material	c_{ig}^{supply}	10, plus a component at 15 in Complex	Assumption
Production cost	c_{ip}^{prod}	35/ S per step, 20/ S in Complex, times the bloc multiplier	Assumption
Bloc multiplier		C^L 1.00, C^T 1.10, C^M 1.30, C^H 1.50	Range from Ceglowski and Golub (2012) and KPMG LLP and The Manufacturing Institute (2020, p. 10)
Commitment cost	f_{ip}^P	15 D / S per qualified facility	Assumption
Margin		40, the residual	Shin et al. (2012) and Damodaran (2026)
Intermediate gross values	v_g	Table 7.6	Set to follow the stage profile of Rungi and Del Prete (2017)
Shipping cost	c_{ijg}	3% of embodied cost across a border, 1% within	Calibrated on observed freight (Barrot et al., 2019)
Minimum throughput	ρ	5 per qualified facility	Assumption
Capacity bound	K_{ip}	150	Assumption, never binds
Baseline tariff	τ_{ijg}^0	0	By construction (Section 7.1.5)
Regret tolerance	θ	15 values from 0.05 to 1.00	Swept (Section 7.1.6)
Regime caps	$\bar{\Delta}$	0.20 to 1.00	Chapter 4
Location rights	q	1 to 4	Varied (Section 7.4.2)

Baseline value division All values are percentages of the final good's value, normalised to 100, as in a common-size income statement, which states every cost as a percentage of sales (Foster, 1986). The final good's value takes the place of sales: it is the good's arm's-length worth rather than the firm's revenue, since the model fixes demand and minimises cost, leaving revenue out of the objective. It is also the final good's tariff base and the scale to which every cost is referred.

Table 7.4 reports the division at a bloc multiplier of one. Raw material, production and commitment are set, the margin is the residual, and shipping, which depends on the design, is paid out of the margin. Complex purchases its strategic component, so its internal production share is lower while cost of goods and margin are unchanged. The shares are declared assumptions, which Appendix A.1 compares with published figures alongside the multiplier and shipping calculations. Chapter 8 replaces them with a division derived from industry costs.

Table 7.4: Baseline division of the final good's value at a bloc multiplier of one.

Component	Serial	Modular	Complex
Raw material	10	10	10
Strategic component	0	0	15
Production	35	35	20
Commitment	15	15	15
Cost of goods	60	60	60
Margin (residual)	40	40	40
Final good	100	100	100

Cost rules Table 7.5 gives the three cost rules and the terms through which they enter the objective. Every architecture divides the same budgets of Table 7.4 evenly over its S production steps, so more steps cut the same production into smaller steps at an unchanged total. Trade models of sequential production work the same way, splitting one fixed production process into stages that can sit in different countries (Baldwin and Venables, 2013). A step of the automotive-like Complex is thus a smaller share of its product's value than a step of the steel-like Serial, although a real car step may cost more in

absolute terms. The results are shares and rates, so this scale does not affect them, and no difference between architectures comes from a difference in total cost. The even split is a simplification, since real steps differ in what they cost.

Table 7.5: Cost rules and their model terms at a bloc multiplier of one, for Serial, Modular and Complex in turn.

Cost	Model term	Type	Scales with	Value
Production	$c_{ip}^{\text{prod}} y_{ip}^{\omega}$	variable	bloc multiplier, output	7.00, 2.69, 1.82 per unit and step
Commitment	$f_{ip}^T a_{ip}$	fixed	facilities committed	300, 115, 136 per facility
Shipping	$c_{ijg} x_{ijg}^{\omega}$	variable	crossings, flow	1.35 per unit of final good crossing a border, in all three

Commitment is incurred in full once a facility is qualified, however much it produces. Splitting a step's output across blocs therefore divides its production cost in proportion, while every additional facility carries its full commitment, so diversification is cheap in production and expensive in commitment.

Shipping is variable. Each movement pays a charge per unit set in proportion to the embodied cost of what it carries, independent of the good's customs value, so the charge rises along the chain as cost accumulates and a crossing late in the chain costs more than one early in it. The rate of 3 percent across a bloc border is calibrated on observed freight: a Serial chain that changes bloc at every step then pays about as much over its five crossings as one shipment into the United States pays in freight, insurance and other charges, a median of 4.7 percent of its free-on-board value (Barrot et al., 2019). Each crossing therefore costs less than observed freight. The rates of 1 percent within a bloc and zero at one location are assumptions.

Unit costs are percentages of the final good's value, and instance totals cover the 100 units of demand. A deterministic network serving C^H therefore costs the 60 of Table 7.4 for each of the 100 units, 6,000, plus shipping of 145.0 in Serial, 158.3 in Modular and 203.3 in Complex. Of that shipping, 135 is the final good crossing into the market at 1.35 per unit (Table 7.5), and the rest moves supplied inputs.

Gross value Gross value is the customs value v_g of a commodity, as a percentage of final-good value. No source observes gross values for constructed chains, so they are set by hand to meet three conditions. Supplied inputs enter at their purchase value, 10 for the raw materials and 15 for the strategic component, and at a merge the values of the inputs add up. The value each step adds follows the property of Section 2.1.1, largest at the first and last steps of a chain and smallest in the middle (Rungi and Del Prete, 2017). Every step adds more than its production cost, since customs value also carries the margin. In Serial the value added follows a symmetric parabola over the five steps, 21.6, 16.2, 14.4, 16.2 and 21.6, so each end step adds half as much again as the middle one while all five cost the same to run. Modular and Complex keep this shape along each path: the increments fall along each stream and rise again towards the closing step, which adds the most, 25 in Modular and 13 in Complex. Table 7.6 gives every value. The values are fixed per commodity wherever it is produced (Section 5.2.2) and are not varied in the experiments, so the magnitudes reported here belong to this profile, which reproduces no particular product.

Table 7.6: Gross-value profiles of the three architectures.

Architecture / stream	Gross value along the production path
Serial	10.0 → 31.6 → 47.8 → 62.2 → 78.4 → 100.0
Modular, each of three streams	3.3 → 11.2 → 16.7 → 20.9 → 25.0
Final assembly	$3 \times 25.0 \rightarrow 100.0$
Complex, streams 1 and 2	3.3 → 11.0 → 15.0
Strategic component	15.0
Merged sub-assembly	45.0 → 49.6 → 60.0
Complex, stream 3	3.3 → 12.0 → 18.0 → 22.6 → 27.0
Final assembly	$60.0 + 27.0 \rightarrow 100.0$

7.1.4. Location rights, minimum throughput and capacity

Because the objective is the lowest baseline cost, the commitment cost limits how many locations a firm qualifies. It does not price the other limits on how many locations a step can hold, such as supplier approval, tooling, certification and the managerial attention needed to keep each one viable. Location rights make that limit an experimental variable, bounding the number of locations at which a step may be qualified,

$$\sum_{i \in V^P} a_{ip} \leq q \quad \forall p \in \mathcal{P}. \quad (7.1)$$

The parameter asks how much geographic flexibility a network needs, the question the process-flexibility literature asks of plants and products, where a few well-placed links deliver most of what full flexibility would (Jordan and Graves, 1995). A second qualified location lets the firm move production once a regime is known, the contingent rerouting that Tomlin (2006) finds is often part of the optimal response to a supply disruption. With one candidate location per bloc, q ranges from 1 to 4, and at $q = 4$ the constraint cannot bind. Every headline result in this chapter is therefore computed on the unrestricted problem, and Section 7.4.2 varies q over $\{1, 2, 3\}$ to report what each additional right buys. The reference networks carry the same constraint, so the committed design and its benchmark share one design space, which makes the resulting floors internally consistent and, as Section 7.4.2 shows, not comparable across levels of q .

Qualified production follows the minimum-throughput constraint of Equation (5.6). In the experiments, $\rho = 5$, equal to five percent of total demand, so every qualified facility produces at least five units in every scenario. Surplus output may be discarded where that is cost efficient.

The capacity bound K_{ip} is one and a half times total demand. Surplus output only adds cost, so no facility needs to produce more than total demand, and the bound cannot bind. Treating location and capacity acquisition separately is defensible where capacity costs do not vary across candidate locations (Verter and Dincer, 1995), which holds here by construction, since the main calibration charges the same commitment cost in every bloc.

7.1.5. Regime instantiation

The baseline tariff τ^0 is the bilateral effectively applied rate of normal trading conditions (Section 5.2.2), and the regimes of Chapter 4 add to it. The blocs have no tariff schedule of their own, so this chapter sets the baseline to zero, the rate the documented solar chain of Chapter 8 carries on every crossing its designs use. A zero baseline makes every premium, floor and regret attributable to the regimes alone, and the deterministic network ignores ordinary duties, which would already favour designs with fewer crossings.

The exposure rules of Chapter 4 apply to the four-bloc network unchanged, reading the bloc of each node, the relation of each pair in Table 7.1 and the commodities. Each scenario charges its regime at a cap that never changes, so the uncertainty lies only in which regime comes and which movements it names (Section 4.3).

Table 7.7 shows the resulting scenarios for a Serial chain serving C^H , each movement written as (origin, destination, commodity). The anti-circumvention protected group, the final good and the direct inputs to the closing step, holds two commodities in Serial, four in Modular and three in Complex. Bundle codes b_0 to b_7 follow Table 6.1, and scenario counts range from 18 for Serial serving one industrial market to 63 for Modular under dispersed demand (Table A.1 in Appendix A.1).

Table 7.7: Scenarios of each bundle for a Serial chain serving C^H .

Bundle	Anchor	Scenarios	Cap	Example scenario
b_0 Baseline	none	one, no regime in force, as in the deterministic model	none	every movement at its baseline tariff of zero
b_1 Bilateral escalation	bloc pair	one per unordered bloc pair	0.54, 0.25 if allied	all movements between C^L and C^H , either direction, any commodity
b_2 Trade defence	origin and market	one per adversarial or neutral origin	0.36	(C^L , C^H , final good)
b_3 Market protection	market	one, all foreign origins	0.30	(C^M , C^H , final good), (C^T , C^H , final good), (C^L , C^H , final good)
b_4 Friendshoring	market	one, non-allied origins, all commodities	0.20	every movement into C^H from C^T or C^L
b_5 Anti-circumvention	hub and market	one per commodity in the protected group, per hub and market	0.54	(C^T , C^H , final good) and (C^T , C^H , direct input), as separate scenarios
b_6 Export control	origin	one per supplied input and its origin	1.00	(C^L , j , raw material) for every $j \neq C^L$
b_7 Composite	as in b_2 and b_5	one per pair of trade-defence and anti-circumvention measures protecting the same market	0.36 and 0.54	(C^L , C^H , final good) at 0.36 and (C^T , C^H , final good) at 0.54, in one scenario

7.1.6. Experimental grid

The main experimental grid combines three architectures, seven demand orientations and four location-rights levels, giving $3 \times 7 \times 4 = 84$ instances. The supplier-origin and supplier-dependence experiments introduced later are targeted perturbations of these instances rather than additional dimensions of the main grid.

Each instance is solved across the tolerance grid

$$\theta \in \{0.05, 0.06, 0.07, 0.08, 0.09, 0.10, 0.12, 0.15, 0.20, 0.25, 0.30, 0.45, 0.60, 0.80, 1.00\},$$

after which the floor is bisected to 10^{-3} as in Algorithm 1. The grid is denser where designs change quickly, and its upper end exceeds the largest deterministic exposure in the main grid, 0.836 for Complex serving C^M . All other parameters are held constant within the grid.

7.1.7. Implementation

The model is implemented in Python 3.13 and solved with Gurobi 13.0.2 on an Apple M3 Pro with 18 GB of memory, at a one-hour limit per solve and a relative gap of 10^{-4} . The robust problem is solved single-threaded with a fixed seed, so that a rerun returns the same design where several designs tie on cost. The largest instance, Modular serving dispersed demand, holds 63 scenarios. The slowest single solve took 350 seconds, and no solve behind a reported result reached the one-hour limit. Appendix C.2 reports the solver record.

7.1.8. Reading the results

The measures used to evaluate committed designs are defined in Chapter 6. The premium compares a design's normal-times cost with the deterministic network, regret compares its worst-case cost under a regime with that regime's reference cost, and the binding regime is the one under which its regret is largest. Table 7.8 defines the terms this chapter adds to those of Table 6.3.

Relative quantities are reported to three decimals. Anti-circumvention reaches no deterministic network in the grid, so its reference cost equals C_0^* in every instance, and by the condition of Section 6.4.1 no

floor falls below its premium. Every floor is therefore premium-bound or exposure-bound. The grid leaves gaps between solved tolerances, so rung counts are lower bounds.

Table 7.8: Additional terms used to interpret the computational results.

Term	Meaning
Premium-bound	A robustness floor at which $\theta_{\min} = \pi$. The remaining regret under each regime does not exceed the normal-times premium paid for robustness.
Exposure-bound	A robustness floor at which $\theta_{\min} > \pi$. Remaining tariff exposure causes the regret under at least one regime to exceed the normal-times premium.
Rung	One design on the frontier, numbered from the loosest tolerance. Each rung is optimal over an observed range of tolerances.
Instance	One combination of architecture, demand orientation and location rights. The main grid contains 84 instances, of which the 21 at $q = 4$ are used for most architecture and regime comparisons.
C^H instance	One of the three instances at $q = 4$ with all demand in C^H , one per architecture, used for worked examples unless stated otherwise.
Midpoint	The middle rung of a frontier by rung count, the tighter of the two central rungs when the count is even. It is counted in rungs, so it need not lie midway in tolerance or premium. On the frontiers of the C^H instances it is rung 3 of 4 in Serial, 4 of 6 in Modular and 5 of 8 in Complex.
Move	A change in where one step is qualified between consecutive rungs: a relocation when its number of locations stays the same, a duplication when it grows and a consolidation when it shrinks.
Shock cost	What a regime adds to a design's cost, its cost under the regime's worst scenario less its baseline cost. A design's worst shock cost is the largest of these across the seven regimes.
Break-even cap	The cap at which the duty the deterministic network pays under a regime equals the increment of the regime's reference cost over C_0^* . Above it, defending against the regime pays (Section 7.2.5).

7.2. The robustness frontier

This section follows the robustness frontier, the sequence of committed designs the model returns as the regret tolerance θ tightens (Section 6.4.1). It shows what changes in the network as protection rises and what each change costs. It first sets out the benchmarks every design is judged against, the deterministic network and each regime's reference cost. It then follows how the committed design changes along the frontier, which regime binds each design and where maximum-regret protection reaches its limits. It closes with the break-even caps, which show how low a regime's cap can be before defending against it stops paying.

7.2.1. Reference networks

The deterministic network The deterministic network is the committed design with the lowest baseline cost, C_0^* , the design the firm would choose if it left the regimes out. In each C^H instance it places every step in C^L , the cheapest bloc, and sends the final good across one border into C^H . Table 7.9 builds its baseline cost from the value division of Table 7.4 for 100 units of demand. Shipping adds the 135 the final good pays to cross into C^H , at 1.35 per unit, and what the supplied inputs pay to reach production (Section 7.1.3). No duty is paid, since the baseline tariff is zero.

Table 7.9: Baseline cost C_0^* of the deterministic network in the three C^H instances by component, in units of final-good value for 100 units of demand.

Component	Serial	Modular	Complex
Raw materials	1,000.0	1,000.0	1,000.0
Strategic component			1,500.0
Production cost	3,500.0	3,500.0	2,000.0
Commitment cost	1,500.0	1,500.0	1,500.0
Shipping of the final good	135.0	135.0	135.0
Shipping of supplied inputs	10.0	23.3	68.3
C_0^*	6,145.0	6,158.3	6,203.3

Reference costs Each regime's reference cost C_b^* is the lowest worst-case cost any committed design can reach against that regime alone (Section 6.3.1). The design that attains it is the regime's reference network, the network the firm would commit if it knew which regime comes but not which scenario it names. Every reference cost is computed with each regime charging its full cap (Table 7.7).

Table 7.10 gives each reference cost and its increment over C_0^* , what defending against the regime costs, including any duty the reference network design still pays. The increment ranges from nothing to 765.7, up to 12 percent of the deterministic cost.

Table 7.10: Reference cost C_b^* of each regime and its increment over C_0^* in the three C^H instances, with every regime at its cap, in units of final-good value for 100 units of demand. The baseline row gives C_0^* .

Bundle	Regime	Serial		Modular		Complex	
		C_b^*	Increment	C_b^*	Increment	C_b^*	Increment
b_0	baseline	6,145.0		6,158.3		6,203.3	
b_1	bilateral escalation	6,854.5	709.5	6,663.9	505.6	6,969.0	765.7
b_2	trade defence	6,469.0	324.0	6,329.9	171.6	6,333.8	130.5
b_3	market protection	6,519.0	374.0	6,329.9	171.6	6,333.8	130.5
b_4	friendshoring	6,469.0	324.0	6,366.0	207.7	6,387.4	184.1
b_5	anti-circumvention	6,145.0	0.0	6,158.3	0.0	6,203.3	0.0
b_6	export control	6,145.0	0.0	6,288.8	130.5	6,452.3	249.0
b_7	composite	6,469.0	324.0	6,329.9	171.6	6,333.8	130.5

Observation 7.1 *Reach, more than rate, determines what a regime costs*

The cost of defending against a regime is determined by the exposure it can reach in the network rather than by its tariff cap alone. Bilateral escalation and anti-circumvention both carry a cap of 0.54, yet bilateral escalation is the most expensive regime to defend against in every architecture while anti-circumvention requires no additional commitment in any of them.

The contrast follows from how the two regimes attach to the network. Bilateral escalation is anchored on a bloc pair and reaches every commodity moving in either direction between that pair. Avoiding one exposed crossing can therefore create another: relocating production from one side of the pair changes where value crosses the border but does not necessarily remove the bilateral exposure. The cheapest defence therefore needs a second location to reroute to when a pair is named. This makes it the largest increment in every architecture.

Anti-circumvention has the same cap but a much narrower effective reach in these instances. It targets protected movements from the neutral hub C^T into the served market, and no deterministic network uses that hub for those movements, so the regime has nothing to act on.

Export control depends on where supplied inputs come from. Serial's raw material comes from C^L , where its reference network uses it, so export control adds no cost. Modular and Complex buy raw materials from C^T and C^M , away from where they would otherwise be used, and rerouting the finished good cannot remove those crossings. The strategic component, worth 15, makes the increment largest in Complex.

The remaining regimes follow the same logic. Trade defence spares allied C^M and the market itself, so the firm escapes it by finishing in C^M or in C^H . Finishing in C^M is cheaper to run, at a bloc multiplier of 1.30 against 1.50, but the final good then crosses a border into C^H , which costs 90 more in shipping than delivering within the market. The production saving is 0.20 times the closing step's production cost of 7.00 per unit in Serial, 2.69 in Modular and 1.82 in Complex (Table 7.5), which is 140, 54 and 36 over 100 units. Only Serial's saving exceeds the 90, so Serial finishes in C^M and Modular and Complex in C^H . Market protection names every foreign origin, C^M included, so only finishing in C^H escapes it. Modular and Complex already finish there under trade defence, so their two reference networks coincide. Serial must give up C^M , and its reference cost rises by the net saving it loses, $140 - 90 = 50$, from 6,469.0 to 6,519.0. The composite matches trade defence in all three architectures, because its anti-circumvention part reaches no further movement in the trade-defence design.

7.2.2. Design changes along the frontier

Tightening the regret tolerance θ changes the committed design in steps. A design is admissible at θ when its regret stays within θ under every regime, and each tolerance returns the admissible design with the lowest baseline cost (Section 6.4.1). The result is a staircase of committed designs, the rungs of the frontier, each costing more in normal times and leaving a smaller largest regret. Table 7.11 shows it for the Serial C^H instance, where every adjustment occurs at the closing step, and its last column gives the tolerances at which each design is optimal.

Regret follows from the reference costs of Table 7.10 (Equation (6.14)). The deterministic network, rung 1, pays 5,400 under bilateral escalation, so its worst-case cost there is $6,145.0 + 5,400 = 11,545.0$ against a reference cost of 6,854.5, a regret of

$$\frac{11,545.0 - 6,854.5}{6,854.5} = 0.684.$$

It is therefore admissible only from $\theta = 0.684$. At a tighter tolerance, such as the grid value 0.60, the firm commits rung 2, whose largest regret is 0.568. At the floor design, rung 4, the binding regimes collect no duty and their reference costs equal C_0^* , so the largest regret equals the premium, $6,771.5/6,145.0 - 1 = 0.102$.

Table 7.11: Frontier of the Serial C^H instance. Only the closing step is shown because the first four steps remain in C^L throughout. The lower end of each tolerance interval is the design's largest regret.

Rung	Closing step	Move	Baseline cost	π	Binding	Optimal for $\theta \in$
1	C^L (100)	none	6,145.0	0.000	bilateral escalation	[0.684, 1.00]
2	C^L (95), C^T (5)	duplication	6,454.2	0.050	composite	[0.568, 0.60]
3	C^T (95), C^H (5)	relocation	6,638.5	0.080	friendshoring	[0.298, 0.45]
4	C^M (95), C^H (5)	relocation	6,771.5	0.102	anti-circumvention, export control	[0.102, 0.25]

The deterministic network has one closing location, in C^L , so any regime reaching the flow it sends from C^L into C^H is paid in full. The first transition qualifies a second closing location in C^T to defend against bilateral escalation. Once the targeted bloc pair is known, the bulk of finishing switches to the location off that pair, so a regime aimed at one relation no longer reaches the whole of output. The minimum-throughput requirement keeps both facilities running, so the facility on the targeted relation still produces its five units. When bilateral escalation targets C^L and C^H , the pair that costs this design most, the design discards those five units rather than ship them across the taxed border and pays no bilateral duty, since at this cap replacing them costs less than the duty (Section 7.3.1). Moving the bulk to the dearer C^T and remaking the five discarded units still costs money, so bilateral escalation adds to the design's cost (Figure 7.3), though far less than paying the duty would.

The next two transitions move that flexibility rather than add to it. The second shifts bulk finishing to C^T and qualifies a small facility in the served market C^H in place of C^L . The third shifts the bulk from C^T to allied C^M and keeps the C^H facility. Neither adds a location, but each changes which exposures the pair of locations can route around. C^T is neutral and remains within the reach of several regimes, while

C^M is allied and exempt from trade defence and friendshoring. Table 7.12 gives what each transition costs in normal times.

Table 7.12: What each transition on the frontier of the Serial C^H instance costs, in units of final-good value.

Transition	Move	Commitment	Production	Shipping	Total
1 to 2	duplicate closing step in C^T	300.0	3.5	5.7	309.2
2 to 3	shift bulk closing from C^L to C^T and add C^H	0.0	80.5	103.8	184.3
3 to 4	shift bulk closing from C^T to C^M , retaining C^H	0.0	133.0	0.0	133.0
	cumulative	300.0	217.0	109.5	626.5

The first transition costs most, 309.2, nearly all of it the commitment cost of the second closing facility. The next two relocate facilities without adding one, so they add only production and shipping cost, 184.3 and 133.0, and the floor design's premium totals 626.5.

Modular and Complex follow the same logic at more steps. Modular adjusts its closing step as Serial does (Table A.2). Complex changes seven of its eleven steps over eight rungs (Table A.3), which Section 7.2.3 follows rung by rung and explains.

7.2.3. Binding regimes along the frontier

Figure 7.2 evaluates every rung design of the frontiers of the three C^H instances under the regime that binds it, so the bars are not a like-for-like comparison between rungs. They show how the source of worst-case exposure changes as robustness is added.

Serial and Modular start the same way. Bilateral escalation binds first, because it reaches the whole final good on one pair, and once both add a second closing facility the composite binds. Serial then moves its bulk finishing to the neutral hub, which friendshoring still reaches, and then to allied C^M , where only anti-circumvention and export control bind and neither collects duty.

Modular keeps its bulk in C^L longer, because finishing in the market would bring its stream outputs from C^L across the same pair. Bilateral escalation therefore binds again at rung 3, and friendshoring at rung 4, once finishing in the market draws intermediates from non-allied C^L . At rung 5 the bulk still finishes in cheap C^L , with small facilities in C^H and C^M as alternatives, and market protection binds on the minimum output of the two facilities outside the market. The floor design drops C^L , which removes the exposed crossing into the market and lowers the largest regret, but moves bulk finishing into the dearer market, raising the premium from 4.1 to 4.7 percent. What remains is export control on a raw material from another bloc (Table A.2).

Complex follows a different pattern. Bilateral escalation binds at seven of its eight rungs, because it can name any pair of blocs and so reaches whichever crossing each rung leaves most exposed (Table A.4). Figure A.1 draws each rung under that scenario, with the taxed movements ringed. Each new rung moves the taxed value without removing it. Most take it off the targeted pair, and others move it onto the allied pair between C^M and C^H , where the cap is 0.25 rather than 0.54, so the duty falls at every rung, from 5,400 to 452. Rungs 3 and 5 both carry the sub-assembly worth 60, but its duty falls from 3,240 to 1,500. Single steps move first. The merge moves only at rung 4 (Table A.3), together with the step after it, which would otherwise receive its input across a border, and in allied C^M it also stops the strategic component from crossing at all.

The three architectures therefore reach different kinds of floor. Serial is premium-bound, its largest regret equal to its premium (Section 7.2.2). Modular and Complex are exposure-bound, both still paying 333 under export control on a supplied input with a fixed origin. Moving the consuming step to the supplier would remove that crossing but raise the premium so far that the largest regret would not fall. Every exposure left at the floor design is of this kind. Removing it would raise the largest regret, so it is cheaper to keep than to remove.




Figure 7.2: Realised cost of every rung design under its binding regime, in units of final-good value, tolerance loosest on the left. Each bar is labelled with its binding regime's cap in percent.

Effect of architecture The architecture acts through its bill of materials and its supplied inputs, which fix where a crossing is forced. The bill of materials joins streams at merges, and streams made in different blocs must cross to meet there. The finished good crosses unless it is made in the market, and a supplied input crosses unless the step that uses it moves to its origin, which pays only for a valuable input. Value rises at every merge and peaks at the closing step, so robustness comes from moving merges and the closing step into the market or an allied bloc, which takes the high-value crossings out of most regimes' reach and leaves the forced crossing upstream, on goods worth less. Serial and Modular can do this at one point, the closing step, and their floor designs add one facility there, at premiums of 10.2 and 4.7 percent. Complex cannot. Its merge combines two streams, whose raw materials come from C^L and C^T , with the strategic component from C^M , so some of its inputs cross a border wherever it sits, and the sub-assembly it produces, worth 60 before final assembly, crosses again unless it is made next to final assembly. Its floor design therefore adds two facilities, one of them a second merge, at a premium of 8.6 percent. Each premium pays for the added facilities and for producing and shipping from the new locations.

As protection tightens, the model changes where steps are qualified. In 155 of the 164 moves on the 21 frontiers at $q = 4$ the step is a merge, the closing step or a step feeding one, or a step receiving a supplied input from another bloc. The other nine are all in Complex, at steps inside a stream. Each follows a neighbouring step, such as the merge, out of the low-cost bloc into a dearer one, because the lower production cost of the low-cost bloc did not outweigh the cost of shipping across the border between the two.

Observation 7.2 *Exposure identifies where robustness is valuable*

A step is worth moving where the architecture forces a crossing, at a merge, the closing step or a step feeding one, and where a supplied input arrives from another bloc. Value rises at every merge and peaks at the closing step, so merges and the closing step move into the market or an allied bloc and leave the crossing upstream, on goods worth less. A step inside a stream moves only to follow a neighbouring step into a dearer bloc, when shipping between them would cost more than its cheaper production saves.

7.2.4. Limits of maximum-regret protection

A tighter tolerance lowers the largest regret, but it need not lower what every regime costs. Figure 7.2 showed each design under the regime that binds it, and Figure 7.3 shows what every regime adds to the cost of each design on the frontiers of the three C^H instances. These costs fall overall as the premium rises, but bilateral escalation becomes far dearer at Modular's third rung, and anti-circumvention and the composite at Complex's third.




Figure 7.3: What each regime adds to the cost of every design on the frontiers of the three C^H instances, loosest tolerance on the left, with the cumulative premium below each panel. Costs are in units of final-good value.

A regime's duty can even rise as protection tightens. When Complex serves C^L , friendshoring's duty rises from 433 at the deterministic network to 503 at the floor design, while the largest regret falls from 0.197 to 0.059. The path between the two ends of a frontier can also carry exposure that neither end has. Anti-circumvention reaches no deterministic network or floor design in the 21 instances, but it reaches 12 midpoint designs and seven of the twelve interior rungs on the frontiers of the C^H instances, which route production through the neutral hub on the way out of C^L .

Observation 7.3 *Maximum-regret protection is not uniform*

Tightening the regret tolerance lowers the largest regret across regimes, but need not lower what every regime costs. Intermediate designs may also introduce exposures that are absent from both the deterministic network and the floor design. Which regime binds therefore depends on the committed design and cannot be fixed in advance.

7.2.5. Break-even caps

Every result so far charges each regime its full cap. The break-even cap asks how low a regime's cap can be before defending against it stops paying, if the firm knew which regime would come. It is the cap at which the duty the deterministic network pays under the regime equals what defending against the regime costs, the increment of Table 7.10. Above it, the regime's reference network costs less than the deterministic network under that regime.

The deterministic network has one location per step, so it cannot move production and pays the cap on all the customs value the regime reaches. The reference network costs the increment of Table 7.10 more, so the break-even cap is the increment divided by the value reached. In Serial, bilateral escalation reaches the whole final good, 10,000 of customs value, and its reference network costs 709.5 more, so the break-even cap is $709.5/10,000 = 0.071$. Every regime that taxes the final good reaches the same 10,000 in all three architectures, while export control reaches a supplied input, worth 333 in Modular and 1,500 in Complex.

Table 7.13 gives every break-even cap beside the regime's cap. For bilateral escalation it is an upper bound, because its reference network would cost less at a lower cap. Both costs are per period, so the comparison leaves out the cost of moving a network that already exists.

Table 7.13: Break-even caps beside each regime's cap, as shares of customs value, in the three C^H instances. None means the regime reaches nothing in the deterministic network. The bilateral escalation entries are upper bounds.

Regime	Cap	Serial	Modular	Complex
Bilateral escalation	0.54	0.071	0.051	0.077
Trade defence	0.36	0.032	0.017	0.013
Market protection	0.30	0.037	0.017	0.013
Friendshoring	0.20	0.032	0.021	0.018
Anti-circumvention	0.54	none	none	none
Export control	1.00	none	0.392	0.166
Composite	0.36, 0.54	0.032	0.017	0.013

Observation 7.4 *Defending against a single regime pays well below its cap*

For the regimes that reach the finished good in the deterministic networks, the break-even caps lie between 0.013 and 0.077, each less than a fifth of the regime's cap. A smaller exposed value raises the break-even cap, so export control's break-even cap is 0.392 in Modular and 0.166 in Complex.

7.3. Regime effects

Section 7.2 changed the design as the tolerance tightened. This section keeps three designs fixed, the deterministic network, the midpoint and the floor design, and asks what each pays under every regime, at that regime's worst scenario for the design. Figure A.2 draws the midpoint and floor designs.

7.3.1. Operational response to each regime

Robustness mainly narrows how much a design's cost depends on which regime comes. The deterministic network's worst shock cost is 5,400 in all three C^H instances, the duty on its whole final good at the bilateral cap.

The floor designs pay more in normal times but are much less sensitive to which regime materialises. In Serial, total cost is 6,771 in normal times and reaches a maximum of 7,028 across the seven regimes, a worst shock cost of only 257 against 5,400 for the deterministic network. The corresponding worst shock costs are 383 for Modular and 452 for Complex.

In 18 of the 21 instances the floor design costs less than the deterministic network under at least four of the seven regimes, and the other three serve demand entirely in C^L . The deterministic network stays cheaper only in normal times and under regimes that either miss it, such as anti-circumvention, or hit the floor design just as hard, such as export control on a raw material both use in the same place.

A qualified alternative protects a design in three ways, which the rest of this subsection takes in turn. It keeps value outside a regime's reach, it lets production move away from a regime once it comes, and

it keeps a scenario from ever becoming the worst case.

Value kept outside a regime's reach

These cost differences come from how the designs protect the value they send across borders, which Table 7.14 shows. Under a regime's worst scenario, each unit of that value ends up in exactly one of three groups: the scenario does not name it, the design moves it away, or it stays exposed. For each design the table adds up the value in each group over the seven regimes and the 20 instances at $q = 4$ with border crossings, and divides by the total of all three groups, so the three shares in a row add up to 100. Only 20 of the 21 instances enter, because Serial serving C^L sends nothing across a border.

Robust designs protect mainly by keeping value out of reach. The share no scenario names rises from about half in the deterministic networks to between 65 and 85 percent at the midpoint and floor designs, while moving value away once a regime comes covers between 5 and 32 percent. The exposed share falls from about half to between 1 and 10 percent.

Table 7.14: How the value each design sends across borders ends up under the worst scenario of each regime, as a share of the total over the seven regimes and the 20 instances at $q = 4$ with border crossings. Flows are counted at their normal-times size.

	Design	Not named	Moved	Exposed
Serial	Deterministic	44	0	56
	Midpoint	65	32	3
	Floor	78	21	1
Modular	Deterministic	47	0	53
	Midpoint	66	26	8
	Floor	72	25	3
Complex	Deterministic	51	0	49
	Midpoint	85	5	10
	Floor	83	12	5

Production moved to the alternative

Table 7.15 shows the second channel for the Serial midpoint and floor designs, giving where the closing step produces under every regime. Both hold two closing locations and differ only in the non-market one, C^T at the midpoint and allied C^M at the floor design. The midpoint moves its bulk out of C^T under six of the seven regimes, because nearly all of them reach it. The floor design keeps its bulk in allied C^M under five of the seven, because most regimes spare it. Same number of locations, opposite behaviour, so the difference is where the alternative sits. The design that exercises recourse more often is the one at the looser tolerance: more frequent reallocation is not evidence of greater robustness.

Table 7.15: Output of the closing step by location under each realisation, Serial midpoint and floor designs, the rung 3 and rung 4 networks of Table 7.11. The first four steps stay in C^L under every realisation. Volumes are units of final good.

Realisation	Midpoint, rung 3		Floor design, rung 4	
	in C^T	in C^H	in C^M	in C^H
Baseline	95	5	95	5
Bilateral escalation	5	100	5	95
Trade defence	5	95	95	5
Market protection	5	95	5	95
Friendshoring	5	95	95	5
Anti-circumvention	5	100	95	5
Export control	95	5	95	5
Composite	5	100	95	5

The table also explains why friendshoring costs the midpoint most: it still pays 1,589.6 in duty, against at most 180 under any other regime and 78.4 for the floor design. Under friendshoring the Serial midpoint

shifts 95 units of final production into C^H . Its first four steps remain in C^L , so the intermediate feeding the closing facility now enters the market from a non-allied origin and is itself taxed. The response that protects the design against several other regimes therefore creates the exposure friendshoring names, and Modular's midpoint has the same weakness.

The table also shows when paying duty on a forced flow becomes dearer than replacing it. The minimum-throughput constraint (5.6) makes the midpoint's C^T facility produce at least five units under every realisation, carrying 500 of customs value into the market. The firm can ship those units and pay duty, or reproduce them at the already qualified C^H facility and discard the C^T output, which the output constraint (5.3) allows because a location may ship less than it produces. Replacing one unit remakes it from the raw material onwards and ships it again, for 48.84 in total (Appendix A.1). Against a customs value of 100, replacement is cheaper than import once the rate exceeds 0.488, which is why C^H produces the full 100 units under bilateral escalation, anti-circumvention and the composite, which tax those five units at the cap of 0.54. The mechanism holds wherever an alternative is already qualified and surplus output may be discarded, while the threshold belongs to this calibration.

Scenarios kept out of the worst case

A committed alternative can also keep a scenario from ever becoming the worst case. A regime's worst case is the scenario that costs the design most after it has responded. An alternative that makes one scenario cheap to absorb therefore removes it from the worst case, without ever running at scale. Complex under export control shows this. One eligible scenario targets the strategic component leaving C^M , 95 units with a customs value of 1,425, and another a raw material worth 333. The floor design merges in C^H , next to final assembly, and holds its merge facility in C^M at its minimum. If the component were targeted, the design would move the merge to the component's origin and remove almost all of that crossing, so the smaller raw-material scenario becomes the worst case.

A government could still target the component. The design would then absorb it at a cost below the reported worst case, provided the facility can scale up when the measure arrives, which the single period assumes. Because the C^M facility runs at its minimum both in normal times and in the worst scenario, its value shows in no reported flow, only in the scenario it neutralises. Which scenario is worst thus depends on the committed design, the reason the reference cost is set per regime rather than per scenario.

Observation 7.5 *Where the alternative sits matters more than how often it is used*

A qualified alternative protects a design in three ways: it keeps value outside a regime's reach, it lets production move away from a regime once it comes, even replacing output a facility must still make, and it keeps a scenario from ever becoming the worst case while running only at its minimum, because the design would move production to it if that scenario occurred.

7.3.2. Premium and protection along the frontier

Each transition from one rung to the next adds premium and buys protection, and this section asks how much. Every rung design on the frontiers of the three C^H instances is held fixed and exposed to every regime at its worst scenario. Its shock cost under a regime is what that regime adds to its cost, its realised cost under the regime less its own baseline cost. Table 7.16 compares, for each transition, the premium it adds, in units of final-good value, with two savings. Worst saved is how much the largest shock cost across the seven regimes falls. Total saved is how much the sum of shock costs over the seven regimes falls. Total saved is usually the larger, because one transition lowers what several regimes cost at once, but it can be smaller or even negative when a transition that lowers the worst case makes other regimes dearer. The sum is a descriptive total, not an expected saving, since the regimes carry no probabilities.

Across all fifteen transitions, the added premium is smaller than the worst shock cost saved, so moving further along these frontiers continues to buy substantial protection. The return is not linear. Single transitions save between 2.1 and 401 times the premium they add, because different transitions require different changes to the committed design. Summed over each frontier (Table 7.16), each unit of premium removes 8.2 units of worst shock cost in Serial, 17.4 in Modular and 9.3 in Complex. Over the three frontiers together the ratio is about ten, a figure that belongs to these three instances and this calibration.

Complex's fifth-to-sixth transition shows the cost of moving linked production steps together. It qualifies two additional facilities to break the sub-assembly crossing, adds 208 in premium and lowers the worst shock cost by 435, a ratio of 2.1. The transitions on either side add less than 50, with ratios of 10.1 and 3.8. The expensive transition is where linked production steps must gain new locations together.

Modular's fourth-to-fifth transition shows the opposite case. It moves one five-unit closing facility from C^T to allied C^M and changes nothing else. With a C^M facility, finishing can occur on the allied side instead of drawing intermediates from non-allied C^L , reducing friendshoring shock cost from 1,696 to 352 and the worst shock cost by 1,080. Because commitment costs are identical across blocs in the main calibration, the transition adds only 2.7, the additional production cost of shifting five units from C^T to C^M .

Table 7.16: Premium added and shock cost saved at each transition on the frontiers of the three C^H instances, in units of final-good value. Worst saved is the fall in the largest shock cost across the seven regimes, total saved the fall in their sum. Ratios are savings divided by premium.

	Transition	Premium	Worst saved	Total saved	Worst ratio	Total ratio
Serial	1 to 2	309.2	1,710	7,939	5.5	25.7
	2 to 3	184.3	1,929	5,983	10.5	32.5
	3 to 4	133.0	1,504	3,148	11.3	23.7
Modular	1 to 2	123.1	1,710	7,983	13.9	64.9
	2 to 3	102.2	739	2,816	7.2	27.5
	3 to 4	21.7	1,256	2,995	57.8	137.8
	4 to 5	2.7	1,080	1,742	401.2	646.9
	5 to 6	39.0	233	1,425	6.0	36.5
Complex	1 to 2	130.5	702	11,162	5.4	85.6
	2 to 3	44.5	1,458	-1,458	32.7	-32.7
	3 to 4	76.4	1,440	6,723	18.9	88.0
	4 to 5	29.7	300	300	10.1	10.1
	5 to 6	207.9	435	269	2.1	1.3
	6 to 7	40.7	154	651	3.8	16.0
	7 to 8	4.8	460	516	95.8	107.5

Complex's second-to-third transition shows that worst-case and aggregate protection are different objectives. It lowers the worst shock cost by 1,458 but raises the sum of shock costs by the same amount, because moving the third stream's last step into C^T removes the limiting exposure and creates an anti-circumvention exposure on that step's output leaving the hub. The regret criterion accepts the transition because it lowers the largest regret. Figure 7.3 shows this redistribution between regimes rung by rung.

7.3.3. Responses by regime type

Floor designs still pay duty under some regimes, and whether a design can remove that duty depends on what the regime names. The seven regimes fall into three types. Trade defence, the composite, market protection and anti-circumvention name one commodity. Bilateral escalation and friendshoring name every commodity on a relation, and export control names a supplied input. Across the 21 floor designs, each type calls for its own response, which gives three rules.

1. **If a regime names one commodity, shift that step's output to a qualified alternative outside the regime's reach.** Trade defence and the composite are cleared in all five floor designs they reach, and market protection in 10 of the 14 it reaches. The other four still pay on the minimum output of a foreign finishing facility, because at the cap of 0.30 replacing that output costs more than the duty.
2. **If a regime names every commodity on a relation, the final good can leave the relation, but something else still crosses it.** Bilateral escalation and friendshoring each reach 20 floor designs and are cleared in none, although 13 and 15 of them shift production in response. What still pays is never the final good but an input that must cross the relation anyway: mostly the five

units of intermediate feeding or leaving a facility held at its minimum output, and otherwise a larger intermediate flow, a raw material or the strategic component from its fixed origin. Responding therefore lowers what a relation-wide regime costs, but in none of these designs does it remove the regime's duty entirely.

3. **If a regime names a fixed-origin supplied input, move the consuming step to the supplier's bloc.** Export control reaches 13 floor designs and is cleared in none, and only one moves the consuming step at all. The move pays only when the input carries enough value. A raw material worth 333 rarely justifies relocating the steps that use it, while the strategic component, worth 1,425 in Complex, is valuable enough for the floor design to hold a second merge in its origin. The input itself cannot be rerouted to another origin, and Section 7.4.1 tests how supplied inputs, through their origins and their value, shape robustness.

Several regimes can reach the same movement, so one relocation can clear several regimes at once. That is why trade defence is cleared wherever it reaches a floor design, although no frontier is ever built against it alone.

Where exposure remains, the reason differs. A facility held at minimum throughput may need an input that can only cross the taxed border, as with the allied facility of the Serial floor design under bilateral escalation, whose five units need the intermediate from C^L across the taxed pair. The committed design may hold no alternative at the relevant step, as with the raw materials of Modular and Complex under export control. Or an alternative may exist but cost more to use than the duty it avoids, as with the Serial floor design under market protection. Appendix A.3 draws the floor designs of the three C^H instances under every regime that taxes them.

Observation 7.6 *The escape depends on what the regime reaches*

Commodity-specific exposure can often be removed by shifting one step's output to a qualified alternative. Relation-wide exposure tends to remain on another movement across the relation, while fixed-origin input exposure requires moving the consuming step to the supplier's bloc. What remains exposed at the floor design reflects forced operation, a missing alternative, or an alternative that costs more to use than the duty it avoids.

7.4. Conditions shaping robustness

This section asks what determines how much tariff exposure the firm must protect against, how much of it a different design can remove, and what prevents further improvement. It varies three conditions in turn: supplier origins, the number of locations each production step may use, and demand geography.

7.4.1. Supplier origins and input dependence

The first condition is the supplied inputs, where they come from and how much of the product's value they carry. Eight of the 21 instances at $q = 4$ are premium-bound and thirteen are exposure-bound (Table A.5). Every exposure-bound floor design still has a supplied input crossing a bloc border, and no premium-bound floor design does. Export control, which taxes a supplied input leaving its origin, collects duty in all thirteen exposure-bound instances but binds in only twelve. In the Modular instance serving C^M , market protection sets the larger regret through the minimum output of its finishing facility in C^H .

The pattern across architectures follows from whether production can be aligned with supplier origins. Serial buys only its raw material, from C^L , where most production is already located, and all seven Serial instances are premium-bound. Modular and Complex draw their supplied inputs from several blocs, which makes complete co-location harder. The exception is the Complex instance serving concentrated demand in C^M : although its supplied inputs originate in C^L , C^T and C^M , the floor design places each consuming step at its input's origin, so no supplied input crosses a bloc border and the instance is premium-bound.

Raw-material origin flexibility

A first experiment frees the raw materials' origins: each can be bought from any bloc at unchanged value. This removes most of the residual exposure, the floor minus the premium. Ten of the thirteen

exposure-bound instances become premium-bound (Table 7.17), and the effect is largest in Modular, whose floors fall most. Complex moves from one premium-bound instance to six, yet its mean floor barely falls. Its floors already lay close to their premiums, and in two instances the reference network gains more from free sourcing than the committed design, which raises the floor. Table A.5 reports every instance.

Table 7.17: Effect of raw-material origin flexibility at $q = 4$, as means over the seven demand orientations.

Architecture	Premium-bound instances		Mean floor		Mean residual exposure	
	fixed	flexible	fixed	flexible	fixed	flexible
Serial	7 of 7	7 of 7	0.083	0.083	0.000	0.000
Modular	0 of 7	5 of 7	0.066	0.036	0.028	0.001
Complex	1 of 7	6 of 7	0.084	0.081	0.009	0.003
Total	8 of 21	18 of 21	0.078	0.066	0.012	0.001

Three instances remain exposure-bound, but none because of a raw-material origin. Complex serving C^L still imports its strategic component from its fixed origin in C^M . In the two Modular instances, the remaining exposure comes from a finishing facility held at its minimum output. Its regret was already there under fixed origins, but it was smaller than the premium and so never set the floor. Origin flexibility lowers the premium below that regret, so the finishing facility's exposure becomes binding. The experiment assigns no cost to supplier switching, certification or contracting, so it measures the potential value of flexible sourcing rather than its net value.

Supplier dependence

A second experiment asks when dependence on a fixed-origin supplier becomes important enough to change production geography. The strategic component is increased from 5 to 30 percent of final-good value under three absorption modes. The component's value is taken out of in-house production under production absorption, out of the margin under margin absorption, and half out of each under the half mode (Table A.7). Production absorption therefore reduces the production-cost penalty of moving to a dearer bloc. Commitment costs are unchanged.




Figure 7.4: Robustness floor and qualified merge locations as strategic-component share increases in the Complex C^H instance. Open markers indicate that the merge is not qualified in the component's origin bloc C^M . The main calibration is the production-absorption series at 15 percent.

Figure 7.4 traces the floor and the merge locations as the component's share rises, and the floor first rises, then falls. At a share of 5 percent the cheapest design keeps the merge in C^L , so the component crosses a border and that crossing sets the floor, 0.074 under production absorption. As the share grows, the same crossing carries more value and the floor rises. From 7 percent under production and half absorption, and from 9 percent under margin absorption, the design also qualifies the merge in the supplier's bloc C^M , an alternative it switches to when the component is targeted.

The floor peaks between 7 and 15 percent across the modes, at 0.096 under production absorption at 15 percent, which is the main calibration. There the component is valuable enough to hurt but not valuable enough to justify concentrating production at the supplier, so the firm pays for the alternative and still carries residual exposure. Beyond about 15 percent the floor follows the premium. Under production absorption the two are equal at 18, 20 and 24 percent and never more than 0.006 apart, so as the premium falls from 0.089 at 18 percent to 0.046 at 30 percent, the floor falls with it, from 0.089 to 0.052. The other modes behave alike (Table A.7).

Only production absorption makes protection cheaper. There the component replaces the firm's own production, less of which has to sit in a dearer bloc, and the premium in cost terms falls from 535 at a 15 percent share to 286 at 30 percent. Under margin absorption the firm pays the dearer component out of its margin while the final good's value stays at 100, and the premium stays near 650 in cost terms. The deterministic cost of the chain rises from 7,748 to 9,338, of which 1,500 is the component itself and 90 shipping, which is charged on embodied cost and so rises with it. The same premium therefore falls from 8.3 to 7.1 percent of that cost.

At roughly a quarter of product value, between 22 and 28 percent across the modes, the merge moves entirely to the supplier's bloc and the component no longer crosses a border, so export control on it has nothing left to tax. Where a supplier's origin is flexible, as for the raw materials in the first experiment, the firm can instead move the source towards production.

Observation 7.7 *Supplied inputs decide what remains at the floor*

In these calibrations, what remains exposed at the floor design is mostly a supplied input crossing a border. Where the input's origin can move, the firm moves it towards production, which turns ten of the thirteen exposure-bound floors premium-bound. Where it cannot, the firm moves production towards the supplier as the input gains value: it accepts the exposure while the input is cheap, qualifies a location near the supplier at a few percent of product value and concentrates production there at about a quarter. The floor is highest in between, and protection becomes cheaper in cost terms only when the input replaces the firm's own production.

7.4.2. Production-location flexibility

This experiment asks the question of Section 7.1.4, how much geographic flexibility a network needs, by varying location rights from $q = 1$, where each step has only one location, to $q = 4$, where the constraint cannot bind. Table 7.18 reports the floor design at each value by its premium, its floor and its worst and total shock cost, as in Section 7.3.2. It also counts the instances whose floor design changes from the preceding value of q .

The floor should not be compared across values of q . A larger q also improves the reference networks, and in seven instances a reference network improves faster than the committed design, so the floor rises with q (Table A.6). Premium and shock cost do not depend on the reference costs and remain comparable.

Table 7.18: Floor designs by location rights, averaged over the seven demand orientations. Premium and shock costs are in units of final-good value, and Changed counts the instances whose floor design differs from the preceding value of q .

	q	Premium	Worst shock	Total shock	Floor	Changed
Serial	1	575.3	1,902.3	3,438.9	0.099	
	2	505.2	412.6	745.2	0.082	6
	3	506.7	286.0	558.0	0.083	1
	4	506.7	286.0	558.0	0.083	0
Modular	1	386.2	1,464.0	3,259.1	0.083	
	2	253.7	508.9	1,125.6	0.069	6
	3	238.1	450.7	1,050.5	0.066	3
	4	238.1	450.7	1,050.5	0.066	0
Complex	1	306.7	1,409.1	3,082.0	0.088	
	2	434.9	758.5	1,559.8	0.087	6
	3	465.2	691.1	1,341.7	0.084	2
	4	465.2	691.1	1,341.7	0.084	0

A second location makes protection cheaper as well as stronger. With one location per step, a design cannot move production once a regime arrives, so it can protect itself only through where it produces in normal times: at $q = 1$ the floor designs move between three and ten steps out of C^L , always into a dearer bloc. Moving to $q = 2$ changes 18 of the 21 floor designs. In Serial and Modular a second line at the closing step takes over that protection, so the upstream steps return to C^L , whose share of production rises from 51 to 85 percent in Serial and from 64 to 92 percent in Modular. The production this saves exceeds the commitment of the extra line in every instance where the premium falls, so the mean premium falls by 70 in Serial and 133 in Modular while the mean worst shock cost falls from 1,902 to 413 and from 1,464 to 509. The exceptions are the instances serving C^L , where nothing changes, and the dispersed instances, whose floor design at $q = 1$ is already the deterministic network in C^L . The premium counts the extra line through its commitment, a cost per period, and leaves out the up-front investment.

Complex pays more for its second location, 128 on average, while its mean worst shock cost falls by about 650. An extra line there costs 2.5 times what moving one step from C^M back to C^L saves, because Complex spends 20 rather than 35 on its own production while its commitment budget is the same 15. The merge, which takes the fixed-origin strategic component, also stays outside C^L , so the added commitment, 175 on average, outweighs the production saved, 39.

Placement also matters. The second location is added at the closing step in 16 of the 21 instances, with 12 inside a served market and four in allied C^M . Eleven instance-regime pairs move from paying duty at $q = 1$ to clearing the regime at $q = 2$, all involving trade defence, market protection or the composite, as rule 1 of Section 7.3.3 predicts for regimes that name one commodity.

The gains fall sharply after the first alternative. The move from $q = 1$ to $q = 2$ delivers between 91 and 94 percent of the reduction in mean worst shock cost that unrestricted location rights achieve. Moving from $q = 2$ to $q = 3$ changes only six floor designs, and no design changes between $q = 3$ and $q = 4$. Third locations still appear in six instances serving multiple markets, where distinct market exposures justify another route. The main value of production-location flexibility therefore comes from the first well-placed alternative, with further locations becoming useful mainly when the design faces several distinct exposures.

Observation 7.8 *The second location creates recourse and can lower the premium*

A second location per step lets the design move production once a regime arrives, and the first one delivers between 91 and 94 percent of the reduction in worst shock cost that unrestricted location rights achieve. In Serial and Modular it also lowers the mean premium, because a second line at the closing step takes over the protection and lets the upstream steps return to the low-cost bloc. Where an extra line costs more than the production it saves, as in Complex, protection becomes stronger but dearer.

7.4.3. Demand geography

Figure 7.5 shows, for every architecture and demand orientation, the paying band: the range from the robustness floor to the tolerance at which the deterministic network first becomes admissible, over which the firm pays to alter its production network.

Each line runs from a cell's floor out to the tolerance at which the deterministic network becomes admissible




Figure 7.5: The paying band by architecture and demand orientation at $q = 4$. Each line runs from the robustness floor, marked with its premium, to the tolerance at which the deterministic network becomes admissible. Ticks indicate distinct committed designs along the frontier.

The effect of demand geography on the unprotected network is large. Mean deterministic exposure falls from 0.72 when demand is concentrated in C^H or C^M to 0.51 across the dual-market orientations and to 0.25 under dispersed demand. The last comparison also places part of demand in C^L and therefore combines market dispersion with production-demand co-location.

The cleaner comparison is within the dual-market cases, which serve the same two markets. Moving from an 80/20 split, averaged over which market is primary, to an even 50/50 split reduces deterministic exposure. The reduction is 43 percent for Serial, 39 percent for Modular and 28 percent for Complex. Diversification therefore works by limiting the share of output that a regime aimed at one market, or at the relation into it, can reach. Because bilateral escalation binds the deterministic network in 18 of the 21 instances, concentrating sales in one market leaves a correspondingly large share of output

exposed to a single relation.

Which market is served matters most in Complex. Its strategic component and one raw material come from C^M , so when C^M is also the market, these inputs travel from C^M to production in C^L and the final good travels back. Bilateral escalation on that pair taxes both directions: the final good's 10,000 of customs value, as when serving C^H , plus 1,500 for the component and 333 for the raw material. The deterministic exposure, the right end of each band in Figure 7.5, therefore rises from 0.665 to 0.836. Serial buys nothing from C^M and stays at 0.684. In Modular the raw material adds a little duty, but its best defence becomes dearer, and the exposure falls slightly, from 0.734 to 0.724.

The robustness floor responds much less. Excluding concentration in the cheap production bloc, floors, each measured against its own reference costs, vary by only a factor of 1.16 for Serial, 1.14 for Modular and 1.37 for Complex across demand orientations. Dispersion does shorten the paying band and the frontier: moving from concentrated C^M demand to dispersed demand reduces the number of rungs from 6 to 4 in Serial, 6 to 4 in Modular and 11 to 6 in Complex (Figure 7.5), and Appendix A.6 shows what each of those rungs pays under its binding regime.

Observation 7.9 *Demand geography determines how much exposure must be protected*

Spreading demand across markets sharply reduces the exposure of the deterministic network, because a regime aimed at one market, or at the relation into it, reaches only the share routed there. The floor changes much less, because the exposure the floor design keeps comes mainly from production steps and supplied inputs. Demand dispersion therefore substitutes for part of production-network flexibility, but not for all of it.

Taken together, the three experiments act on different parts of the exposure.

Observation 7.10 *Three margins of robustness*

Demand geography sets how much exposure enters the problem. Location rights set how much of it the firm can move once a regime is known, with the first alternative doing most of the work and, in Serial and Modular, lowering the premium. Supplier origins account for much of what remains, and the floor is set by the exposure that is cheaper to keep than to remove, not by the tariff rate alone.


7.5. Managerial implications

The experiments give seven priorities for a firm that designs a production network under trade-policy uncertainty. Each follows from a mechanism of the model, and the numbers show how large it is in this calibration.

1. **Value a location by the responses it makes possible.** A location protects through what the network can do once a measure hits, not through what it produces in normal times. A facility in the market lets the firm respond to any measure on foreign final goods entering it (Section 7.3.1). An alternative near a supplier can protect while running at its minimum: because the firm could switch to it, a measure on that input stops being the worst case. What a measure names decides the response. Shift output away from a named commodity, accept that a named relation can only be reduced, and move production to the supplier of a valuable named input (Section 7.3.3).
2. **Check what a move exposes, not only what it removes.** A relocation shifts exposure rather than removing it: moving final production into a market brings its inputs across the border, and routing through a neutral hub creates anti-circumvention exposure (Section 7.2.4). What a measure costs depends on the value it reaches more than on its rate, so map the crossings and the value on them before choosing a move (Section 7.2.1).
3. **Compare the cost of a defence with the duty it avoids.** The cheapest defence against a measure costs less than the deterministic network as soon as the duty the measure would charge exceeds the defence's extra cost. For a regime known in advance that reaches the final good, that happens from a break-even cap of 0.013 to 0.077, far below the regimes' caps (Section 7.2.5), so the decision does not hinge on the exact rate. Export control on a cheap supplied input has a far higher break-even cap, up to 0.392 here. The comparison is per period: it counts the extra facility through its commitment cost, not the up-front investment needed to build it or the cost of

moving an existing network.

4. **Add a second location where exposure enters.** Exposure concentrates where flows meet: a merge or closing step combines inputs that often come from different blocs, so several crossings a measure can name converge on it, and a step that uses a supplied input from another bloc adds one more. That is where 155 of the 164 moves happen (Section 7.2.3). A second location there delivers 91 to 94 percent of the drop in worst shock cost that full freedom of location gives (Section 7.4.2). In Serial and Modular it also makes protection cheaper, because the rest of the chain can move back to the low-cost bloc. More locations pay only for a new, separate exposure, and a taxed flow can only be moved where the alternative was set up in advance.
5. **Sell in more than one market.** This lowers exposure without changing production at all, because a measure aimed at one market, or at the trade into it, only reaches the output sold there. Selling 50/50 to two markets instead of 80/20 cuts the exposure of an unprotected network by 28 to 43 percent, so less protection is needed (Section 7.4.3). Once the network is protected, where it sells matters far less, because what remains comes from production and supply.
6. **Decide on suppliers and production together.** If a supplied input crosses a border, either move the source to production or move production to the source. Free choice of raw-material origin removes most remaining exposure (Section 7.4.1). For an input with a fixed origin, the response grows with its value: first accept the exposure, then set up an alternative near the supplier, and at about a quarter of product value move production there.
7. **Remove only the exposure worth removing.** Even the best-protected design keeps some exposure, for one of three reasons: a facility must keep running on an input that crosses a taxed border, there is no alternative at the step that matters, or the alternative costs more than the duty it avoids (Section 7.3.3). Each points to a different remedy: supply the facility differently, add an alternative, or keep paying the duty because that is the cheaper option. The frontier shows what each further step of protection costs and what it buys, and where to stop is the firm's choice.



Case study: the solar photovoltaic value chain

This chapter applies the framework to a crystalline silicon photovoltaic manufacturer serving the United States. The United States is chosen because Trina Solar and Hanwha Qcells both built footprints to serve it under the same measures, and because its measures and the sector's costs are documented in public sources (U.S. Department of Commerce, International Trade Administration, 2023; Basore and Feldman, 2022). The formulation, exposure rules, caps and criterion are unchanged. The case supplies production options and regional costs, calibrated from public sources or declared as assumptions where none exist (Appendix B), and the scenarios its locations generate. It tests whether the mechanisms of Chapter 7 hold where geography, costs and trade measures are documented (RQ1 and RQ3), establishes what robustness costs a photovoltaic manufacturer and where further protection stops paying (RQ4), asks what producing without China costs, and sets the committed designs beside the footprints the two firms built.

8.1. Case framework

8.1.1. Case setting

The decision maker is a multinational crystalline silicon photovoltaic manufacturer, a single firm rather than the whole industry that Cui et al. (2025) plan under a known ban (Section 2.3.1). The industry combines sequential production with concentrated manufacturing and trade between steps, a concentration the International Energy Agency (IEA) identifies as a supply-chain vulnerability (International Energy Agency, 2022).

The firm may qualify each step at any location in the main production geography that supplies the United States, and the model chooses among every combination of the resulting facilities. That geography runs from China, where almost all upstream inputs are made, through Vietnam, Malaysia, Thailand and Korea, the main importers of Chinese wafers and the source of over 60 percent of traded cells (International Energy Agency, 2022), to the United States itself, and it shifted towards Southeast Asia as United States measures expanded (Pitt et al., 2026). Europe enters only as a polysilicon supplier. Which steps each location can perform follows the documented operating or committed capability of Trina Solar and Hanwha Qcells, whose plants together cover this geography (Table 8.1), and industry evidence supplies the economic coefficients. The two firms return only after solving, as footprints set beside the designs.

Concentration is greatest upstream. China holds 98 percent of global ingot capacity and 97 percent of wafer capacity, against 81 percent of cell capacity and 77 percent of module capacity (Basore and Feldman, 2022). These are 2021 shares of industry capacity, not of the firm's production, and they do not enter the model (Section 8.1.3). Figure 8.1 shows the chain. The case starts where the firm buys polysilicon and ends at the module.




Figure 8.1: The crystalline silicon photovoltaic production chain. Adapted from International Energy Agency (2022).

8.1.2. Production structure

The firm buys polysilicon and performs wafer, cell and module production. Ingot growth and wafer sawing form one step because their capital and labour requirements are reported jointly (Basore and Feldman, 2022). Each step consumes one unit of its predecessor per unit of output, with every quantity counted in gigawatts of finished modules. The remaining module inputs are bought on world markets at prices that vary little between assembly locations (International Energy Agency, 2022), so they present no location decision and sit in c_{ip}^{prod} .

Polysilicon procurement defines the upstream boundary of the firm. Polysilicon trades as a commodity at a market price (International Energy Agency, 2022), and as a supplied input it is the upstream input on which export control acts (Section 4.2). China, the European Union and the United States are admitted as documented producing regions (International Energy Agency, 2022), representing adversarial, allied and domestic supply relative to the served market. Supplier quantities are unbounded and sourcing can adjust after a regime is realised. Malaysia, also a documented producer, is omitted because it offers no cheaper or less exposed supply to any admitted wafer location.

Table 8.1 records the admitted facilities and their evidence. The feasible set covers capability operating or committed by 2024, while the cost coefficients are 2022 values. This matters for the American wafer and cell steps, since the source records no American ingot, wafer or cell production in the calibration year (Basore and Feldman, 2022), and for Vietnamese wafering, which began at Thai Nguyen in August 2023 (Trina Solar, 2023). Wafering has two alternatives to China against five for each downstream step, so the choice of location is widest for cells and modules.

Table 8.1: Admitted facilities and their evidence. A tick marks a step documented as operating or committed by 2024 at that location, and the evidence column says which. CN China, VN Vietnam, MY Malaysia, TH Thailand, KR Korea, US United States.

Location	Wafer	Cell	Module	Evidence
CN	✓	✓	✓	Integrated wafer, cell and module production (Trina Solar, 2024a)
VN	✓	✓	✓	Ingot, wafer, cell and module production at Thai Nguyen (Trina Solar, 2023, 2024b)
MY		✓	✓	Cell and module production (Hanwha Solutions, 2022)
TH		✓	✓	Cell and module production at Rayong (Trina Solar, 2016, 2024a)
KR		✓	✓	Cell and module production (Hanwha Solutions, 2022)
US	✓	✓	✓	Committed integrated ingot, wafer, cell and module capacity (Qcells, 2023)

Europe has no row in Table 8.1. Its polysilicon already supplies Chinese wafering (International Energy Agency, 2022), while its module plants depend almost entirely on imported cells (International Energy Agency, 2022) and produced around 2 GW in 2023 (Chatzipanagi et al., 2024), for a market that imports 84 percent of the modules it installs (International Energy Agency, 2022). For a firm serving the United States, Europe therefore offers no cell capability and no module output beyond its own market.

8.1.3. Economic parameterisation

Table 8.2 lists the economic inputs of the solar case, each with its value and source. Industry costs replace the value shares of Chapter 7, and inputs without case data keep that chapter's conventions. Production costs are per GW of realised output. Commitment costs are per period for a qualified one-gigawatt facility and equal the IEA's depreciation per GW produced at full use, so they need no annualisation.

Table 8.2: Economic inputs of the solar case, with units in the value column.

Parameter	Notation	Value	Source
Demand	D	1 GW of modules in the United States	Reference scale
Production cost	c_{ip}^{prod}	Table 8.3, MUSD per GW of output	IEA stage cost (International Energy Agency, 2022) less commitment cost
Commitment cost	f_{ip}^P	Table 8.3, MUSD per period for a one-gigawatt facility	IEA depreciation (International Energy Agency, 2022) allocated by DOE investment shares (Basore and Feldman, 2022)
Polysilicon price	c_{ig}^{supply}	98 MUSD per GW at every origin	IEA price (International Energy Agency, 2022) times DOE intensity (Basore and Feldman, 2022). Common price assumed
Gross margin		15 percent of customs value	NREL benchmark (Smith et al., 2021)
Gross values	v_g	Table 8.4, MUSD per GW	Chinese reference cost plus margin (World Trade Organization, 1994a). Polysilicon at purchase price
Shipping cost	c_{ijg}	3% of the shipping base across a border, 1% within a country, 0 between co-located steps	As in Chapter 7
Minimum output	ρ	0.05 GW per qualified facility, 5% of demand	As in Chapter 7
Capacity bound	K_{ip}	1.5 GW, 1.5 times demand	Assumption, never binds
Baseline tariff	τ_{ijg}^0	0 on wafers, cells and modules, 4 percent on silicon into China	Bilateral effectively applied rates in WITS (World Bank, 2024), Table B.4
Regret tolerance	θ	Swept, with bisection between values	Choice (Section 8.1.5)
Regime caps	$\bar{\Delta}$	0.20 to 1.00 of customs value	Chapter 4
Location rights	q	Not restricted	Set by the admitted geography (Table 8.1)

Stage costs by location The IEA reports stage costs and a regional decomposition into cost types (International Energy Agency, 2022). Regional depreciation serves as the proxy for the commitment cost f_{ip}^P . It is allocated across steps in proportion to their equipment investment (Basore and Feldman, 2022), which is an assumption, and the remainder of each stage cost gives c_{ip}^{prod} , so the two sum back to the reported stage cost. No geographic multiplier is applied on top, because regional differences already enter the stage coefficients. The cost rules and model terms of Table 7.5 apply unchanged, and only the coefficients differ. Section B.3 gives the allocation in full, its limits and two reconciliations against the IEA, and Table B.1 gives the coefficients at the stored precision.

Table 8.3: Production cost c_{ip}^{prod} in MUSD per GW of output and commitment cost f_{ip}^P in MUSD per period for a one-gigawatt facility, by location. Blank cells mark steps a location does not admit.

Location	c_{ip}^{prod}			f_{ip}^P		
	Wafer	Cell	Module	Wafer	Cell	Module
CN	28.1	34.1	124.5	4.9	4.9	3.5
VN	28.9	36.9	127.6	6.1	6.1	4.4
MY		36.9	127.6		6.1	4.4
TH		36.9	127.6		6.1	4.4
KR		46.6	124.6		14.4	10.4
US	35.1	49.1	130.4	11.9	11.9	8.6

Vietnam, Malaysia and Thailand share one Southeast Asian calibration, and shipping does not vary with distance, so as module locations the three are identical in cost. Each is a separate neutral jurisdiction that every regime treats alike, so a design naming one of them should be read as naming an ASEAN hub.

The split between the two costs matters most at module assembly, where production costs differ by less than five percent across locations while commitment costs differ by a factor of three. Module assembly draws on processed materials traded globally, which limits cost differences between locations (International Energy Agency, 2022), and the depreciation that does differ moves into f_{ip}^P . Production costs differ more upstream, with Korean cells at 46.6 against Chinese cells at 34.1.

Gross values Duties apply to gross commodity values. Polysilicon is bought, so its value is its purchase price of 98. The three produced commodities have no observed transaction value, so they take a computed value, the method the WTO valuation agreement provides for that case. It adds a gross margin of 15 percent of the resulting value to the cumulative cost along the Chinese reference chain (World Trade Organization, 1994a). The margin is the one at which the NREL photovoltaic benchmark defines a minimum sustainable price (Smith et al., 2021). The values therefore rise in step with cost along the chain, so the case does not use the smile property on which Chapter 7 sets its values (Section 2.1.1). As in Chapter 5, each commodity has one value wherever it is produced, so locating a step in a dearer country raises its production cost without changing its customs value. Section B.2 sets out the valuation rule and why observed spot prices could not be used.

Table 8.4: Gross values v_g , the base of every ad valorem duty, in MUSD per GW and as a share of the module's value.

Commodity	Gross value v_g	Share of module value
Polysilicon	98.00	28.0 percent
Wafer	154.12	44.0 percent
Cell	200.00	57.0 percent
Module	350.59	100.0 percent

Shipping Shipping is a charge per unit on a second base, the supplied input plus the production cost accumulated along the Chinese reference chain, fixed by commodity as the customs value is. It leaves out commitment cost, which is fixed per facility, and the margin, so it sits below the customs value. The rates are those of Chapter 7, 3 percent of the base across a border, 1 percent within a country and nothing between co-located steps (Section 7.1.3). Solar freight falls mainly on modules and is negligible for wafers and cells (Goodrich et al., 2013), so the uniform rule overstates what moving cells costs. The United States is a production node and a customer node separately, so a domestic module delivery is a movement at the domestic rate. Table B.3 gives both bases at every commodity.

Baseline tariff The baseline tariff is the bilateral effectively applied rate of Section 5.2.2. On wafers, cells and modules it is zero. WITS reports zero on these products for every importer in the network from every partner in 2020 to 2023 (World Bank, 2024), since they fall under the Information Technology

Agreement, whose participants bind and eliminate their duties on them (World Trade Organization, 1996). High-purity silicon falls outside the agreement, and the only positive rate WITS reports on a modelled lane is 4 percent on silicon from the United States into China (Section B.4).

Every tariff above this baseline comes from a measure, and the case represents measures through its regimes at the caps of Chapter 4 (Section 8.1.4). Measures on solar trade change in rate and reach, as when the United States extended its orders to completion in Southeast Asia in 2023 (U.S. Department of Commerce, International Trade Administration, 2023), so the design is chosen against every scenario the regimes generate, whatever measures are in force when it is committed.

Scale and capacity Demand is one gigawatt of modules at the American customer node, a reference volume rather than a sales estimate, and its geography is fixed, so the case varies production and sourcing within one served market. Each qualified facility has a minimum output of 0.05 GW, and surplus output may be discarded. Scaling demand, the reference facility, the minimum output and the capacity bound together changes every term of the objective by the same factor, so the designs and every premium, floor and regret do not depend on the reference volume. At one gigawatt the coefficients read directly, since one dollar per watt is a thousand MUS\$ per gigawatt.

The capacity bound of 1.5 GW per facility cannot bind, since no facility needs to produce more than total demand. As in Chapter 7, the case therefore asks where to qualify facilities and does not establish whether the selected locations could take that volume in practice. Commitment costs differ by location, so location and capacity decisions interact (Verter and Dincer, 1995) and the results hold for the one-gigawatt reference facility. Location rights are not restricted. The admitted geography already determines which facilities each step can open (Table 8.1), so the case adds no limit on how many of them a step may use.

8.1.4. Regime instantiation

The rules and caps are those of Chapter 4, and the case supplies each rule's commodity scope, its anchors, the protected group and the alignment of the jurisdictions it reaches. Relative to the United States market, Korea and the European Union are allied, Vietnam, Malaysia and Thailand are neutral, and China is adversarial. Korea's classification rests on its security alliance with the United States and on their trade agreement (Office of the United States Trade Representative, 2012).

Three regimes have documented solar counterparts, the WTO trade measures and the extension that Section 2.2.1 describes. Trade defence corresponds to the anti-dumping and countervailing duties imposed on Chinese solar cells and modules in 2012 (Table 2.1), market protection to the 2018 safeguard (Bown, 2021) and anti-circumvention to the 2023 findings on Southeast Asian completion (U.S. Department of Commerce, International Trade Administration, 2023).

Table 8.5 gives the resulting scenarios, with "into the United States" covering movements into the production node and into the market. Bundle codes b_0 to b_7 follow Table 6.1.

Table 8.5: Scenarios of each bundle in the solar case. Cap is the rate as a share of customs value, and Count is the number of scenarios in the bundle.

Bundle	Anchor	Scenarios	Cap	Count
b_0 Baseline	none	one, no regime in force, as in the deterministic model	none	1
b_1 Bilateral escalation	country pair	one per unordered pair of the seven jurisdictions, all commodities, either direction, less three reaching no feasible movement	0.54, 0.25 if allied	18
b_2 Trade defence	origin and market	one per adversarial or neutral module origin into the market	0.36	4
b_3 Market protection	market	one, all foreign module origins into the market	0.30	1
b_4 Friendshoring	market	one, non-allied origins, all commodities, into the United States	0.20	1
b_5 Anti-circumvention	hub and market	one per hub and protected-group commodity, modules or cells, into the United States	0.54	6
b_6 Export control	origin	one per polysilicon origin, polysilicon leaving that origin for every other country	1.00	3
b_7 Composite	as in b_2 and b_5	one per pair of trade-defence and anti-circumvention measures protecting the same market	0.36 and 0.54	24

Four choices are specific to this chain. The lower bilateral cap of 0.25 applies to the two allied pairs, Korea and the United States, and the European Union and the United States. The anti-circumvention hubs are Vietnam, Malaysia and Thailand, and Cambodia appears in the 2023 proceeding but is not a production option here. The protected group of Chapter 4 holds the finished good and the input to its final step, here modules and cells, because an extension can follow a firm that finishes modules in a hub and a firm that makes cells there and finishes the modules in the United States. Each anti-circumvention scenario names either modules or cells leaving one hub, six scenarios in all. Export control designates polysilicon and generates one scenario per origin. The composite puts one trade-defence and one anti-circumvention measure in force together, which gives $4 \times 6 = 24$ joint scenarios, and the uncertainty set holds $33 + 24 + 1 = 58$ scenarios in all.

The caps are Chapter 4's, and this sector has since run well above them, the 2025 determinations reaching country maxima far beyond any rate in the regime set (ING, 2025). Section 8.3 sets out where the regimes' reach differs from the measures in force.

8.1.5. Experiments

The chapter solves the model deterministically, computes a reference network for each regime, sweeps the regret tolerance, and evaluates committed designs against every realisation. Two controls then vary one feature of the instance, each with its own scenario libraries and recomputed reference costs. The first cuts the supplier origins to China alone, asking whether alternative origins matter. The second excludes Chinese production, asking what leaving it costs. Section 8.2.6 reports both.

The premium compares a design's normal-times cost with the deterministic network, regret compares its worst-case cost under a regime with that regime's reference cost, and the binding regime is the one with the largest regret. Relative quantities are reported to three decimals. With a single instance, the bisection that Chapter 7 uses to locate the floor is also applied between every pair of solved tolerances, which refines the frontier down to the bisection resolution.

The implementation and solver settings are those of Chapter 7. Every logged solve behind a reported objective or design terminated at proven optimality within the configured gap. Section C.2 gives the solve accounting.

8.2. Results

8.2.1. Reference networks

The deterministic network The deterministic network performs every step in China and exports finished modules, at a baseline cost of 307.52 MUSD per GW. Table 8.6 gives the components, and every qualified facility runs at the full reference volume.

Table 8.6: Components of the deterministic network's baseline cost, in MUSD per GW and as a share of that cost.

Component	Cost	Share
Supplied polysilicon	98.00	31.9 percent
Production	186.76	60.7 percent
Commitment cost	13.24	4.3 percent
Shipping	9.52	3.1 percent
Duty	0	0.0 percent
Total	307.52	100 percent

The design crosses a border only when the finished module leaves China. Polysilicon is consumed at its origin and the production steps are co-located. No solution in this chapter sends silicon into China, because Chinese wafer production buys Chinese polysilicon at the same price without crossing a border, so the 4 percent on silicon into China never applies. That single crossing carries the full customs value of 350.59 per GW, so every regime reaching it taxes the same movement, and no qualified alternative exists to reroute through.

Reference networks Each regime's reference cost C_b^* is the lowest worst-case cost any committed design can reach against that regime alone (Section 6.3.1). The design that attains it is the regime's reference network. Every reference cost is computed with each regime charging its full cap (Table 8.5), and the caps are never varied. Table 8.7 gives each reference cost and its increment over the deterministic cost C_0^* of 307.52, what defending against the regime costs, including any duty the reference network still pays. Every reference network keeps wafer and cell production in China, so the table shows only where each reference network assembles modules.

Table 8.7: Reference cost C_b^* of each regime and its increment over C_0^* , with every regime at its cap, in MUSD per GW, and the module assembly of each reference network with its baseline output in GW. The baseline row gives C_0^* .

Bundle	Regime	C_b^*	Increment	Reference module assembly
b_0	baseline	307.52		China (1.00)
b_1	bilateral escalation	328.91	21.39	China (0.95), ASEAN (0.05)
b_2	trade defence	317.64	10.11	United States (1.00)
b_3	market protection	317.64	10.11	United States (1.00)
b_4	friendshoring	319.33	11.81	Korea (1.00)
b_5	anti-circumvention	307.52	0.00	China (1.00)
b_6	export control	307.52	0.00	China (1.00)
b_7	composite	317.64	10.11	United States (1.00)

The increment ranges from nothing to 21.39 MUSD per GW. Anti-circumvention and export control cost nothing, since the deterministic network uses no hub and makes its wafers next to its polysilicon supplier. Trade defence, market protection and the composite share one defence, domestic assembly.

Friendshoring has the lowest cap but the second-highest increment, because it names every commodity entering the United States from a non-allied origin, the Chinese cells of domestic assembly included. Korean assembly escapes it, since its cells stop in Korea and its modules enter from an allied origin. It costs 1.70 more than domestic assembly, because shipping modules across the border, 5.70, and Korea's dearer commitment, 1.76, outweigh its cheaper production, 5.76.

8.2.2. Design changes along the frontier

The robustness floor The floor design attains a largest regret of 0.066 at a premium of 6.6 percent, 20.42 MUSD per GW in cost terms. It is premium-bound. Its binding regimes, anti-circumvention and export control, collect no duty and their reference costs equal the deterministic cost, so its largest regret is its premium.

In normal times it keeps Chinese wafer and cell production and assembles 0.95 GW of modules in Korea and 0.05 GW in the United States. The two locations answer different regimes, an allied origin for trade defence and friendshoring and domestic delivery for market protection, so the floor design holds both. This is the floor design of Serial in Chapter 7, bulk at an allied location and a minimum facility in the market, and it responds the same way, moving 0.90 GW from Korea to the United States under bilateral escalation and market protection and nothing under the other regimes (Table 7.15). Figure 8.2 sets it against the deterministic network.




Figure 8.2: Committed facilities and baseline flows in GW, for the deterministic network and for the floor design. Rings mark a facility held at its minimum output.

The frontier Table 8.8 reports the eight designs and their attained regret. Each is optimal from its own attained regret up to the next looser rung's.

Table 8.8: The frontier of committed designs, loosest first. Only module assembly is shown, with baseline output in GW, because wafer and cell production stay in China throughout. π is the premium as a share of the deterministic network's cost, and regret is the largest regret, attained under the binding regime.

Rung	Module assembly (GW)	π	Regret	Binding
1	CN 1.00	0.000	0.511	bilateral escalation
2	CN 0.95, ASEAN 0.05	0.016	0.391	composite
3	CN 0.95, US 0.05	0.029	0.316	bilateral escalation
4	US 1.00	0.033	0.294	bilateral escalation
5	CN 0.90, ASEAN 0.05, US 0.05	0.044	0.154	friendshoring
6	ASEAN 0.05, US 0.95	0.048	0.139	friendshoring
7	CN 0.90, KR 0.05, US 0.05	0.063	0.076	market protection
8	KR 0.95, US 0.05	0.066	0.066	anti-circumvention, export control

Figure 8.3 gives the frontier in cost terms. Each bar shows a rung's cost in normal times and under the regime that binds it, so the bars are not a like-for-like comparison between rungs. The cost in normal times rises from 307.52 to 327.94 MUSD per GW, while the cost under the binding regime falls from 496.84 to 327.94.




Figure 8.3: Cost of every rung design in normal times and under its binding regime, in MUSD per GW, loosest tolerance on the left. The binding regime is named above each bar.

Two thirds of the floor design's premium removes four fifths of the fall in regret from the deterministic exposure to the floor, the pattern Snyder and Daskin (2006) report, where large cuts in maximum regret cost little (Section 2.3). Rung 5 sits at 0.044 and 0.154 against the floor design's 0.066 and 0.066, so the last fifth costs a third of the premium. No design qualifies more than one ASEAN hub, at any rung or in any reference network. A second hub would add no capacity, since capacity never binds, and no escape from friendshoring or market protection, which name every hub at once, so the premium goes on Korea and the United States instead.

No design on the frontier holds more than three module locations, and from rung 5 further protection changes which of them it holds rather than adding a fourth, as in Chapter 7.

Table 8.9 decomposes the premium. Procurement does not change and no duty is paid in normal times, so commitment, production and shipping account for all of it.

Table 8.9: Premium added at each transition on the frontier, in MUSD per GW, split into commitment, production and shipping. The last row is the floor design's premium in cost terms. Each transition is rounded separately, so the transitions sum to within 0.02 of it.

Transition	Move	Commitment	Production	Shipping	Total
1 to 2	duplication into ASEAN	4.39	0.16	0.24	4.79
2 to 3	relocation from ASEAN to the US	4.22	0.14	-0.28	4.07
3 to 4	consolidation out of China	-3.51	5.61	-0.84	1.25
4 to 5	duplication into China and ASEAN	7.91	-5.45	1.08	3.53
5 to 6	consolidation out of China	-3.51	5.32	-0.80	1.00
6 to 7	duplication into China, relocation from ASEAN to Korea	9.49	-5.46	0.80	4.82
7 to 8	consolidation out of China	-3.51	0.13	4.33	0.95
	Floor	15.46	0.44	4.52	20.42

Commitment accounts for 76 percent of the floor design's premium, shipping 22 percent and production 2 percent. The transitions come in three pairs, rungs 3 and 4, 5 and 6, and 7 and 8, each looser design keeping Chinese module assembly and each tighter one dropping it. Dropping it saves its commitment cost of 3.51 but gives up the cheapest module production in the set, and it removes a module crossing

valued at 350.59 per GW, leaving a cell crossing at 200.00. Where the bulk goes decides the price. It moves to the United States at rungs 4 and 6, where production cost rises by 5.61 and 5.32, and to Korea at the floor design, where module production costs almost the same as in China and the rise is 0.13. All seven moves change module assembly, the closing step, as most moves do in Chapter 7 (Observation 7.2), and the frontier never moves cell production. No regime requires it, because module assembly in Korea and the United States already limits the duty on Chinese cells to minimum flows, and cells are cheapest in China, by the widest margin in the chain.

8.2.3. Break-even caps

The break-even cap is calculated as in Chapter 7 (Section 7.2.5). It is the cap from which a regime's reference network costs less than the deterministic network.

The deterministic network pays the cap on its one module crossing, 350.59 per GW, under every regime that reaches it, and anti-circumvention and export control reach nothing. The break-even cap is therefore the increment of Table 8.7 divided by 350.59. Against bilateral escalation, for example, the reference network costs 21.39 more, so the break-even cap is $21.39/350.59 = 0.061$. This is an upper bound, because that reference network still pays 9.47 in duty. Table 8.10 gives every break-even cap beside the regime's cap.

Table 8.10: Break-even caps beside each regime's cap, as shares of customs value. None means the regime reaches nothing in the deterministic network. The bilateral escalation entry is an upper bound.

Regime	Cap	Break-even cap
Bilateral escalation	0.54	0.061
Trade defence	0.36	0.029
Market protection	0.30	0.029
Friendshoring	0.20	0.034
Anti-circumvention	0.54	none
Export control	1.00	none
Composite	0.36, 0.54	0.029

Every break-even cap lies between 0.029 and 0.061, each less than a fifth of the regime's cap, as in Chapter 7 (Observation 7.4). Above it, the reference network costs less. That is a duplicate module facility in an ASEAN hub against bilateral escalation, domestic assembly against trade defence, market protection and the composite, and Korean assembly against friendshoring.

The floor design, which defends against every regime at once, costs less than the deterministic network against a single regime that reaches the module crossing from a cap of about 0.058, its premium of 20.42 over the 350.59 that crossing carries.

8.2.4. Solar value chain strategies

The frontier follows five strategies, concentrated production at rung 1, an ASEAN hub at rung 2, domestic assembly at rungs 3 and 4, an ASEAN hub with domestic assembly at rungs 5 and 6, and allied and domestic assembly at rungs 7 and 8. Each combines the defences of Table 8.7 and escapes the regimes by the rules of Chapter 7 (Observation 7.6), and its binding regime shows what the combination still leaves exposed. Figure 8.4 shows four of them in normal times and under a regime that reaches them.

An ASEAN hub Adding an ASEAN module facility is the cheapest first move, at 4.79 MUSD per GW, and the same move the record shows after the 2012 orders, when assembly relocated to Southeast Asia (Table 2.1). It cuts bilateral duty from 189.32 to 9.47, since the hub lies off the China and United States pair. The composite then binds. Trade defence names Chinese modules and anti-circumvention names hub modules, so both module origins are taxed at once, and neither regime binds alone because each leaves the other origin free. The hub is taxed at the higher cap, 0.54 against 0.36, so the firm keeps bulk assembly in China and the composite collects 129.37.




Figure 8.4: Four strategies in normal times, left, and under the regime named, right, with flows in GW and a ring on each taxed movement.

Domestic assembly Domestic assembly answers trade defence, market protection and the composite, the defence Table 8.7 finds for all three, and bilateral escalation binds again. The American facility draws its cells from China, so the module leaves the relation but the cells still cross it. The duty falls from a module base of 350.59 to a cell base of 200.00 but is never removed, as in Chapter 7 (Observation 7.6). At rung 3 bilateral escalation still collects 112.07, of which 102.60 falls on those cells and 9.47 on the minimum output of the Chinese module facility.

An ASEAN hub with domestic assembly Holding both locations answers bilateral escalation and market protection together, and friendshoring binds. It reaches both, the hub as a non-allied origin and the American facility through its Chinese cells, so moving the bulk between them changes only what friendshoring taxes, hub modules at 350.59 or Chinese cells at 200.00 per GW. At rung 5 it collects 43.01, of which 36.00 falls on the Chinese cells feeding the American facility and the rest on the two module minimums the design still ships. Rung 6, which has dropped Chinese module assembly, binds on friendshoring for the same reason.

Allied and domestic assembly Korean assembly answers friendshoring, so market protection binds at rung 7, which still keeps Chinese module assembly, and the floor design, which drops it, is premium-bound. Friendshoring's duty falls from 43.01 at rung 5 to 2.00 at the floor design.

8.2.5. The solar floor design under each regime

Worst shock cost The floor design cuts the worst shock cost from 189.3 MUSD per GW for the deterministic network to 5.5, about thirty-five to one, for a premium of 20.42 in cost terms. What it still pays falls on minimum-output flows, as in Chapter 7 (Observation 7.6). Bilateral escalation and friendshoring tax the Chinese cells that the Korean and the American facility consume to meet their minimum output, 5.40 and 2.00. Market protection taxes the Korean facility's minimum output, 5.26, because remaking it in the United States would cost 14.54.

The floor design's option to switch The option to switch between its Korean and American module facilities protects the floor design, rung 8. When a regime reaches one of them, the bulk sits at or moves to the other. Only bilateral escalation and market protection make it move, 0.90 GW of assembly from Korea to the United States in each case. Which facility carries the bulk in normal times therefore matters little. Cost settles it, and Korea wins by 0.06 MUSD per GW, because its lower production cost and higher shipping almost cancel. A cross-border shipping rate of 3.02 percent instead of 3 would put the bulk in the United States without changing how the floor design meets any regime.

Value kept out of reach Table 8.11 splits the value each rung sends across borders as Table 7.14 does in Chapter 7, and divides the moved value by whether the regime still reaches it after the shift. Each of the last three strategies spans two rungs, and the two protect in different ways. The looser rung keeps Chinese module assembly, the cheapest in normal times, and protects mainly by moving production once a regime names it, 60 to 65 percent of its crossing value at rungs 3, 5 and 7. The shift is never complete. The Chinese facility must still produce its minimum output, which keeps paying duty, and the moved volume pays the dearer production of the alternative. The tighter rung drops Chinese module assembly for about 1 MUSD per GW more in normal times (Table 8.9) and keeps 71 to 86 percent of its value out of reach. Keeping the bulk where the regimes reach it therefore leaves a residue that recourse cannot remove, and the tighter rung pays a small premium to avoid it.

At rungs 3 and 5 part of the moved value stays within reach, 26 and 12 percent of their crossing value, because moving assembly to the United States brings the Chinese cells across the relation that bilateral escalation or friendshoring names. Rung 7, with Korea as an alternative, moves the same value fully out of reach. Counting the moved value still named, rungs 3 and 5 leave about as large a share within reach as rungs 4 and 6, which hold the same cells from the start.

In the terms of Bruckler et al. (2024), the looser rung protects through adaptive capacity and the tighter rung through absorptive capacity (Section 2.3). Tightening the tolerance within a strategy therefore replaces protection that depends on moving production after a regime arrives with protection built into the design, which holds even if moving production takes longer than the single-period model allows.

Table 8.11: How the value each rung sends across borders ends up under the worst scenario of each regime, as a share of the total over the seven regimes, in percent. Moved value is still named when its production moves to a location the same regime reaches through its inputs. Flows are counted at their normal-times size, and the crossing value differs by design, so the shares compare how each design's exposure is composed, not its size.

Rung	Strategy	Not named	Moved		Exposed
			Out of reach	Still named	
1	Concentrated production	29	0	0	71
2	ASEAN hub	31	25	0	44
3	Domestic assembly	30	39	26	5
4	Domestic assembly	71	0	0	29
5	ASEAN hub with domestic assembly	33	48	12	7
6	ASEAN hub with domestic assembly	81	0	0	19
7	Allied and domestic assembly	35	60	0	5
8	Allied and domestic assembly	86	13	0	1

8.2.6. Leaving China and the reference firms

Every design on the frontier makes its wafers and cells in China. This section asks two questions this leaves open. Can a firm serving the United States produce without China, and at what cost? Does the model find the strategies of the two firms that built footprints for this market?

Leaving China A firm serving the United States can produce without China. The cheapest network without Chinese production performs every step in Vietnam, at 319.48 MUSD per GW, 3.9 percent above the deterministic network, since Vietnam is admitted at every step and the chain moves together. Protected without China, the floor design costs 335.90, 9.2 percent above the base deterministic network against 6.6 percent for the base floor. The two floor designs differ only upstream. Both assemble the bulk in Korea with an American minimum, and the one without China makes its wafers and cells in Vietnam, 6.00 dearer, and imports its polysilicon, adding 1.96 of freight. Leaving China therefore adds 7.96 per GW, 2.6 percentage points, to the cost of protection. These figures rest on one IEA reading, which places Southeast Asian stage costs about five percent above China's, and on capacity that never binds, while China holds 97 percent of global wafer capacity (Section 8.1.1). On these costs, leaving China is affordable wherever capacity allows.

The exit lands in a hub. Vietnam is one of the hubs the case's anti-circumvention names, whatever the origin of their inputs, so that regime's reference cost rises most, by 19.77, and the floor design without China answers it by assembling the bulk in Korea from Vietnamese cells.

Export control without China While wafering stays in China, polysilicon never crosses a border, so the other two origins are never used, and restricting supply to China leaves every reference cost and the floor design unchanged. Leaving China creates the crossing, since Vietnam has no polysilicon supplier. With other origins the design switches when export control names its supplier, at no cost, because the model gives suppliers no commitment and one price. With Chinese polysilicon alone, a Chinese export control at its cap of 1.00 adds the full value of the polysilicon, 98 MUSD per GW, to every admitted design, nearly twice the 15 percent margin of 52.59 a module carries. It still does not bind, because the reference network pays the same duty, and regret measures loss against the best design under the same regime (Chapter 6).

Trina and Qcells The frontier finds the strategies of both firms that built footprints for this market. Trina combined production in China and Vietnam with cells and modules in Thailand and a module plant in Texas, sold in December 2024 (Trina Solar, 2024a; FREYR Battery, 2024), which is an ASEAN hub with domestic assembly, as at rungs 5 and 6 (Section 8.2.4). Qcells combines cells and modules in Korea and Malaysia (Hanwha Solutions, 2022) with wafer, cell and module production in the United States (Qcells, 2023), which is the floor design's allied and domestic assembly with a Malaysian hub added, closest to the floor design without China. The match is one of mechanism. Capacity never

binds, so the case cannot judge how many facilities either firm holds, and every frontier design keeps wafers and cells in China, where Qcells makes its own (Section 8.2.2).

8.3. Case synthesis

For a solar PV manufacturer serving the United States, the floor design costs 20.42 MUSD per GW more under baseline tariffs, 6.6 percent of the deterministic network's baseline cost, and its worst shock cost falls from 189.3 to 5.5. It keeps wafer and cell production in China and assembles modules in Korea, with a small American facility at its minimum.

What carries over The solar floor design has the same shape as the floor design of Chapter 7's Serial chain (Figure 8.2 and Table 7.11). Most module assembly sits in an allied location, a small facility in the market runs at its minimum and takes the bulk under two regimes (Section 8.2.5), and further improvement stops at the premium. Serial was constructed to isolate these mechanisms, while the solar case was built from industry costs and the documented supply geography and reaches the same shape, so the shape does not rest on the constructed costs of Chapter 7, although both instances share the formulation and its conventions. Table 8.12 sets each finding of Chapter 7 against its solar evidence.

Table 8.12: Chapter 7's findings in the solar case.

Finding	Solar evidence	Result
Reach, more than rate, determines what a regime costs (Observation 7.1)	The two regimes that name every commodity cost most to defend, bilateral escalation at a cap of 0.54 and friendshoring at the lowest cap, 0.20, while anti-circumvention, also capped at 0.54, costs nothing (Section 8.2.1)	holds
Defending against a single regime pays well below its cap (Observation 7.4)	Break-even caps lie between 0.029 and 0.061, under a fifth of every cap (Section 8.2.3)	holds
Exposure identifies where robustness is valuable (Observation 7.2)	Every change along the frontier is at module assembly, and wafers and cells stay in China (Section 8.2.2)	holds
Maximum-regret protection is not uniform (Observation 7.3)	Anti-circumvention collects nothing from the deterministic network or the floor design but reaches every design between them that holds a hub (Section 8.2.4)	holds
Where the alternative sits matters more than how often it is used (Observation 7.5)	The floor design keeps 86 percent of its crossing value out of reach, against 78 in Serial, and its American facility runs at its minimum yet takes the bulk under two regimes (Section 8.2.5)	holds
The escape depends on what the regime reaches (Observation 7.6)	Moving assembly to the United States escapes trade defence and market protection, which name the module, but bilateral escalation and friendshoring also name the Chinese cells it imports, and the floor design pays only on minimum outputs (Sections 8.2.4 and 8.2.5)	holds
Supplied inputs decide what remains at the floor (Observation 7.7)	Polysilicon is bought at world prices from any origin, and from a single origin the reference network would pay the same duty (Section 8.2.6)	does not arise
The second location creates recourse and can lower the premium (Observation 7.8)	Location rights are not varied	not tested
Demand geography determines how much exposure must be protected (Observation 7.9)	One served market	not tested

Six findings hold. The supplied-input finding does not arise, because export control on a single origin would tax the reference network equally and so could not bind, and the case does not test the last two, since it fixes location rights and serves one market.

What the case adds Four results go beyond Chapter 7. Within each strategy, the looser rung protects by moving production once a regime arrives and pays a residue, the Chinese minimum output and the dearer production it moves to, while the tighter rung removes that residue for about 1 MUSD per GW (Section 8.2.5). In the floor design, holding two module facilities is what protects, and where the bulk sits in normal times hardly matters (Section 8.2.5). Leaving China adds 7.96 per GW, 2.6 percentage points, to the cost of protection and moves the exposure into the hubs (Section 8.2.6). A dependence no design can avoid leaves no regret, so a Chinese export control on a single polysilicon origin would add 98 per GW to every China-free design without binding (Section 8.2.6).

Limitations The regimes simplify a tariff environment that layers several measures on the same products and changes faster than a fixed set of regimes at fixed caps can follow. Since 2018 the United States has imposed and extended a safeguard on cells and modules (Bown, 2021; Basore and Feldman, 2022), extended its orders on Chinese cells to completion in Southeast Asia (U.S. Department of Commerce, International Trade Administration, 2023), imposed new duties on four Southeast Asian countries (U.S. Department of Commerce, 2025) and raised its Section 301 duty on Chinese cells (Office of the United States Trade Representative, 2024).

Each product takes the origin of the country that last processed it, so a Korean module counts as Korean whatever cells it contains. A measure that followed Chinese content would reach the floor design's Korean assembly, as United States manufacturers asked Commerce to do in June 2026 for Korean cells made from Chinese wafers (Pitt et al., 2026). The case therefore scores the floor design's protection more generously than such a measure would. Anti-circumvention also names one hub at a time, while the 2023 findings named four countries at once (U.S. Department of Commerce, International Trade Administration, 2023). This leaves the frontier unchanged, since every reported design and reference network holds at most one hub.

The case models no production subsidies, since the IEA costs exclude them. The case also serves a single market. Capacity never binds, so the same facilities could serve others, but the effect of spreading demand on robustness (Observation 7.9) lies outside the case.

Discussion and conclusion

This chapter answers the four sub-questions and the main question, using the terms defined in the introduction (Section 1.2). It then states the contribution and relevance of the thesis, its limitations and further research. The findings are mechanisms. The constructed chains isolate them and the solar case demonstrates them on documented costs, so the numbers that illustrate them belong to these two calibrations. The constructed chains come in three architectures. Serial is a single stream of production steps, and Modular joins three streams at the closing step. Complex joins two streams and a purchased strategic component at a merge partway along the chain, and a third stream at the closing step. Three of the 21 constructed instances, one per architecture, serve only the high-cost market and are reported in most detail.

9.1. Answers to the research questions

RQ1. What the design determines

How does the configuration of a fragmented value chain determine its exposure to trade measures, and which of that exposure can production-network design remove?

A value chain is exposed through the customs value on its border crossings that a regime can reach. The configuration decides which goods cross which borders at what value, and design can remove most of that exposure. What a regime reaches matters more than its cap. Bilateral escalation against a non-allied country and anti-circumvention both add 0.54 to the tariff of every movement they reach. In the three chains serving only the high-cost market, the reference network for bilateral escalation, the design the firm would choose if it knew that regime, costs up to 12 percent more than the deterministic network. The reference network for anti-circumvention costs nothing extra, because the deterministic networks produce nothing in the neutral hub that anti-circumvention reaches (Observation 7.1). In solar, friendshoring has the lowest cap, 0.20, yet costs more to defend against than trade defence or market protection, because it also reaches the Chinese cells that module assembly in the United States imports (Section 8.2.1).

Design removes exposure mainly by keeping value outside a regime's reach. Under the worst scenario of each regime, each unit of value a design sends across borders is either outside the regime's reach, moved once the measure comes, or exposed. The deterministic networks leave about half of that value exposed, and the floor designs 1 to 5 percent. The floor designs keep 72 to 83 percent outside the regime's reach and move a further 12 to 25 percent once a measure comes. In solar, the exposed share falls from 71 percent for the deterministic network to 1 percent for the floor design (Tables 7.14 and 8.11).

What remains depends on what the regime reaches (Observation 7.6). A regime that names one commodity is usually escaped by an alternative location. A regime that names every good crossing a given border, such as bilateral escalation or friendshoring, usually keeps taxing some input that must still cross it. A regime on a supplied input is escaped only by producing at the input's origin, which pays

only when the input is valuable.

RQ2. What the design is decided against

How can uncertainty about future trade policy be represented for robust production network design when no probabilities can be assigned to it?

The thesis represents this uncertainty as a finite set of regimes, each a documented class of measure. Trade rules fix what each regime can reach, and recorded measures fix whom it targets and how much it charges (Chapter 4). Trade defence (anti-dumping and countervailing duties), market protection (safeguards) and anti-circumvention, which extends trade defence to goods completed in a third country, build on the WTO trade measures. Bilateral escalation (tariffs raised under domestic law), friendshoring (origin conditions attached to preferences and subsidies) and export control (restrictions on exports of critical inputs) follow what governments have done outside them. The composite of trade defence and anti-circumvention makes a seventh regime.

A two-stage model (Atamtürk and Zhang, 2007) judges every design by its regret under each regime separately, which needs no probabilities. A design's regret under a regime is how far its worst-case cost exceeds the regime's reference cost, the lowest worst-case cost any design can attain, as a share of that reference cost. The model accepts only designs whose regret stays within a chosen tolerance under every regime. Among them, it commits the one that is cheapest in normal times.

The set meets the double unknown, which measure comes and how severe it is, by enumerating the first and fixing the second at each regime's cap. The cap is the fixed tariff increment the regime adds in full to the baseline tariff of every movement it reaches. Charging the full cap is safe for cost, since a milder rate can only lower what a committed design pays (Section 4.2.3). Applied to the candidate network, the regimes generate 18 to 63 scenarios in the constructed instances and 58 in solar (Tables A.1 and 8.5).

Regret counts only the cost that a better design could avoid, so a duty that every design must pay does not shape the choice. If a firm producing outside China could buy polysilicon only from China, a Chinese export control at its cap of 1.00 would add 98 MUSD per GW to every design it could commit. That is nearly twice a module's margin. A bound on worst-case cost would let this measure decide the design. Under regret it does not bind, because the reference network pays the same duty (Section 8.2.6). Each design is judged by its own worst scenario under each regime, against the best design for that regime, and its regret is bounded under every regime at once (Section 6.1). A design therefore answers for the exposure it creates as well as the exposure it removes. This matters because anti-circumvention reaches the designs halfway along 12 of the 21 frontiers while missing both ends (Section 7.2.4).

RQ3. How the design changes

How does uncertainty about trade policy change the cost-minimising production network, and what determines the resulting network configuration?

Under trade-policy uncertainty the firm adds or moves locations where the architecture forces a crossing. The bill of materials joins streams of production steps at merges, and streams made in different blocs must cross a border to meet there. The finished good crosses unless it is made in the market, and a supplied input crosses unless it is used at its origin. Value rises at every merge and peaks at the closing step. The firm therefore moves merges and the closing step into the market or an allied bloc and leaves the crossing upstream, on goods worth less. Demand geography sets how much exposure enters, the locations each step may use set how much of it can move, and supplier origins account for much of what remains (Observation 7.10).

In the three chains serving only the high-cost market and in solar, the deterministic network has one location per step. It therefore cannot move production when a measure comes and pays the cap on all the value the regime reaches. Its largest regret, the deterministic exposure, is the tolerance below which the firm commits a different network. Bilateral escalation sets it in all four cases. It is 0.684 in Serial, 0.734 in Modular and 0.665 in Complex. A value of 0.684 means that the deterministic network's worst-case cost lies 68.4 percent above the regime's reference cost. In solar the deterministic exposure is 0.511 (Sections 7.4.3 and 8.2.2).

The break-even cap is an output of the model. It shows how low a regime's cap can be before defending

against that regime stops paying, if the firm knew which regime would come. It is the cap at which the duty the deterministic network pays under the regime equals the amount by which the regime's reference cost exceeds the deterministic network's cost. For a regime that reaches the deterministic network's finished good, it lies between 0.013 and 0.077 in the constructed chains. In solar it lies between 0.029 and 0.061, so in both applications it stays under a fifth of the regime's cap. The solar floor design, which defends against every regime at once, costs less than the deterministic network once any single regime on the module charges about 0.058 (Sections 7.2.5 and 8.2.3). Both comparisons count each extra facility at its per-period commitment cost, without the up-front investment or the cost of moving an existing network.

The 21 constructed frontiers contain 164 changes in where a step is qualified. Of these, 155 fall at a merge, the closing step or a step feeding one, or at a step receiving a supplied input from another bloc (Observation 7.2). The other nine, all in Complex, are steps inside a stream. Each follows a neighbouring step, such as the merge, out of the low-cost bloc into a dearer one. The cheaper production of the low-cost bloc did not outweigh the cost of shipping across the border between the two (Section 7.2.3). In solar, a single stream like Serial, every change falls at module assembly, its closing step (Section 8.2.2).

A regime aimed at one market reaches only the output sold there. Mean deterministic exposure therefore falls from 0.72 when all demand sits in the high-cost or the allied market to 0.51 when demand is split between two markets. It falls to 0.25 when demand is spread over three markets. Between two markets, an even split instead of 80/20 cuts it by 28 to 43 percent (Section 7.4.3). The robustness floor changes much less. Across the demand patterns, except demand served only in the low-cost bloc, it varies by a factor of at most 1.37, because the exposure left at the floor design comes from production and supply (Observation 7.9).

A design's worst shock cost is the most that any regime adds to its cost. Restricted to one location per step, a design can protect itself only by producing in dearer blocs in normal times. Allowing a second location per step achieves 91 to 94 percent of the cut in worst shock cost that allowing every location achieves. In Serial and Modular it also lowers the premium, because the upstream steps can return to the low-cost bloc (Section 7.4.2).

RQ4. What the design costs, and its limit

What does a robust network cost in normal trading conditions and under the trade measures it may face, and where does further protection stop paying?

A robust network, here the floor design, costs more in normal times, but its cost hardly moves with the measures. Costs in the constructed chains are in units of final-good value for 100 units of demand. In the three chains serving only the high-cost market, the floor design cuts the worst shock cost from 5,400 to between 257 and 452. Its premium is 4.7 to 10.2 percent. In solar, the floor design cuts the worst shock cost from 189.3 to 5.5 MUSD per GW for a premium of 6.6 percent. Facility commitment makes up 76 percent of that premium, shipping 22 percent and production 2 percent (Sections 7.3.1, 7.2.3, 8.2.5 and 8.2.2).

At the full caps, far above the break-even caps, the floor design costs less than the deterministic network under at least four of the seven regimes. This holds in every constructed instance except the three that serve only the low-cost bloc. There the deterministic network already produces inside its only market, so regimes can reach only its supplied inputs. Elsewhere the deterministic network stays cheaper only in normal times and under regimes that miss it or hit the floor design as hard (Section 7.3.1).

On the frontiers of the three chains serving only the high-cost market, each change of design removes 2.1 to 401 times as much worst shock cost as it adds in premium. Over a whole frontier, a unit of premium removes 8.2 units of worst shock cost in Serial, 17.4 in Modular and 9.3 in Complex. In solar it removes about nine (Sections 7.3.2 and 8.2.5).

Further protection keeps paying up to the floor design, where the remaining exposure is cheaper to keep than to remove. It stops there for one of two reasons. A floor is premium-bound when the regimes under which the floor design's regret is largest collect no duty. Its largest regret then equals its premium, and only a design cheaper in normal times could lower it. A floor is exposure-bound when a regime still

collects duty that is cheaper to keep than to remove, so the largest regret stays above the premium.

Eight of the 21 constructed floors and the solar floor are premium-bound. Each of the other 13 keeps a supplied input crossing a bloc border. When raw materials can be bought from any bloc, 10 of these 13 become premium-bound (Section 7.4.1). The other three stay exposure-bound for a reason that buying raw materials from any bloc cannot remove. In Complex serving the low-cost bloc, the strategic component still crosses a border, because it can be bought only from its one fixed origin. In two Modular instances, a finishing facility must still produce its minimum output, and the duty on that output sets the largest regret. That regret was already there, but it stayed below the premium until buying raw materials from any bloc lowered the premium. In all three, removing the remaining exposure would raise the premium by more than it lowers the largest regret.

The exposure a floor design keeps comes from minimum output a facility must produce, from a step without an alternative, or from an alternative that costs more than the duty it avoids (Observation 7.6). Low production cost also ties a step to its location. The Serial floor design keeps its first four steps in the low-cost bloc (Section 7.3.1). The solar floor design keeps wafers and cells in China, where cells are cheapest by the widest margin in the chain. Leaving China would raise its premium from 6.6 to 9.2 percent (Sections 8.2.2 and 8.2.6).

Main question

How should a firm design a multi-stage global value chain so that it remains cost-efficient under normal trading conditions and limits its exposure to trade policy it cannot foresee?

A firm should commit the design that is cheapest in normal times among those whose regret stays within a chosen tolerance under every regime (RQ2). Choosing the cheapest design keeps it cost-efficient in normal times, and the regret bound limits its exposure without probabilities on which measure comes. Such a design holds a few well-placed alternative locations where the architecture forces a crossing and keeps its most valuable crossings outside the regimes' reach (RQ1 and RQ3). It accepts the exposure that is cheaper to keep than to remove, and the firm chooses on the frontier how much regret it will bear (RQ4). The firm can do so without knowing which measure will come or how severe it will be. What a measure costs depends more on the value it reaches than on its rate, and the firm's own network decides how much value a measure can reach (RQ1). Four conclusions support this answer.

1. **Protection belongs where the architecture forces a crossing (RQ1 and RQ3).** Merges, the closing step, the steps feeding them and the steps receiving supplied inputs account for 155 of the 164 design changes along the constructed frontiers. Moving merges and the closing step into the market or an allied bloc leaves the crossing upstream, on goods worth less. Allowing a second location per step achieves 91 to 94 percent of the cut in worst shock cost that allowing every location achieves. In 16 of the 21 instances the model places that second location at the closing step. Mapping the crossings and the value on them therefore shows a firm where its alternatives belong (Section 7.5).
2. **Protection is cheap relative to what it removes (RQ4).** In the constructed chains, the floor design costs around 4 to 10 percent more in normal times than the deterministic network, and in solar 6.6 percent more. Serial and Modular serving only the low-cost bloc need no premium (Figure 7.5 and Table A.5). On the frontiers of the three chains serving only the high-cost market, each unit of premium removes about ten units of worst shock cost (Table 7.16). In solar, three quarters of the premium pays for holding facilities.
3. **An alternative protects through the response it makes possible (RQ1 and RQ4).** A facility running at its minimum output can be the one that protects the design. The American facility of the solar floor design produces only its minimum in normal times, yet takes over most assembly under bilateral escalation and market protection. Moving production once a measure comes, the adaptive capacity of Bruckler et al. (2024), keeps normal-times cost down. It leaves some regret, because the facility the measure reaches must still run at its minimum output and pay duty on it. Keeping value outside the regimes' reach, their absorptive capacity, removes that remaining duty, in solar for about 1 MUSD per GW of extra normal-times cost (Section 8.2.5).
4. **Markets and suppliers decide how much protection is needed (RQ3).** An even split of demand between two markets instead of 80/20 cuts the deterministic exposure by 28 to 43 percent.

Buying raw materials from any bloc turns 10 of the 13 exposure-bound floors premium-bound (Sections 7.4.3 and 7.4.1). Where a firm sells and where it buys therefore belong in the same decision as where it produces.

9.2. Contribution and relevance

The thesis makes the four contributions named in the introduction (Section 1.3).

The first is a rule for choosing a network without probabilities, a min-max p -robust rule. Its min-max half sets, for each regime, a reference cost, the lowest cost any design can guarantee against every scenario of that regime (Equation (6.8)). Its p -robust half admits a design only if its worst-case cost stays within a tolerance of that reference cost under every regime. Among the admissible designs, the rule commits the one that is cheapest in normal times (Equation (6.9)). The rule extends the stochastic p -robustness of Snyder and Daskin (2006), with normal-times cost in place of expected cost and the reference cost per regime in place of the best design per scenario. Regret then prices not knowing which regime will come, without assuming the firm knows which origin a measure will name. Tightening the tolerance traces what each level of protection costs, down to the floor design, whose largest regret is the lowest any design attains. To the author's knowledge, this pairing of a p -robust bound with min-max reference costs and normal-times cost is new.

The second is a way to represent trade-policy uncertainty that follows trade law. The closest tariff study gives tariffs a range through the budgeted sets of robust optimisation (Bertsimas and Sim, 2004), in a model with a single production stage (Ahmadi Dighehsara et al., 2025). This thesis enumerates the regime set from the trade rules. Each regime takes its reach from those rules and its cap from recorded measures. A scenario therefore charges one rate on every movement its legal criterion covers (Section 4.1). The set can be checked against the measures it encodes, at the price of saying nothing about likelihood.

The third is three findings on how a firm can protect itself. The first is that it can do so without forecasting tariff rates, because which movements a measure reaches matters more than the rate it charges. If the firm knew which regime would come, defending against one that reaches the finished good would pay from a cap far below any the regimes charge (Sections 7.2.5 and 8.2.3). The second is that a few alternatives where the architecture forces a crossing deliver most of the protection (Section 7.4.2). This carries two results on supply disruptions over to trade policy. In process flexibility, a few well-placed links between plants and products deliver most of what full flexibility would under uncertain demand (Jordan and Graves, 1995). Here a second location per step achieves 91 to 94 percent of the cut in worst shock cost that allowing every location achieves, and the well-placed links are the steps where the architecture forces a crossing. Contingent rerouting to a backup source is often part of the optimal response to a supply disruption (Tomlin, 2006). Here the alternative location is that backup, running at its minimum output until a measure comes. Both results change in two ways. The disruption is a policy decision whose reach is set by law, so an alternative protects against a regime only if it lies outside that regime's reach. The floor designs therefore place it in the market or an allied bloc. Protection is also judged by worst shock cost and regret under every regime, without the probabilities on which both earlier results rest. The third is that protection keeps paying up to the floor design. There the remaining exposure is cheaper to keep than to remove, and in these calibrations it sits mostly on supplied inputs and forced minimum output (Observations 7.6 and 7.10).

The fourth is a framework that joins trade economics and network design, the two literatures that ask how tariffs change the cost of production spread across borders (Miller et al., 2025). From trade economics it takes the gross-value tariff base, under which duty falls on the full customs value at every crossing (Yi, 2003; Rouzet and Miroudot, 2013; Blanchard et al., 2026). From network design it takes the firm's choice of where to place each step (Melo et al., 2009; Dong and Kouvelis, 2020; Cui et al., 2025), so that the tariff base follows the design, as in Zhou et al. (2017). The framework combines the robust model, the regimes and a simple cost structure. The cost structure consists of a bill of materials, a customs value for each commodity, a production cost and a per-period commitment cost for each step at each location, and a shipping rate. All of them can be stated as shares of the final good's value, as in a common-size income statement (Foster, 1986). Given a chain's candidate locations, supplier origins and markets, the framework returns its frontier, which gives the cheapest design at each level of

protection, and the exposure left at the floor design (Chapters 4, 5 and 6). To the author's knowledge, it is the first to combine multi-stage design, measure-specific uncertain policy and a tariff base that follows the design.

The results also matter outside the literature. For a firm, they turn trade-policy uncertainty into a design question it can answer before any measure comes (Section 7.5). For governments, they show that the reach of a measure decides which production it moves. In both applications, a measure on the finished good is answered by moving the closing step, and upstream production stays where it is cheapest. The relocation of solar module assembly to Southeast Asia after the 2012 orders shows the same pattern (Section 8.2.4). For the design of complex systems, the thesis shows how an institutional environment can enter the design of a technical system without probabilities on political decisions. Here that environment is the trade rules and the measures governments take under them.

9.3. Limitations

Each design is tested against one regime at a time, apart from the composite, while measures accumulate in practice, as the United States measures on solar since 2018 show (Sections 4.3 and 8.3). The tolerance therefore leaves regret under two regimes together unbounded. The caps, taken from several industries so that one set serves every chain, rest on evidence of uneven strength and are not varied, so the results belong to them (Section 4.2.3). Without probabilities, the model cannot say which level of protection is worth its premium in expectation.

Every measure enters as a charge on the customs value of a crossing, with quotas and licences entering through their ad valorem equivalent (Section 2.1). Route closures such as bans, charges on a base other than customs value and duty drawback lie outside the model. A commodity takes the origin of the location that last processed it, while extensions of trade defence can test what a good contains (U.S. Department of Commerce, International Trade Administration, 2023). Modelled duty can therefore be too high or too low. The solar floor design, which assembles Chinese cells into modules in Korea, is scored more generously than such a test would score it (Section 8.3).

The model covers a single period and designs the network from scratch. It leaves out the up-front investment and its financing, the time to build or qualify a facility, the cost of moving an existing network and recovery over time. Production moved after a measure is therefore assumed to move within the period, and the frontier is not an investment sequence. With fixed demand and a cost objective, the model also leaves out how far duty passes on to customers and how a measure shifts the firm's position against competitors.

Both applications start from a baseline tariff of zero on every crossing their designs use, and measures already in force enter only as regimes. In solar this matches the bilateral effectively applied rates WITS reports for 2020 to 2023 (World Bank, 2024), and in the constructed chains it holds by construction (Sections 8.1.3 and 7.1.5). Positive baseline tariffs would already favour designs with fewer crossings.

Shipping is a share of embodied cost and ignores weight, volume, distance and transport mode, so it can understate the gain from making a bulky good close to its market (Section 7.1.3). Unit production costs do not fall with volume. Capacity could be made binding, since the model carries a capacity bound for each step at each location, but both applications set it high enough never to bind (Sections 5.2.5 and 7.1.4). Production can therefore move in full to an alternative once a measure comes, and the results leave open whether capability exists at the scale required. The minimum output of 5 percent of demand per facility is itself one source of the exposure the floor designs keep (Observation 7.6). No systematic sensitivity analysis varies the commitment costs or the minimum output, so the deterministic exposures and break-even caps belong to this calibration.

The constructed chains and their stage values are set by hand, so they support no statistical generalisation. Their blocs make the adversarial bloc the cheapest and the high-cost market the dearest, which sharpens the trade-off between cost and exposure that a cheap allied producer would soften (Section 7.1.2). The solar case covers one industry and one market, with costs from one IEA decomposition, and compares its designs with the footprints of Trina Solar and Hanwha Qcells by mechanism only. It demonstrates the mechanisms without validating the designs against observed decisions (Sections 8.1 and 8.3).

9.4. Further research

The first step is to vary the caps, commitment costs, minimum output and stage values. They are parameters of the implementation, so this needs no new model and would show which deterministic exposures, break-even caps and floors hold beyond this calibration. Two assumptions behind the protection the floor designs report need testing next. Tracking the origin of every input would let a measure test what a good contains. Admitting several regimes at once, with a budget on how many come together in the manner of Bertsimas and Sim (2004), would test the floor designs against measures that accumulate.

Three extensions keep the current two-stage model. Starting from an existing network, with a cost for each facility the firm opens or closes, would add the cost of moving a network that already exists, which the break-even caps leave out. Capping how far production at an alternative can rise once a measure comes would show how much protection remains when production cannot move in full within the period. Making the split of demand across markets a first-stage decision, with a cost of entering each market, would test how far spreading sales helps.

Applying the framework to one industry and one firm would test it on that firm's chain structure, costs, capacities and the bilateral effectively applied rates it faces. Value added by stage would come from inter-country input-output tables (Rouzet and Miroudot, 2013), together with customs unit values and a shipping cost based on distance and weight in place of the assumed shares. Caps and targets built from the measures that industry has faced would replace the cross-industry set. Solar already needs this, since its duties in 2025 reached rates far beyond any in the regime set (ING, 2025). Such a study could also check the designs against the footprints that firms such as Trina Solar and Hanwha Qcells built, using their own costs.

Bibliography

- Ahmadi Digebsara, A., Nejati, M., Ardestani-Jaafari, A., Rastani, S., Aflaki, S., 2025. Geopolitical disruptions in global supply chains: The role of strategic alliances. *Transportation Research Part E: Logistics and Transportation Review* 204, 104448. doi:10.1016/j.tre.2025.104448.
- Amiti, M., Redding, S.J., Weinstein, D.E., 2019. The impact of the 2018 tariffs on prices and welfare. *Journal of Economic Perspectives* 33, 187–210. doi:10.1257/jep.33.4.187.
- Anderson, J.E., van Wincoop, E., 2004. Trade costs. *Journal of Economic Literature* 42, 691–751. doi:10.1257/0022051042177649.
- Atamtürk, A., Zhang, M., 2007. Two-stage robust network flow and design under demand uncertainty. *Operations Research* 55, 662–673. doi:10.1287/opre.1070.0428.
- Baldwin, R., Venables, A.J., 2013. Spiders and snakes: Offshoring and agglomeration in the global economy. *Journal of International Economics* 90, 245–254. doi:10.1016/j.jinteco.2013.02.005.
- Barrot, J.N., Loualiche, E., Sauvagnat, J., 2019. The globalization risk premium. *The Journal of Finance* 74, 2391–2439. doi:10.1111/jofi.12780.
- Basore, P., Feldman, D., 2022. Solar Photovoltaics: Supply Chain Deep Dive Assessment. Technical Report DOE/OP-0012. U.S. Department of Energy. Washington, DC. URL: <https://www.energy.gov/sites/default/files/2022-02/Solar%20Energy%20Supply%20Chain%20Report%20-%20Final.pdf>, doi:10.2172/1871588. report prepared in response to Executive Order 14017, America's Supply Chains.
- Ben-Tal, A., Goryashko, A., Guslitzer, E., Nemirovski, A., 2004. Adjustable robust solutions of uncertain linear programs. *Mathematical Programming, Series A* 99, 351–376. doi:10.1007/s10107-003-0454-y.
- Bertsimas, D., Sim, M., 2004. The price of robustness. *Operations Research* 52, 35–53. URL: <https://www.jstor.org/stable/30036559>, doi:10.1287/opre.1030.0065.
- Blanchard, E.J., Boice, M.W., Johnson, R.C., 2026. Effective Protection with Global Value Chains. NBER Working Paper 34670. National Bureau of Economic Research. URL: <https://www.nber.org/papers/w34670>, doi:10.3386/w34670.
- Blanchard, E.J., Bown, C.P., Johnson, R.C., 2016. Global Supply Chains and Trade Policy. NBER Working Paper 21883. National Bureau of Economic Research. Cambridge, MA. URL: <https://www.nber.org/papers/w21883>, doi:10.3386/w21883.
- Bown, C.P., 2021. The US-China Trade War and Phase One Agreement. Working Paper 21-2. Peterson Institute for International Economics. Washington, DC. URL: <https://www.piie.com/publications/working-papers/2021/us-china-trade-war-and-phase-one-agreement>.
- Bown, C.P., Crowley, M.A., 2013. Import protection, business cycles, and exchange rates: Evidence from the Great Recession. *Journal of International Economics* 90, 50–64. doi:10.1016/j.jinteco.2012.12.001.
- Bridgeford, E.W., Campbell, I.D., Chen, Z., Lin, Z., Ritz, H., Vandekerckhove, J., Poldrack, R.A., 2026. Twelve quick tips for AI-assisted coding in science. *PLOS Computational Biology* 22, e1014428. doi:10.1371/journal.pcbi.1014428.
- Bruckler, M., Wietschel, L., Messmann, L., Thorenz, A., Tuma, A., 2024. Review of metrics to assess resilience capacities and actions for supply chain resilience. *Computers & Industrial Engineering* 192, 110176. doi:10.1016/j.cie.2024.110176.


- Ceglowski, J., Golub, S.S., 2012. Does China still have a labor cost advantage? *Global Economy Journal* 12. doi:10.1515/1524-5861.1874.
- Charnovitz, S., Hoekman, B., 2013. US–Tyres: Upholding a WTO accession contract – imposing pain for little gain. *World Trade Review* 12, 273–296. doi:10.1017/S1474745612000602.
- Chatzipanagi, A., Jaeger-Waldau, A., Letout, S., Mountraki, A., Gea Bermudez, J., Georgakaki, A., Ince, E., Schmitz, A., 2024. Clean Energy Technology Observatory: Photovoltaics in the European Union – 2024 Status Report on Technology Development, Trends, Value Chains and Markets. Technical Report JRC139297. European Commission, Joint Research Centre. Luxembourg. URL: <https://publications.jrc.ec.europa.eu/repository/handle/JRC139297>, doi:10.2760/1812909. eUR 40073.
- Chen, Z., Hu, Z., Li, K., 2021. The spillover effect of trade policy along the value chain: Evidence from China's rare earth-related sectors. *The World Economy* 44, 3550–3582. doi:10.1111/twec.13172.
- Cleland, J., Driessen, E., Masters, K., Lingard, L., Maggio, L.A., 2026. When and how to disclose AI use in academic publishing: AMEE Guide No.192. *Medical Teacher* 48, 542–553. doi:10.1080/0142159X.2025.2607513.
- Cui, C., Lonergan, K.E., Sansavini, G., 2025. Policy-driven transformation of global solar PV supply chains and resulting impacts. *Nature Communications* 16, 6742. doi:10.1038/s41467-025-61979-5.
- Damodaran, A., 2026. Margins by sector (US). NYU Stern School of Business, data as of January 2026. URL: https://pages.stern.nyu.edu/~adamodar/New_Home_Page/datafile/margin.html.
- Dong, L., Kouvelis, P., 2020. Impact of tariffs on global supply chain network configuration: Models, predictions, and future research. *Manufacturing & Service Operations Management* 22, 25–35. doi:10.1287/msom.2019.0802.
- European Commission, 2018. Notice of the expiry of certain anti-dumping measures (2018/C 310/06) and Notice of the expiry of certain anti-subsidy measures (2018/C 310/07). URL: [https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52018XC0903\(01\)](https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52018XC0903(01)). official Journal of the European Union, C 310, 3 September 2018, pp. 4–5 (CELEX 52018XC0903(01) and 52018XC0903(02)). Measures on crystalline silicon photovoltaic modules and key components (i.e. cells) originating in or consigned from the People's Republic of China; expired at midnight on 3 September 2018.
- European Commission, 2023. Critical raw materials: Ensuring secure and sustainable supply chains for EU's green and digital future. URL: https://ec.europa.eu/commission/presscorner/detail/en/ip_23_1661. press release IP/23/1661, 16 March 2023. Accessed 30 July 2026.
- European Commission, 2024a. Carbon Border Adjustment Mechanism. URL: https://taxation-customs.ec.europa.eu/carbon-border-adjustment-mechanism_en. directorate-General for Taxation and Customs Union. Accessed 30 July 2026.
- European Commission, 2024b. Commission Implementing Regulation (EU) 2024/2754 of 29 October 2024 imposing a definitive countervailing duty on imports of new battery electric vehicles designed for the transport of persons originating in the People's Republic of China. URL: https://data.europa.eu/eli/reg_impl/2024/2754/oj. official Journal of the European Union, L, 2024/2754, 29 October 2024.
- European Union, 2024. Regulation (EU) 2024/1735 of the European Parliament and of the Council of 13 June 2024 on establishing a framework of measures for strengthening Europe's net-zero technology manufacturing ecosystem and amending Regulation (EU) 2018/1724. Official Journal of the European Union, L 2024/1735, 28 June 2024. URL: <http://data.europa.eu/eli/reg/2024/1735/oj>.
- Fajgelbaum, P.D., Goldberg, P.K., Kennedy, P.J., Khandelwal, A.K., 2020. The return to protectionism. *The Quarterly Journal of Economics* 135, 1–55. URL: <https://www.jstor.org/stable/10.2307/26864946?seq=1&cid=pdf->, doi:10.1093/qje/qjz036.
- Feldmann, A., Olhager, J., 2019. A taxonomy of international manufacturing networks. *Production Planning & Control* 30, 163–178. doi:10.1080/09537287.2018.1534269.

- Foster, G., 1986. *Financial Statement Analysis*. 2nd ed., Prentice-Hall, Englewood Cliffs, NJ.
- FREYR Battery, 2024. FREYR Battery Closes Transformative Acquisition of Trina Solar's U.S. Manufacturing Assets. URL: <https://ir.t1energy.com/news-releases/news-release-details/freyr-battery-closes-transformative-acquisition-trina-solars-us>. investor release.
- Goodrich, A.C., Powell, D.M., James, T.L., Woodhouse, M., Buonassisi, T., 2013. Assessing the drivers of regional trends in solar photovoltaic manufacturing. *Energy & Environmental Science* 6, 2811–2821. doi:10.1039/c3ee40701b.
- Goron, C., 2018. Fighting against climate change and for fair trade: finding the EU's interest in the solar panels dispute with China. *China-EU Law Journal* 6, 103–125. URL: <https://doi.org/10.1007/s12689-018-0080-z>, doi:10.1007/s12689-018-0080-z.
- Handley, K., Limão, N., 2015. Trade and investment under policy uncertainty: Theory and firm evidence. *American Economic Journal: Economic Policy* 7, 189–222. doi:10.1257/pol.20140068.
- Hanwha Solutions, 2022. Green Bond Allocation and Impact Report. Report. Hanwha Solutions. URL: https://www.hanwhasolutions.com/static/en/data/HSC_Allocation_and_Impact_Report_2022.PDF. accessed: 2026-09-10.
- Hou, J., Zeng, A.Z., Sun, L., 2017. Backup sourcing with capacity reservation under uncertain disruption risk and minimum order quantity. *Computers & Industrial Engineering* 103, 216–226. doi:10.1016/j.cie.2016.11.011.
- Hummels, D., Ishii, J., Yi, K.M., 2001. The nature and growth of vertical specialization in world trade. *Journal of International Economics* 54, 75–96. doi:10.1016/S0022-1996(00)00093-3.
- Ignatenko, A., Lashkaripour, A., Macedoni, L., Simonovska, I., 2025. Making America great again? The economic impacts of Liberation Day tariffs. *Journal of International Economics* 157, 104138. doi:10.1016/j.jinteco.2025.104138.
- ING, 2025. What does it mean to have up to 3,521% US tariffs on solar imports? URL: <https://think.ing.com/articles/what-does-it-mean-to-have-up-to-3521-us-tariffs-on-south-east-asian-solar-imports/>. Author: Coco Zhang; ING Think economic and financial analysis.
- Internal Revenue Service, 2026. Guidance to apply interim safe harbors for purposes of determining a taxpayer's material assistance from a prohibited foreign entity; other prohibited foreign entity guidance. Notice 2026-15. URL: <https://www.irs.gov/pub/irs-drop/n-26-15.pdf>. implementing Public Law 119-21, 139 Stat. 72 (4 July 2025). Accessed 1 August 2026.
- International Energy Agency, 2022. Solar PV Global Supply Chains. Special Report. International Energy Agency. Paris. URL: <https://www.iea.org/reports/solar-pv-global-supply-chains>, doi:10.1787/9e8b0121-en.
- Javorcik, B., Kitzmüller, L., Schweiger, H., Yıldırım, M.A., 2024. Economic costs of friendshoring. *The World Economy* 47, 2871–2908. doi:10.1111/twec.13555.
- Jordan, W.C., Graves, S.C., 1995. Principles on the benefits of manufacturing process flexibility. *Management Science* 41, 577–594. URL: <https://www.jstor.org/stable/2632880>, doi:10.1287/mnsc.41.4.577.
- Kowalski, P., Legendre, C., 2023. Raw materials critical for the green transition: Production, international trade and export restrictions. OECD Trade Policy Papers 269. OECD Publishing. Paris. doi:10.1787/c6bb598b-en.
- KPMG LLP, The Manufacturing Institute, 2020. Cost of Manufacturing Operations Around the Globe. Report. KPMG LLP. URL: <https://kpmg.com/kpmg-us/content/dam/kpmg/pdf/2024/cost-of-manufacturing-operations-around-the-globe-whitepaper.pdf>. Accessed 3 August 2026.
- Küpper, D., Nordemann, J., 2025. The new cost calculus for manufacturers. Boston Consulting Group. URL: <https://www.bcg.com/publications/2025/new-cost-calculus-for-manufacturers>. bCG Henderson Institute survey of 1,000 manufacturing executives. Accessed 3 August 2026.

- Liang, D., Wilhelm, W.E., 2013. Dual-ascent and primal heuristics for production-assembly-distribution system design. *Naval Research Logistics* 60, 1–18. doi:10.1002/nav.21515.
- Madansky, A., 1960. Inequalities for stochastic linear programming problems. *Management Science* 6, 197–204. doi:10.1287/mnsc.6.2.197.
- Mancheri, N.A., 2015. World trade in rare earths, Chinese export restrictions, and implications. *Resources Policy* 46, 262–271. doi:10.1016/j.resourpol.2015.10.009.
- Melo, M.T., Nickel, S., da Gama, F.S., 2009. Facility location and supply chain management – A review. *European Journal of Operational Research* 196, 401–412. doi:10.1016/j.ejor.2008.05.007.
- Miller, J.W., Jin, Y.H., Ortega, D.L., 2025. Shock and awe: A theoretical framework and data sources for studying the impact of 2025 tariffs on global supply chains. *Journal of Supply Chain Management* 61, 3–15. doi:10.1111/jscm.12350.
- Mudambi, R., 2008. Location, control and innovation in knowledge-intensive industries. *Journal of Economic Geography* 8, 699–725. doi:10.1093/jeg/lbn024.
- Nizovtsev, D., Skiba, A., 2020. Non-equivalent pass-through of anti-dumping duties, countervailing duties and import tariffs. *The World Economy* 43, 170–190. doi:10.1111/twec.12849.
- Office of the United States Trade Representative, 2012. U.S.–Korea Free Trade Agreement. Office of the United States Trade Representative. URL: <https://ustr.gov/trade-agreements/free-trade-agreements/korus-fta>. Entered into force 15 March 2012. Accessed 10 September 2026.
- Office of the United States Trade Representative, 2024. Notice of modification: China's acts, policies and practices related to technology transfer, intellectual property and innovation. *Federal Register*, Vol. 89, p. 76581 (18 September 2024); FR Doc. 2024-21217; Docket USTR-2024-0007. URL: <https://www.federalregister.gov/documents/2024/09/18/2024-21217/notice-of-modification-chinas-acts-policies-and-practices-related-to-technology-transfer>. section 301 Four-Year Review final modification. Additional duty on solar cells (whether or not assembled into modules) raised to 50%, effective 27 September 2024; solar-cell determination at p. 76585.
- Ossa, R., Redding, S.J., 2026. The Economics of Tariffs. NBER Working Paper 34915. National Bureau of Economic Research. Cambridge, MA. URL: <https://www.nber.org/papers/w34915>, doi:10.3386/w34915.
- Pitt, H., Movalia, S., Jain, K., 2026. Solar Manufacturing in Asia at a Crossroads. Note. Rhodium Group. URL: <https://rhg.com/research/solar-manufacturing-in-asia-at-a-crossroads/>. based on Rhodium Group Clean Investment Monitor data. Accessed 14 September 2026.
- Qcells, 2023. Qcells invests \$2.5 billion in building complete solar supply chain in U.S. URL: <https://us.qcells.com/blog/qcells-invests-2-5-billion-in-building-complete-solar-supply-chain-in-u-s/>. company press release.
- Read, R., 2005. The political economy of trade protection: the determinants and welfare impact of the 2002 US emergency steel safeguard measures. *The World Economy* 28, 1119–1137. doi:10.1111/j.1467-9701.2005.00722.x.
- Rosenhead, J., Elton, M., Gupta, S.K., 1972. Robustness and optimality as criteria for strategic decisions. *Operational Research Quarterly* 23, 413–431. URL: <https://www.jstor.org/stable/3007957>, doi:10.1057/jors.1972.72.
- Rouzet, D., Miroudot, S., 2013. The Cumulative Impact of Trade Barriers along the Value Chain: An Empirical Assessment Using the OECD Inter-Country Input-Output Model. GTAP Resource 4184. Global Trade Analysis Project (GTAP), Purdue University. West Lafayette, IN. URL: https://www.gtap.agecon.purdue.edu/resources/res_display.asp?RecordID=4184. presented at the 16th Annual Conference on Global Economic Analysis, Shanghai, China.

- Rungi, A., Del Prete, D., 2017. The “Smile Curve”: where Value is Added along Supply Chains. IMT Lucca EIC Working Paper Series 05/2017. IMT School for Advanced Studies Lucca. Lucca. URL: <https://ssrn.com/abstract=2936710>, doi:10.2139/ssrn.2936710.
- Savage, L.J., 1951. The theory of statistical decision. *Journal of the American Statistical Association* 46, 55–67. doi:10.1080/01621459.1951.10500768.
- Shapiro, A., Dentcheva, D., Ruszczyński, A., 2009. *Lectures on Stochastic Programming: Modeling and Theory*. SIAM. doi:10.1137/1.9780898718751.
- Shin, N., Kraemer, K.L., Dedrick, J., 2012. Value capture in the global electronics industry: Empirical evidence for the “smiling curve” concept. *Industry and Innovation* 19, 89–107. doi:10.1080/13662716.2012.650883.
- Smith, B.L., Woodhouse, M., Horowitz, K.A.W., Silverman, T.J., Zuboy, J., Margolis, R.M., 2021. Photovoltaic (PV) Module Technologies: 2020 Benchmark Costs and Technology Evolution Framework Results. Technical Report NREL/TP-7A40-78173. National Renewable Energy Laboratory. Golden, CO. URL: <https://www.osti.gov/biblio/1829459>, doi:10.2172/1829459.
- Snyder, L.V., Daskin, M.S., 2005. Reliability models for facility location: The expected failure cost case. *Transportation Science* 39, 400–416. URL: <https://www.proquest.com/scholarly-journals/reliability-models-facility-location-expected/docview/203515217/se-2>, doi:10.1287/trsc.1040.0107.
- Snyder, L.V., Daskin, M.S., 2006. Stochastic p-robust location problems. *IIE Transactions* 38, 971–985. doi:10.1080/07408170500469113.
- Tomlin, B., 2006. On the value of mitigation and contingency strategies for managing supply chain disruption risks. *Management Science* 52, 639–657. URL: <https://www.jstor.org/stable/20110543>, doi:10.1287/mnsc.1060.0515.
- Trina Solar, 2016. Trina Solar Launches Operations at Thailand Manufacturing Facility and Signs a US\$143 million Syndicated Financing Facilities Agreement. URL: <https://www.trinasolar.com/en-glb/20160327-1627/>. press release. Accessed 3 October 2026.
- Trina Solar, 2023. Trinasolar begins producing 210mm monocrystalline wafers in Vietnam. URL: <https://www.trinasolar.com/us/resources/newsroom/Trina-Solar-210mm-monocrystalline-wafers-Vietnam>. press release. Accessed 3 October 2026.
- Trina Solar, 2024a. Power Beyond Solar. Corporate brochure. Trina Solar Co., Ltd. URL: https://static.trinasolar.com/sites/default/files/Trinasolar_Corporate_Brochure_EN_20240514.pdf. accessed: 2026-09-10.
- Trina Solar, 2024b. Trinasolar brings advanced solar innovations to Vietnam, demonstrates commitment to empowering renewable sector. URL: <https://www.trinasolar.com/en-apac/resources/newsroom/aptrinasolar-brings-advanced-solar-innovations-vietnam-demonstrates-commitment/>. Accessed: 2026-09-10.
- U.S. Department of Commerce, 2025. Final affirmative determinations in the AD and CVD investigations of crystalline silicon photovoltaic cells from Cambodia, Malaysia, Thailand, and Vietnam. URL: <https://www.trade.gov/final-affirmative-determinations-antidumping-and-countervailing-duty-investigations-crystalline>. international Trade Administration, April 2025.
- U.S. Department of Commerce, International Trade Administration, 2023. Antidumping and countervailing duty orders on crystalline silicon photovoltaic cells, whether or not assembled into modules, from the People’s Republic of China: Final scope determination and final affirmative determinations of circumvention with respect to Cambodia, Malaysia, Thailand, and Vietnam. URL: <https://www.govinfo.gov/content/pkg/FR-2023-08-23/pdf/2023-18161.pdf>. 88 Fed. Reg. 57419, 23 August 2023, FR Doc. 2023-18161, cases A-570-979 and C-570-980.

- Verter, V., Dincer, M.C., 1995. Facility location and capacity acquisition: An integrated approach. *Naval Research Logistics* 42, 1141–1160. doi:10.1002/1520-6750(199512)42:8<1141::AID-NAV3220420803>3.0.CO;2-B.
- World Bank, 2024. World Integrated Trade Solution (WITS). URL: <https://wits.worldbank.org>. accessed: 2025.
- World Bank, n.d. Technical Note on Bilateral Applied Tariff. World Integrated Trade Solution (WITS). URL: <https://wits.worldbank.org/Bilateral-Tariff-Technical-Note.html>. accessed 5 October 2026.
- World Customs Organization, n.d. What is the Harmonized System (HS)? URL: <https://www.wcoomd.org/en/topics/nomenclature/overview/what-is-the-harmonized-system.aspx>. accessed 5 October 2026.
- World Trade Organization, 1994a. Agreement on Implementation of Article VII of the General Agreement on Tariffs and Trade 1994. Annex 1A to the Marrakesh Agreement Establishing the World Trade Organization. URL: https://www.wto.org/english/docs_e/legal_e/20-val.pdf. Customs Valuation Agreement. Accessed 1 August 2026.
- World Trade Organization, 1994b. Marrakesh Agreement Establishing the World Trade Organization. URL: https://www.wto.org/english/docs_e/legal_e/marag_e.htm. done at Marrakesh on 15 April 1994, in force from 1 January 1995.
- World Trade Organization, 1996. Ministerial declaration on trade in information technology products. WT/MIN(96)/16. Adopted in Singapore, 13 December 1996.
- World Trade Organization, 2025. Anti-dumping, countervailing and safeguard measures: Initiations and measures by reporting member. WTO Data Portal. URL: https://data.wto.org/dataset/wto_trdb_adp. countervailing: https://data.wto.org/dataset/wto_trdb_cvl, safeguards: <https://data.wto.org/dataset/safeguards>, accessed 30/09/2026.
- World Trade Organization, n.d.a. Understanding the WTO: Anti-Dumping, Subsidies, Safeguards: Contingencies, etc. URL: https://www.wto.org/english/thewto_e/whatis_e/tif_e/agrm8_e.htm. accessed 22 September 2026.
- World Trade Organization, n.d.b. Understanding the WTO: Principles of the Trading System. URL: https://www.wto.org/english/thewto_e/whatis_e/tif_e/fact2_e.htm. accessed 22 September 2026.
- Yi, K.M., 2003. Can vertical specialization explain the growth of world trade? *Journal of Political Economy* 111, 52–102. doi:10.1086/344805.
- Zeng, B., Zhao, L., 2013. Solving two-stage robust optimization problems using a column-and-constraint generation method. *Operations Research Letters* 41, 457–461. doi:10.1016/j.orl.2013.05.003.
- Zhou, Y., Gong, D.C., Huang, B., Peters, B.A., 2017. The impacts of carbon tariff on green supply chain design. *IEEE Transactions on Automation Science and Engineering* 14, 1542–1555. doi:10.1109/TASE.2015.2445316.
- Zhu, L., 2024. Reconfiguring globalisation: A review of tariffs, industrial policies, and the global solar PV supply chain. OIES Paper CE15. Oxford Institute for Energy Studies. Oxford. URL: <https://hdl.handle.net/10419/309954>.



Computational detail

A.1. Instance calibration

Value division

The raw-material share of 10 in Table 7.4 has no direct counterpart. Firm-level purchased inputs are far larger, because they include intermediates that the constructed chains make in-house: among producers of intermediate inputs, median value added is 37 percent of output (Rungi and Del Prete, 2017). Commitment is a per-period cost of holding capacity qualified, so capital intensity gives only a rough comparison. At 15, a quarter of the cost of goods, it lies above the 8 to 10 percent of production cost that capital expenditure represents in the solar chain (Basore and Feldman, 2022). The margin of 40 lies above the gross margins of 31 to 32 percent that Shin et al. (2012) report for electronics lead firms and component suppliers, and close to the all-market average of 37.8 percent (Damodaran, 2026).

Bloc multipliers

Two measured quantities bound the multiplier. Chinese manufacturing unit labour costs, net of productivity, stood at 68.3 percent of United States levels in 2008 on a census-based measure and at 32.9 percent in 2009 on a national-accounts measure, so United States labour cost per unit of output is 1.46 to 3.04 times China's (Ceglowski and Golub, 2012). Labour is about 16 percent of United States manufacturing cost and up to 30 percent in electronics and apparel, on Morgan Stanley figures that KPMG LLP and The Manufacturing Institute (2020) reports. With non-labour inputs common to both, total cost is $1 + 0.16(1.46 - 1) = 1.07$ to $1 + 0.16(3.04 - 1) = 1.33$ times China's at the average labour share and 1.14 to 1.61 times at a share of 0.30, so the multiplier of 1.50 for C^H sits in the upper part of the combined band of 1.07 to 1.61. Producing a car in a low-cost location rather than domestically cuts its landed cost in a high-cost market by 15 to 20 percent, freight included, a ratio of 1.18 to 1.25 (Küpper and Nordemann, 2025).

Shipping charges

Shipping costs 3 percent of embodied cost across a bloc border, 1 percent within a bloc and nothing between consecutive steps at the same location. Embodied cost is the supplied inputs plus the production cost before the bloc multiplier, up to the commodity shipped, so the charge rises along the chain. The cross-border rate is calibrated so that a Serial chain that changes bloc at every step pays $0.03(17 + 24 + 31 + 38 + 45) = 4.65$ over its five crossings, close to the median freight, insurance and other charges on one shipment into the United States, 4.7 percent of its free-on-board value across 439 manufacturing industries (Barrot et al., 2019). Each crossing therefore costs less than observed freight: the final good pays 1.35 percent of its customs value. The within-bloc rate and the zero rate at a single location are assumptions.

The replacement cost in Section 7.3.1 follows from these rules. Remaking one unit of the Serial final good at the C^H facility costs $10 + 4 \times 7 + 1.5 \times 7 = 48.5$ in materials and production. Its net shipping of 0.34 adds the raw material within C^L (0.01×10), the intermediate from C^L to C^H (0.03×38) and delivery

within C^H (0.01×45), and subtracts the import from C^T it replaces (0.03×45). The intermediate shipped into C^T is paid either way. The total of 48.84, against a customs value of 100, gives the threshold rate of 0.488.

Scenario counts

Table A.1 gives the number of scenarios by architecture and demand orientation. Counts include the baseline scenario. The three dual-market orientations serve the same two markets and so share one count.

Table A.1: Scenario counts by architecture and demand orientation.

Architecture	Concentrated C^H or C^M	Concentrated C^L	Dual	Dispersed
Serial	18	21	28	41
Modular	26	31	42	63
Complex	24	28	37	54

A.2. Full design frontiers

Tables A.2 and A.3 give the frontiers of the Modular and Complex C^H instances in the form of Table 7.11. PF is the final step, PM the Complex merge, $siPj$ step j of stream i counted from the raw material, and u a step between the merge and final assembly. Each move is the change against the rung above, and volumes are units out of a demand of 100. The last interval starts at the floor, so its endpoints can coincide after rounding.

Table A.2: Frontier of the Modular C^H instance.

Rung	PF	Move	Baseline cost	π	Binding	Optimal for $\theta \in$
1	C^L	<i>deterministic network</i>	6,158.3	0.000	bilateral escalation	[0.734, 1.00]
2	C^L 95 C^T 5	duplication PF into C^T	6,281.4	0.020	composite	[0.575, 0.60]
3	C^L 95 C^H 5	relocation PF C^T to C^H , s3P4 to C^T	6,383.6	0.037	bilateral escalation	[0.401, 0.45]
4	C^L 90 C^H 5, C^T 5	duplication PF into C^T , relocation s3P4 back to C^L	6,405.4	0.040	friendshoring	[0.273, 0.30]
5	C^L 90 C^H 5, C^M 5	relocation PF C^T to C^M	6,408.1	0.041	market protection	[0.083, 0.25]
6	C^H 95 C^M 5	consolidation PF out of C^L	6,447.1	0.047	export control	[0.078, 0.08]

Table A.3: Frontier of the Complex C^H instance.

Rung	PF	PM	Move	Baseline cost	π	Binding	Optimal for $\theta \in$
1	C^L	C^L	<i>deterministic net-work</i>	6,203.3	0.000	bilateral escalation	[0.665, 1.00]
2	C^H	C^L	relocation PF to C^H	6,333.8	0.021	bilateral escalation	[0.583, 0.60]
3	C^H	C^L	relocation s3P4 to C^T	6,378.3	0.028	bilateral escalation	[0.380, 0.45]
4	C^H	C^M	relocation PM and uP4 to C^M , s3P4 back to C^L	6,454.7	0.041	bilateral escalation	[0.184, 0.30]
5	C^H	C^M	relocation s2P1 and s2P2 to C^T	6,484.4	0.045	bilateral escalation	[0.146, 0.15]
6	C^H 95 C^M 5	C^T 95 C^M 5	duplication PF into C^M and PM into C^T , relocation uP4 C^M to C^T , s2P1 and s2P2 back to C^L	6,692.3	0.079	bilateral escalation	[0.113, 0.12]
7	C^H	C^L 95 C^M 5	consolidation PF out of C^M , relocation PM C^T to C^L , uP4 C^T to C^H , s3P1 to C^T , duplication s3P4 into C^M	6,733.0	0.085	bilateral escalation	[0.097, 0.10]
8	C^H	C^H 95 C^M 5	relocation PM to C^H , s3P1 back to C^L	6,737.8	0.086	export control	[0.096, 0.096]

The Complex frontier under bilateral escalation

For each rung of the frontier of the Complex C^H instance, Table A.4 gives the pair that the worst bilateral-escalation scenario names, the cap on that pair, the customs value it taxes and the resulting duty, and whether the regime binds.

Table A.4: Frontier of the Complex C^H instance under bilateral escalation, in units of final-good value. Figure A.1 rings the movements taxed.

Rung	Named pair	Cap	Value taxed	Duty	Binds
1	C^L and C^H	0.54	10,000	5,400	yes
2	C^L and C^H	0.54	8,700	4,698	yes
3	C^L and C^H	0.54	6,000	3,240	yes
4	C^L and C^M	0.54	3,333	1,800	yes
5	C^M and C^H	0.25	6,000	1,500	yes
6	C^T and C^M	0.54	1,973	1,065	yes
7	C^L and C^M	0.54	1,688	912	yes
8	C^M and C^H	0.25	1,808	452	no





Figure A.1: Frontier of the Complex C^H instance under its worst bilateral-escalation scenarios, one panel per rung. Rings mark the taxed movements.

Figure A.2 draws the midpoint and floor designs of the three C^H instances with their baseline flows.**Figure A.2:** Committed networks and baseline flows, one row per architecture: the midpoint design of each frontier and the design at the robustness floor.

A.3. Regime flows by architecture

Figures A.3 to A.5 draw the floor design of each C^H instance under baseline tariffs and under every regime that taxes it. Arc width ranks volume among crossings, within-bloc movements appear light, and rings mark the taxed movements, at the cap shown.

Serial finishes its bulk in allied C^M with a small facility in the market. It moves finishing into C^H under bilateral escalation and market protection, which then tax different movements, and holds its flows under friendshoring, which spares the allied origin.



Figure A.3: The Serial floor network under baseline and under the three regimes that tax it.

Modular moves production only under friendshoring. Under bilateral escalation, market protection and export control it keeps its baseline flows and pays on a facility held at minimum throughput or on a raw material from its designated origin.




Figure A.4: The Modular floor network under baseline and under the four regimes that tax it.

Complex keeps its closing step in the market. Under friendshoring it moves its merge and the third stream's last step, the two steps it holds in a second location, into C^M . Under bilateral escalation it moves nothing and pays its largest duty, because the named relation between C^M and C^H carries the strategic component and two facilities at minimum throughput. Under export control it pays on a raw material with no alternative consuming location.




Figure A.5: The Complex floor network under baseline and under the three regimes that tax it.

A.4. The robustness floors in full

Table A.5 gives the floors and premiums of the 21 instances at $q = 4$ to four decimals, against three in Chapter 7, with the raw-material origins fixed, as in the main grid, and flexible, as in Section 7.4.1. Instances are classified on unrounded values. The floor $\theta_{\min}$ is the regret the committed design attains. Bold floors are exposure-bound, the others premium-bound with $\theta_{\min} = \pi$, and tied regimes, coded by their bundles as in Table 6.1, are listed together. Flexible origins raise the floor in two Complex instances because the reference network gains more than the design. Table A.6 gives the floors by location rights. At $q = 1$ the dispersed Serial and Modular floors of 0.0048 and 0.0324 are low only because their reference networks, confined to one location per step as well, avoid little more exposure than the designs, so they do not measure robustness.

Table A.5: Floors, premiums at the floor and binding regimes for the 21 instances at $q = 4$, with raw-material origins fixed and flexible.

Orientation	Fixed origins			Flexible origins		
	$\theta_{\min}$	π	Binds	$\theta_{\min}$	π	Binds
<i>Serial</i>						
Conc. C^H	0.1020	0.1020	b_5, b_6	0.1020	0.1020	b_5, b_6
Conc. C^M	0.0888	0.0888	b_5, b_6	0.0888	0.0888	b_5, b_6
Conc. C^L	0.0000	0.0000	all seven	0.0000	0.0000	all seven
Dual balanced	0.0946	0.0946	b_5, b_6	0.0946	0.0946	b_5, b_6
Dual asym. C^H	0.0990	0.0990	b_5, b_6	0.0990	0.0990	b_5, b_6
Dual asym. C^M	0.0902	0.0902	b_5, b_6	0.0902	0.0902	b_5, b_6
Dispersed	0.1030	0.1030	b_5, b_6	0.1030	0.1030	b_5, b_6
<i>Modular</i>						
Conc. C^H	0.0782	0.0469	b_6	0.0470	0.0470	b_5, b_6
Conc. C^M	0.0687	0.0647	b_3	0.0435	0.0391	b_3
Conc. C^L	0.0327	0.0000	b_6	0.0000	0.0000	all seven
Dual balanced	0.0712	0.0398	b_6	0.0399	0.0399	b_5, b_6
Dual asym. C^H	0.0712	0.0398	b_6	0.0420	0.0399	b_1
Dual asym. C^M	0.0710	0.0396	b_6	0.0397	0.0397	b_5, b_6
Dispersed	0.0716	0.0400	b_6	0.0401	0.0401	b_5, b_6
<i>Complex</i>						
Conc. C^H	0.0959	0.0862	b_6	0.0863	0.0863	b_5
Conc. C^M	0.0748	0.0748	b_5	0.0704	0.0704	b_5
Conc. C^L	0.0586	0.0472	b_6	0.0662	0.0473	b_6
Dual balanced	0.0821	0.0718	b_6	0.0868	0.0868	b_5
Dual asym. C^H	0.0972	0.0875	b_6	0.0886	0.0886	b_5
Dual asym. C^M	0.0805	0.0661	b_6	0.0675	0.0675	b_5
Dispersed	0.1022	0.0926	b_6	0.0978	0.0978	b_5

Table A.6: Attained floors by location rights q .

Orientation	Serial				Modular				Complex			
	$q=1$	$q=2$	$q=3$	$q=4$	$q=1$	$q=2$	$q=3$	$q=4$	$q=1$	$q=2$	$q=3$	$q=4$
Conc. C^H	0.1467	0.1020	0.1020	0.1020	0.1173	0.0782	0.0782	0.0782	0.0951	0.0959	0.0959	0.0959
Conc. C^M	0.1466	0.0888	0.0888	0.0888	0.1207	0.0687	0.0687	0.0687	0.0748	0.0748	0.0748	0.0748
Conc. C^L	0.0000	0.0000	0.0000	0.0000	0.0327	0.0327	0.0327	0.0327	0.1966	0.0937	0.0586	0.0586
Dual balanced	0.1144	0.0946	0.0946	0.0946	0.0926	0.0736	0.0712	0.0712	0.0595	0.0821	0.0821	0.0821
Dual asym. C^H	0.1464	0.0990	0.0990	0.0990	0.0951	0.0762	0.0712	0.0712	0.0642	0.0972	0.0972	0.0972
Dual asym. C^M	0.1333	0.0902	0.0902	0.0902	0.0904	0.0710	0.0710	0.0710	0.0556	0.0805	0.0805	0.0805
Dispersed	0.0048	0.1022	0.1030	0.1030	0.0324	0.0817	0.0716	0.0716	0.0735	0.0863	0.1022	0.1022

A.5. Supplier dependence in full


Table A.7 gives every point behind Figure 7.4. Bold floors are exposure-bound. Share and Production are the component's value and the in-house production budget, in percent of final-good value, and Merge at lists the locations where the merge is qualified, not only where its bulk runs.


Table A.7: Supplier dependence in the Complex C^H instance.

Mode	Share	Production	C_0^*	$\theta_{\min}$	π	Binds	Merge at
Production	5	30	6,173	0.0736	0.0504	b_6	C^L
Production	7	28	6,179	0.0902	0.0796	b_1	C^L+C^M
Production	9	26	6,185	0.0889	0.0777	b_1	C^L+C^M
Production	11	24	6,191	0.0924	0.0859	b_1	C^L+C^M
Production	13	22	6,197	0.0956	0.0888	b_6	C^L+C^M
Production	15	20	6,203	0.0959	0.0862	b_6	C^H+C^M
Production	18	17	6,212	0.0886	0.0886	b_5	C^L+C^M
Production	20	15	6,218	0.0814	0.0814	b_5	C^L+C^M
Production	22	13	6,224	0.0774	0.0742	b_1	C^L+C^M
Production	24	11	6,230	0.0695	0.0695	b_5	C^L+C^M
Production	26	9	6,236	0.0622	0.0604	b_6	C^L+C^M
Production	28	7	6,242	0.0580	0.0563	b_4	C^M
Production	30	5	6,248	0.0518	0.0458	b_4	C^M
Margin	5	35	6,688	0.0648	0.0505	b_6	C^L
Margin	7	35	6,900	0.0843	0.0427	b_6	C^L
Margin	9	35	7,112	0.0830	0.0761	b_1	C^L+C^M
Margin	11	35	7,324	0.0850	0.0850	b_5	C^L+C^M
Margin	13	35	7,536	0.0850	0.0850	b_5	$C^L+C^M+C^T$
Margin	15	35	7,748	0.0830	0.0830	b_5	$C^L+C^M+C^T$
Margin	18	35	8,066	0.0802	0.0802	b_5	$C^L+C^M+C^T$
Margin	20	35	8,278	0.0785	0.0785	b_5	$C^L+C^M+C^T$
Margin	22	35	8,490	0.0769	0.0769	b_5, b_6	$C^L+C^M+C^T$
Margin	24	35	8,702	0.0775	0.0753	b_6	$C^L+C^M+C^T$
Margin	26	35	8,914	0.0763	0.0755	b_5	C^M
Margin	28	35	9,126	0.0731	0.0731	b_5	C^M
Margin	30	35	9,338	0.0711	0.0711	b_5	C^M
Half	5	32.5	6,431	0.0690	0.0505	b_6	C^L
Half	7	31.5	6,540	0.0888	0.0788	b_1	C^L+C^M
Half	9	30.5	6,649	0.0842	0.0769	b_1	C^L+C^M
Half	11	29.5	6,758	0.0878	0.0877	b_1	C^L+C^M
Half	13	28.5	6,867	0.0879	0.0879	b_5	$C^L+C^M+C^T$
Half	15	27.5	6,976	0.0861	0.0861	b_5	$C^L+C^M+C^T$
Half	18	26	7,139	0.0848	0.0835	b_6	$C^L+C^M+C^T$
Half	20	25	7,248	0.0831	0.0820	b_6	C^H+C^M
Half	22	24	7,357	0.0804	0.0749	b_1	C^M
Half	24	23	7,466	0.0763	0.0717	b_3	C^M
Half	26	22	7,575	0.0737	0.0686	b_3	C^M
Half	28	21	7,684	0.0712	0.0655	b_3	C^M
Half	30	20	7,793	0.0690	0.0625	b_6	C^M

A.6. Cost under shock along the frontier

Figures A.6 and A.7 repeat Figure 7.2 for the two orientations that bracket the C^H instance, demand concentrated in C^M , which gives the longest frontiers, and dispersed demand, which gives the shortest. Each bar is one rung under its binding regime, from the loosest tolerance on the left to the floor on the right, labelled with the bundle, its worst scenario and its cap. The dashed line is the design's baseline cost and π its premium.

Figure A.6: Cost of each rung under its binding regime, demand concentrated in C^M , $q = 4$.

Figure A.7: Cost of each rung under its binding regime, dispersed demand, $q = 4$.

B

Solar PV instance calibration

B.1. Network

The instance holds three supplier nodes, $V^S = \{S_{CN}, S_{EU}, S_{US}\}$, which supply polysilicon only, six production nodes, $V^P = \{CN_P, VN_P, MY_P, TH_P, KR_P, US_P\}$, and one customer node, $V^C = \{US_C\}$, in seven jurisdictions. The feasible movement set X holds 69 movements, of which 9 are procurement, 18 wafer to cell, 36 cell to module and 6 module deliveries. Co-located steps use arcs whose origin equals their destination and pay no shipping, and three arcs between distinct nodes in the same country pay the domestic rate.

B.2. Gross values

No source observes transaction values for the firm's wafers, cells and modules, and the IEA's aggregate module selling price and regional manufacturing costs do not establish them (International Energy Agency, 2022). Each produced commodity therefore takes the computed value of Section 8.1.3, its valuation base in Table B.3 divided by 0.85, so that the margin is 15 percent of the resulting value, as the benchmark defines it. For the module, $298/0.85 = 350.59$. Polysilicon keeps its purchase price of 98, since a constructed margin would add the supplier's margin a second time. Realised industry gross margins have ranged from 5 to 25 percent since 2010 (Smith et al., 2021).

B.3. Cost calibration

Sources

The IEA supplies regional stage costs, which are incremental, and regional depreciation, all excluding subsidies, transport, company profits, taxes and tariffs (International Energy Agency, 2022). The DOE supplies investment requirements and silicon intensity for a one-gigawatt facility (Basore and Feldman, 2022). Neither reports the split into commitment and production cost, so depreciation, the observable cost of installed capability, stands in for commitment. The proxy leaves maintenance and overhead that do not vary with throughput in production cost, ignores differences in asset lives across the chain, and treats the IEA's overhead as plant-level, a boundary the report does not define.

Depreciation and its allocation

The IEA puts depreciation at 20 MUSD per GW in China, 25 in Southeast Asia, 49 in the United States and 59 in Korea, digitised from Figure B.1 to about ± 2 MUSD per GW. The DOE puts the investment in a one-gigawatt facility at 110 to 140 MUSD for polysilicon, 80 to 100 for ingots and wafers, 50 to 130 for cells and 50 to 80 for modules (Basore and Feldman, 2022). The midpoints, 125, 90, 90 and 65, sum to 370, and each region's depreciation is split across the steps in these proportions, so the wafer and cell steps each carry $90/370$ of it as commitment cost and the module step $65/370$. The investments themselves do not enter the objective. The polysilicon share of $125/370$ drops out because the firm buys polysilicon, so $245/370$, about two thirds, of each region's depreciation enters the instance.

Cost coefficients

Table B.1 subtracts the allocated commitment cost from each IEA stage cost, digitised from Figure B.2, to give the production cost, so the two reproduce the IEA figure at the stored precision. Stage and production costs are in MUSD per GW of output, commitment cost in MUSD per period for a qualified one-gigawatt facility, and Vietnam, Malaysia and Thailand share the Southeast Asian calibration.

Table B.1: Derivation of the cost coefficients.

Region	Step	IEA stage cost	Commitment f_{ip}^P	Production c_{ip}^{prod}
China	wafer	33.0	4.865	28.135
	cell	39.0	4.865	34.135
	module	128.0	3.514	124.486
Southeast Asia	wafer	35.0	6.081	28.919
	cell	43.0	6.081	36.919
	module	132.0	4.392	127.608
Korea	cell	61.0	14.351	46.649
	module	135.0	10.365	124.635
United States	wafer	47.0	11.919	35.081
	cell	61.0	11.919	49.081
	module	139.0	8.608	130.392

Supplied polysilicon

The firm buys polysilicon at the 2022 market price of 35 USD per kilogram (International Energy Agency, 2022) and the DOE's silicon intensity of 2.8 grams per watt (Basore and Feldman, 2022), so it pays $35 \times 0.0028 = 0.098$ USD per watt, or 98 MUSD per GW, at each of the three supplier origins.

Reconciliation

Table B.2 sets the sum of the four IEA stage costs against each reported regional total, both digitised to about ± 2 MUSD per GW from Figures B.2 and B.1. The largest gap, 2.1 percent, lies within the reading precision, so the check tests the digitisation. The Korean and European rows deviate most, mainly in polysilicon production costs that do not enter the instance. The regional totals also match the IEA's text, from 0.24 USD per watt in China to 0.33 in Europe (International Energy Agency, 2022). A second check replaces the IEA's polysilicon manufacturing cost of 38 with the firm's purchase price of 98, so the Chinese chain costs $98 + 33 + 39 + 128 = 298$ MUSD per GW, against the IEA's statement that Chinese module manufacturing cost approaches 0.30 USD per watt at a polysilicon price of 35 USD per kilogram (International Energy Agency, 2022).

Table B.2: Sum of the four stage costs against the reported regional total, in MUSD per GW.

Region	Poly	Wafer	Cell	Module	Sum	Reported	Deviation
China	38	33	39	128	238	240	−0.8 percent
Southeast Asia	39	35	43	132	249	252	−1.2 percent
United States	45	47	61	139	292	295	−1.0 percent
Korea	62	46	61	135	304	310	−1.9 percent
European Union	73	50	61	139	323	330	−2.1 percent

Two base vectors

Table B.3 gives the two costs accumulated along the Chinese reference chain. The objective charges shipping on the shipping base. The valuation base adds cumulative commitment cost, the difference column, and divided by 0.85 gives the gross values of the produced commodities. Both bases are fixed by commodity, so a Korean cell carries the same bases as a Chinese one, which keeps the shipping coefficient c_{ijg} exogenous.

Table B.3: The shipping base and the valuation base, in MUSD per GW.

Commodity	Shipping base	Valuation base	Difference
Polysilicon	98.000	98	0
Wafer	126.135	131	4.865
Cell	160.270	170	9.730
Module	284.756	298	13.244

B.4. Baseline tariff environment

The baseline tariff of the case is the bilateral effectively applied rate of Section 5.2.2, the lower of the importer's MFN rate and any preferential rate it grants the exporting country (World Bank, n.d.). Table B.4 gives this rate from WITS/World Bank TRAINS (World Bank, 2024) for every product the modelled chain moves across a border, as the simple average across national tariff lines for 2020 to 2023, the latest years reported. Products carry their HS codes (World Customs Organization, n.d.): high-purity silicon 280461, wafers 381800, and cells and modules 854140, which the 2022 revision of the HS divided, giving cells 854142 and modules 854143. The case assumes that modules enter under 854143, the code under which Commerce measured module imports in its 2025 investigations (U.S. Department of Commerce, 2025), and leaves out modules declared under heading 8501, which the Information Technology Agreement does not cover.

Wafers, cells and modules enter at zero for every importer in the query, from every partner, in every year reported. All three fall under the Information Technology Agreement, whose participants bind and eliminate their duties on them (World Trade Organization, 1996), so the MFN rate is zero and no preference can go lower. The bilateral query does not cover the United States as an importer. A query against all partners together gives zero on every tariff line for United States imports of all four products, and for Vietnamese imports of high-purity silicon, in every year from 2015 to 2023.

High-purity silicon falls outside the agreement. China charges 4 percent on it from the United States, which holds no preference, between 1.67 and 2.11 percent from Korea and nothing from Malaysia and Vietnam. Of these lanes the modelled chain admits only the one from the United States, and its baseline takes the observed 4 percent. The query does not cover silicon from the European Union into China. The two have no preferential agreement, so that lane takes China's MFN rate, the same 4 percent. Neither lane carries flow in any solution, since Chinese wafer production buys Chinese polysilicon at the same price without crossing a border.

The effectively applied rate leaves out the measures the regimes represent. The United States rate on cells and modules stays at zero through the 2018 safeguard and the Section 301 duties of the same year (Bown, 2021), and China's 4 percent on United States silicon leaves out its anti-dumping and countervailing duties on that product (Basore and Feldman, 2022).

Table B.4: Bilateral effectively applied rates on the products the modelled chain moves across borders, simple averages across national tariff lines in percent, for 2020 to 2023 where WITS reports a rate. Each zero holds on every line, and partners exclude the importer itself.

Product	Importer	Partner	Rate
Wafers, HS 381800	CN, KR, MY, TH, VN	CN, KR, MY, TH, US, VN	0
Cells and modules, HS 854140 (2020 to 2021), HS 854142 and 854143 (2022 to 2023)	CN, KR, MY, TH, VN	CN, KR, MY, TH, US, VN	0
High-purity silicon, HS 280461	KR, MY, TH, VN	CN, KR, MY, TH, US, VN	0
	CN	US	4.00
	CN	KR	1.67 to 2.11
	CN	MY, VN	0
All four, 2015 to 2023	US	all partners together	0

B.5. Source exhibits

Figures B.1 and B.2 are the IEA exhibits from which Section B.3 reads regional depreciation and stage costs. Both figures include India, which no coefficient of the instance uses.




Figure B.1: Production cost of mono-PERC crystalline-silicon components by input and region, 2022, excluding subsidies, transport, company profits, taxes and tariffs. Source: International Energy Agency (2022), *Special Report on Solar PV Global Supply Chains*. All rights reserved.





Figure B.2: Production cost of mono-PERC crystalline-silicon components by supply-chain segment, 2022, excluding subsidies, transport, company profits, taxes and tariffs. Source: International Energy Agency (2022), *Special Report on Solar PV Global Supply Chains*. All rights reserved.



Implementation and reproducibility

The results of Chapters 7 and 8 come from one Python implementation of the model of Chapter 6, solved with Gurobi under the settings of Section 7.1.7.

C.1. Solution procedure

Algorithm 1 implements the solution procedure of Section 6.4 for one instance. The reference costs come first because every regret divides by one of them.

C.2. Solver record

The 84 instances of the main grid required 2,016 mixed-integer solves: 672 reference costs, 1,260 sweep solves and 84 floor designs. The code accepts only an optimal or infeasible status and stops on any other, so all ended optimal (1,700) or infeasible (316). None reached the one-hour limit, and the longest took 350 seconds. The 513 floor-bisection probes are not among them. The targeted experiments ran their reference solves at Gurobi's default settings, without a time limit, and no instance took longer than 599 seconds.

The case study's three reported instances, the base instance and the two controls of Chapter 8, required 252 mixed-integer solves: three deterministic solves, 24 reference costs, 83 tolerance masters and 142 solves of the bisection that locates each step in the premium. All ended optimal, and none took longer than 1.4 seconds. Their floor-bisection probes are not among them. A probe stops at the first design that meets a trial tolerance, may legitimately prove infeasibility, and leaves no log entry. The run records hold no time or status for the recourse problems behind the case study's fixed-design evaluations, the two footprints among them. Every reported value is finite, which rules out an infeasible evaluation but is not a record of solver status.

Algorithm 1 Solution procedure for one instance**Input:** the tolerance grid of Section 7.1.6**Output:** reference costs C_b^* , designs a with their regrets $R^b(a)$, the floor

```

1: for  $b \in \{b_0, \dots, b_7\}$  do                                     ▷ reference costs
2:    $C_b^* \leftarrow \min_a [C^{\text{fixed}}(a) + Q^b(a)]$ 
3: end for
4: for each  $\theta$  on the grid, from tightest to loosest do           ▷ frontier
5:   solve the robust problem at  $\theta$  to optimality to get  $a$ , or prove it infeasible
6: end for
7: upper end  $\leftarrow$  tightest feasible  $\theta$  on the grid             ▷ bracket
8: lower end  $\leftarrow$  grid value below it, or 0
9: search the robust problem at the lower end for a first admissible design
10: if an admissible design is found then                           ▷ possible only at 0
11:   upper end  $\leftarrow$  lower end
12: end if
13: while upper end – lower end  $> 10^{-3}$  do                       ▷ floor bisection
14:    $\theta \leftarrow (\text{lower end} + \text{upper end})/2$ 
15:   search the robust problem at  $\theta$  for a first admissible design
16:   if an admissible design is found then
17:     upper end  $\leftarrow \theta$ 
18:   else if infeasibility is proved then
19:     lower end  $\leftarrow \theta$ 
20:   else
21:     stop the instance                                           ▷ unresolved probe
22:   end if
23: end while
24: floor design  $a \leftarrow$  robust problem at the upper end, solved to optimality
25: for each frontier design and the floor design  $a$  do             ▷ evaluation with  $a$  fixed
26:   for  $b \in \{b_0, \dots, b_7\}$  do
27:      $Q^b(a) \leftarrow \max_{\omega \in \mathcal{U}(b)} Q^\omega(a)$                  ▷ one operating problem per  $\omega$ 
28:      $R^b(a) \leftarrow [C^{\text{fixed}}(a) + Q^b(a)]/C_b^* - 1$          ▷  $R^{b_0}(a) = \pi(a)$ 
29:   end for
30:   binding regimes of  $a \leftarrow \arg \max_{b \in \{b_1, \dots, b_7\}} R^b(a)$ , all of them where regimes tie
31: end for
32: return the floor  $\max_{b \in \{b_1, \dots, b_7\}} R^b(a)$  of the floor design

```
