The Wrong Tradecraft

Coups, Movements, and the Democratic Failure to Understand Legitimacy

COMPLETE MANUSCRIPT · FIRST FULL DRAFT

CONTENTS

The Wrong Tradecraft

Coups, Movements, and the Democratic Failure to Understand Legitimacy

CHAPTER ONE

The Coup as a Product Category

In March 1954, seven months after the fact, a man sat down to write up what had happened in Tehran. He was a historian by training, working under contract, and the document he produced was not a memoir or an apologia. It was a manual. It had numbered sections. It described the planning cycle, the recruitment of local assets, the management of the press, the use of paid crowds, the coordination with the palace, and the improvisations forced by the near-total failure of the first attempt. It assessed what had worked. It identified what should be done differently. It was classified, it was circulated internally, and it was written for the explicit benefit of officers who would be asked to do something similar somewhere else.

That document is the actual subject of this chapter. Not the operation it describes — the operation has been told many times and told well — but the fact that within seven months of the event, the event had been converted into a method. A thing that happened once, contingently, in a specific city, under conditions that very nearly produced the opposite result, had been abstracted into a set of transferable procedures. Somebody had decided it was repeatable. That decision, more than anything that occurred in Tehran in August 1953, is what this book is about.

Consider what was purchased, and by whom, and for how much.

• • •

The commercial position at issue was old and, by the standards of the postwar world, indefensible. A British company, majority-owned by the British government since before the First World War, held a concession over Iranian oil on terms that returned to Iran a minority share of a profit figure that Iran was not permitted to audit. The refinery at Abadan was among the largest in the world. The company's tax payments to the British Treasury in some years exceeded its royalty payments to the Iranian state. When the Saudi government concluded an arrangement with the American consortium there on a fifty-fifty basis in 1950, the terms in Iran became not merely unpopular but conspicuously anomalous, and the company declined to match them.

It is worth stopping at Abadan before going further, because the argument of this book will spend a great deal of time on abstractions and the abstractions were imposed on a place. Abadan was a company town of the classic kind, built around the refinery, housing tens of thousands of Iranian workers in accommodation that was separated by grade from the accommodation provided to British staff, with the amenities distributed accordingly. The oil workers had struck before, in numbers, over pay and conditions and the visible arithmetic of who lived where. When the Majlis voted to nationalize, it was not an abstraction to those men; it was the proposition that the thing they worked inside might belong to the country they lived in. And when the embargo closed the refinery and the British technical staff went home, it was those same men who lost the work. The nationalization was popular in Iran for reasons that had very little to do with the Cold War and everything to do with a set of housing blocks in Khuzestan, and no assessment written in Washington or London in 1953 records this as a fact bearing on the durability of anything.

The Iranian parliament voted to nationalize the industry in the spring of 1951. Mohammad Mossadegh, who had led the parliamentary campaign for nationalization, became prime minister within weeks. What followed was not a Soviet advance. It was a commercial dispute, escalated: an embargo on Iranian oil, the withdrawal of British technical staff, the shuttering of the refinery, the freezing of Iranian assets, and a legal campaign to make Iranian crude untouchable in world markets. The strategy was effective. By 1953 the Iranian government was broke and Mossadegh's domestic coalition — always a coalition, holding together clerics, nationalists, merchants, and the left on a single issue — was fraying under the strain.

The British had asked the Truman administration to help remove him and had been refused. The refusal was not moral fastidiousness; it was an assessment. Mossadegh was a nationalist, not a communist, and the American view was that removing him would strengthen the Tudeh rather than weaken it. What changed thereafter was not the evidence. The facts available to Washington in early 1953 had not become simpler or more alarming than the facts available in 1952. What changed was the administration, the weight assigned to those facts by the people now assigning it, and the range of remedies considered acceptable. A proposal that had been rejected as strategically counterproductive was put again to a different set of principals and found actionable, and the interpretation that made it actionable was thereafter the operative one.

The operation ran for roughly a month. Its case officer arrived in Tehran in July. The first attempt, on the night of 15 August, collapsed: the officer carrying the royal decree dismissing Mossadegh was arrested, the plot was exposed, and the Shah left the country. By any ordinary standard the operation had failed, and Washington's instruction was to withdraw. It was not withdrawn. Over the following three days a second attempt was assembled out of the wreckage of the first, using paid demonstrators, sympathetic officers, and a press campaign that made the dismissal decree public. On 19 August, army units moved. Mossadegh's house was shelled. He surrendered the following day.

The financial outlay was small — the figures cited internally are in the low hundreds of thousands of dollars, with a budget authorization not much above a million. The relevant number is not the cost but the return. In 1954, the nationalized industry was restored to foreign operation under a consortium. The British company, renamed that year, retained forty per cent. Five American majors, who before 1951 had held no position in Iranian oil whatsoever, received forty per cent between them. Shell took fourteen, the French six.

This was not the whole of the transaction, and it would weaken the argument to pretend otherwise. British imperial interest, American Cold War anxiety, the Shah's own ambition, the calculations of a fractured Iranian officer corps, clerical hostility to Mossadegh, and the ordinary opportunism of a collapsing coalition are all part of the causal picture, and Chapter Six is an attempt to weigh them against each other with more discipline than indignation permits. What can be stated plainly here is the narrower part: a commercial injury became an actionable state problem, and the remedy delivered a new arrangement of foreign access to a national industry, on terms substantially better for American firms than anything they had held before the dispute began. That much is not contested and does not need to be inferred. It is in the consortium agreement.

• • •

Set aside, then, the question of whether this was justified, and ask a narrower one: what kind of thing had just been invented?

Four properties define it, and all four are properties of a product rather than of a policy.

It had a *delivery timeline*. From authorization to result was a matter of months. This is extraordinarily rare in statecraft, where the normal relationship between an action and its consequence is measured in administrations, and it is the single most seductive feature of the method. A principal who authorizes a covert removal will personally see whether it worked.

It had a *cost structure*. The outlay was trivial against the value of the asset in dispute — a ratio no other instrument of foreign policy could approach. Not war, not aid, not diplomacy, not the slow construction of institutional relationships. Nothing else on the menu offered a comparable return on a comparable expenditure, and the accounting was not even close.

It had a *client*. Not in a conspiratorial sense; the point is more banal and more damaging than conspiracy. There existed an identifiable set of interests that could specify what it wanted, that could articulate the injury it had suffered, that had access to the officials making the decision, and that was in a position to be satisfied by a particular outcome. Nationalization creates such a client automatically. It converts a diffuse foreign policy question into a specific commercial grievance held by a specific firm with specific representation in the capital. Diffuse interests do not brief principals. Injured firms do.

And it had a *success metric* that could be evaluated. Mossadegh was out. Zahedi was in. The concession was restored. Each of these is a fact that can be checked on a given morning by a person who was not present, and reported upward without qualification. It is worth pausing on how unusual this is. Almost nothing else an intelligence service does produces a verifiable binary outcome attributable to a named team within a fiscal year.

Those four properties, taken together, describe a product. And a product, unlike a decision, has a natural institutional tendency: it wants to be sold again.

• • •

What the metric did not capture is the entire remainder of the ledger, and the remainder is where this book will spend most of its time.

It did not capture what the operation had done to the basis of the restored monarchy's authority. The distinction wanted here is one this book will make repeatedly and should make carefully the first time. The Shah after 1953 was not a ruler without support: he retained a substantial elite coalition, considerable coercive capacity, the loyalty of much of the officer corps, and periods of real popularity, particularly during the boom years. What he had lost was narrower and more corrosive. He had acquired a permanent and accurate answer to the question of how he had been returned to power, and the answer belonged to somebody else. A ruler restored by foreign action against the expressed position of his own parliament holds an office whose claim can always be disputed on grounds he cannot refute, and every subsequent year of that reign was spent attempting to manufacture domestically what the manner of his restoration had made it harder to earn. Approval, compliance, coercive capacity, and legitimacy are four different quantities that move independently, which is a distinction Chapter Twelve will make the centre of the book's theory, and 1953 damaged one of them very badly while leaving the others substantially intact. That combination is more dangerous than general unpopularity, because it is invisible until it is not.

It did not capture that the arrangement would fail in a manner nobody was watching for, twenty-six years later, when the apparatus that had been the regime's guarantee stopped functioning over a period of months while the assessments coming out of Tehran continued to describe a stable situation. Chapter Nine will argue that this was not a coincidence of timing but a consequence of method: an institution that has learned to see politics as a matter of levers and the men who hold them will not perceive a legitimacy collapse until the collapse is complete, because it is not looking at the quantity that is changing.

And it did not capture the thing that is hardest to put a number against, which is that the operation was a gift to every ruler who would subsequently face an inconvenient domestic movement and need an explanation for it. The gift was not propaganda. It was a true story, generalized: that Western intelligence services remove governments they find commercially inconvenient, and that the crowds in the street may not be what they appear. After 1953 that claim did not have to be invented anywhere. It only had to be applied.

None of these costs appeared in any assessment, because none of them was legible on the timescale the metric operated on. They were not concealed. They were not measured. Chapter Seven will argue that this distinction — between a cost hidden and a cost simply outside the accounting frame — is the mechanism by which a democratic state's foreign intelligence capability came to serve interests too small to have justified it, without anyone involved having to be corrupt.

The manual was written in March 1954. Within four months a second operation of the same design was underway in Guatemala, against a government whose principal offence was an agrarian reform statute affecting the uncultivated landholdings of a single American fruit company. The case officer from Tehran was reportedly offered a role and reportedly declined. It hardly mattered. The document had been written. The method no longer depended on the man.

CHAPTER TWO

Obedience as a Supply

In 1993 a small book was published in Bangkok, in an edition of a few thousand copies, written for Burmese exiles who had asked an American academic for something practical. It was ninety-odd pages. It had numbered sections. It described the analysis of a regime's dependencies, the selection of methods appropriate to a given stage of a campaign, the management of repression, the problem of negotiating from weakness, and the errors that had destroyed comparable movements elsewhere. It assessed what had worked. It identified what should be done differently. It was published openly, it was copied and translated without permission or payment, and it was written for the explicit benefit of people who would be asked to do something similar somewhere else.

Two manuals, forty years apart. The first was classified and the second was photocopied, and they are addressed to the same problem.

• • •

The idea underneath the second manual is much older than the manual, and older than the discipline that produced it. A French law student in the sixteenth century, writing an essay he never published, asked why a population obeys a tyrant who is, physically, one man. His answer was that the tyrant has no power the population does not furnish him: he sees with their eyes, walks on their feet, strikes with their hands. Two centuries later Hume put the same observation in the register of political science, noting that since the many are governed by the few, and force lies with the many, government must ultimately rest on opinion. *Force is always on the side of the governed.*

This is not a moral observation, though it has usually been read as one. It is a claim about mechanism, and it can be stated in operational terms. A ruler's orders do not execute themselves. Between the decision and the effect there stands a long chain of people — a clerk who files it, an officer who transmits it, a soldier who acts on it, a magistrate who ratifies it, a broadcaster who announces it, a banker who funds it — and every link in that chain is a person who could, in principle, decline. Power, on this account, is not a substance a ruler possesses. It is a flow that others sustain, and it can be interrupted at any point along its length.

The systematic version of this argument was assembled over three decades in the second half of the twentieth century, and its central analytical device is the disaggregation of a regime into the institutions that supply it. Not "the state" but the specific pillars: the armed forces, the police and internal security organs, the civil bureaucracy, the courts, the state and private media, organized business, the religious authorities, the schools and universities, the trade unions. Each of these supplies something the regime cannot manufacture for itself — enforcement, administration, legality, legitimation, revenue, quiescence — and each can withdraw it, at a cost to itself that varies enormously across pillars and across time.

The strategic consequence is that a campaign against a regime is not primarily a contest with the regime. It is a contest for the pillars. Every action is evaluated by what it does to the calculations of the people staffing them: whether it makes the cost of continued compliance higher, whether it makes the cost of defection lower, and above all whether it changes what each of those people believes the others are about to do. Read through that framework, the canonical catalogue of several hundred techniques of protest, non-cooperation, and intervention becomes an engineering inventory rather than a list of gestures — forms of action that impose different costs, admit different participants, and act on different supports of rule. The sorting is an analytical operation performed on the catalogue, not a feature the catalogue announces about itself, and this book performs it because the comparison in Chapter Three requires it.

A caution belongs here, because the engineering register is load-bearing for this book and is also the register most likely to be misread. Describing collective action as a mechanism says how power is withdrawn. It says nothing whatever about who is entitled to initiate a campaign, who may define its ends, or who bears its risks. Those are separate questions, they are not technical questions, and the answers this book gives to them are the reason it does not treat the two lineages as interchangeable. The mechanism is common property. The authority to set it in motion is not.

The same framework yields the tradition's most counterintuitive proposition, which is that a regime's repression can be made to operate against it. Violence applied to demonstrators who are visibly not fighting back imposes a cost on the people applying it: it strains the loyalty of the security forces asked to do it, it alienates the constituencies watching, and it converts a political dispute into a moral one in front of an audience the regime needs. The mechanism is not that repression is wrong. It is that repression is expensive, and the expense is paid in exactly the currency the regime cannot afford to spend.

• • •

Belgrade, 2000, is the case where this stopped being theory.

The Serbian student movement that formed in 1998 in response to a university law was, by its own later account, an unusually deliberate organization. It was decentralized by design, so that arrests of leaders did not disable it. It was branded, using a single mark that could be reproduced with a stencil by anyone. It used humour heavily and consciously, on the reasoning that a regime can suppress a threat but is damaged by ridicule, and that a movement people find funny is a movement people are less afraid to join. And it treated the police not as the enemy but as the objective: its members were instructed to be courteous to officers, to speak to them about their salaries and their families, to make defection socially available rather than shameful. This was not sentimentality. It was the pillar analysis, applied.

The movement's organizers had read the Bangkok manual, and had worked through the material with a retired American army officer at a series of workshops abroad — a fact that will matter in Chapter Eight and should be stated plainly here rather than discovered later. They had also received funding, through intermediary organizations, that originated with the American government. On 5 October 2000, after a stolen election, several hundred thousand people converged on the federal parliament, the police declined in large numbers to clear them, and Slobodan Milošević left office two days later.

What happened in Belgrade was not a foreign operation, and the argument of this book will be that it could not have been. But it was assisted, it was trained, it was partly financed from abroad, and the assistance came through channels that terminate, several steps back, in the same government that ran the operation in Chapter One. That is the seam this book has to work along, and it will not be resolved by insisting that the two things were different in spirit.

• • •

The tradition became quantitative in 2011, when a dataset of several hundred campaigns for regime change, secession, and expulsion of occupation across the twentieth century was assembled and analysed. The headline finding, drawn from campaigns between 1900 and 2006, was that the nonviolent ones had succeeded in roughly fifty-three per cent of cases against roughly twenty-six per cent for the violent ones — an outcome so contrary to intuition that it has been re-tested repeatedly and has broadly held. It has also, and this matters more than the headline, deteriorated sharply. Extending the same dataset through the 2010s drops the nonviolent success rate to something closer to a third, even as the frequency of nonviolent campaigns reached its historical peak. The same researcher attributes the decline partly to movement-side failures — thinner organization, less training, over-reliance on the internet for coordination, less use of strikes — and partly to the fact that regimes have learned. Among the learned responses she lists first: blaming foreigners and outsiders. Chapter Eight is about where they learned it.

The finding that matters more for this book is the mechanism. Nonviolent campaigns succeed more often because they are larger, and they are larger because they are easier to join. Armed insurgency selects for young men who are physically capable, ideologically committed, and willing to kill. A general strike, a boycott, a stay-at-home, a mass funeral procession admits the middle-aged, the professional, the devout, the frightened, and the merely fed up. The resulting participation is not just greater in number but broader in composition — and the breadth is the operative variable, because a movement that includes shopkeepers, civil servants, teachers, and clergy is a movement whose participants are the relatives, neighbours, and creditors of the people staffing the pillars.

This is where the two halves of the theory close. Mass participation matters not because numbers are inherently persuasive but because participation is what produces defection. When a conscript is ordered to fire on a crowd containing people from his own street, the order is being given under conditions the regime cannot control. The much-cited figure of three and a half per cent should be handled with corresponding care: it is not a threshold built into political life, and nothing about political physics changes at that number. It is a regularity observed within a particular historical dataset — in those cases, campaigns that achieved sustained active participation on that scale did not fail — and its interest lies in what it estimates rather than in what it prescribes, namely the rough point at which a society's pillars can no longer be staffed by people insulated from the movement.

Two limits should be stated now rather than conceded later, because the book's central claim is about durability and would be self-refuting if it overstated what this tradition delivers. The first is that removing a ruler, achieving a transition, achieving democratization, and achieving durable democratic consolidation are four separate accomplishments, and the dataset that supports the headline finding is measuring the first two far more reliably than the last. A campaign can succeed completely on its own terms and leave the security apparatus, the patronage networks, the courts, and the armed factions entirely intact, which is a description of a great many of the cases in the last twenty years. The second is that the tradition's own literature is clearest on how to force a transition and least developed on what happens in the eighteen months afterward, when the properties that made a movement effective become liabilities. Chapter Seventeen returns to this, and does not solve it.

Which returns us to the chain of hands. A campaign of this kind does not defeat a regime. It makes the regime undeliverable, by acting on the people through whom delivery occurs — which is, stated with unhelpful precision, the same thing the operation in Chapter One did, on a different scale, to different people, by different means.

Unhelpful, but not the whole story, and the remainder is the bridge to the next chapter. In Tehran, the people whose behaviour changed were mostly not aware of what they were participating in; they were paid, or deceived, or moved by a script written in another country by people who would not be present for the consequences. In Belgrade, everyone in the square knew precisely what they were doing and would live in the country that resulted. Foreign assistance in the second case could lower costs, supply technique, and widen what seemed possible. It could not supply the thing the method actually runs on, which was several hundred thousand people independently deciding that the risk was now worth taking. That is not a difference of degree.

CHAPTER THREE

Two Theories of Defection

We have been looking at two traditions that are conventionally understood as moral opposites: one covert, coercive, foreign, and undertaken on behalf of commercial interests; the other open, civic, indigenous, and undertaken by people acting for themselves. The moral distinction is real and this book does not propose to dissolve it. But it has functioned, for seventy years, as an obstacle to a question that ought to have been asked much earlier, which is what these two things are as *methods* — and whether, examined as methods, they are addressing the same problem.

They are. Both are theories of defection. Both begin from the observation that a ruler's power runs through other people, and both ask the same operational question: under what conditions do the people who execute a regime's will stop executing it?

The claim needs stating with some precision, because in its loose form it is the sort of claim that gets a book dismissed. At the level of political mechanics, these are rival answers to a single question, and can be compared as such. At the level of legitimacy, agency, and consequence they are not interchangeable at all, and the rest of this book depends on their not being. Who initiates, who bears the risk, who defines the objective, whether the people whose behaviour changes are participants or instruments, and what kind of political order can plausibly stand on the far side — these are not implementation details that fall out in the wash. They are the substance. What follows is a comparison of mechanisms, offered on the understanding that mechanism is not the whole of the object.

With that said, the two methods separate cleanly along three dimensions.

• • •

How many defections, and of what weight. The coup seeks a very small number of highly decisive defections: an armoured brigade commander, a chief of the general staff, the head of a palace guard, a handful of colonels who control specific units in a specific city on a specific night. Perhaps a dozen consequential decisions determine the outcome. Civil resistance seeks a very large number of individually inconsequential defections — a shopkeeper who closes, a clerk who works slowly, a broadcaster who resigns, a sergeant who takes his time relaying an order — none of which decides anything, and whose aggregation eventually decides everything. The coup buys weight. The movement accumulates it.

What induces the defection. This is the deeper difference. In a coup, the defector is moved by an inducement offered to him individually and privately: money, promotion, ideological affinity, a settled grudge, fear of what happens to him if the plot succeeds without him. His decision is, in principle, independent of what anyone outside the plot believes. In civil resistance, almost nothing is offered to the defector at all. What changes is his estimate of what other people are going to do. The sergeant who slows down does so because he has concluded that the arrangement is ending and that he does not wish to be conspicuously associated with its final days. Nobody paid him. He revised a forecast.

What the method requires in order to function. The coup requires secrecy, not incidentally but constitutively. A coup announced in advance is a coup suppressed in advance; exposure is not a risk to the operation but its terminal condition, which is why the first Tehran attempt collapsed the moment the decree-bearer was arrested. Civil resistance requires the exact opposite. Its entire mechanism is informational: it works by making private disaffection publicly visible, so that people who each believed themselves isolated can observe that they are not. A movement conducted in secret is not a movement. It is a conspiracy, and it will fail as a movement even if it succeeds as a conspiracy.

That third difference is not a detail of tradecraft. It is the hinge of this entire book, and Part Four will spend seven chapters on it. A method whose effect does not depend on anyone else's beliefs can be applied to a society from outside, against its inclinations, on a schedule. A method whose effect *consists* *entirely* of other people's beliefs about each other cannot be. The first can be imposed; that is its attraction and, on the evidence of the next hundred pages, its failure mode. The second can only be amplified where the underlying belief is already shifting, which means that where nothing latent exists, the method does nothing at all. This is the asymmetry the doctrine in Part Five rests on, and it is worth noticing that it falls out of a comparison of mechanisms rather than from any premise about the virtue of the people involved.

The failure modes follow from the same structure and confirm it. A coup fails by exposure — someone talks, and the plot is rolled up. A civil resistance campaign fails by isolation — nobody talks, and the participants conclude they are alone. One method is killed by information and the other is starved of it. It would be difficult to design two techniques more precisely opposed in their relationship to the same variable.

• • •

Before drawing the conclusion, one further difference has to be entered on the record, because a reader who has followed the comparison this far is entitled to worry that the book is building toward an equivalence it will then exploit.

It is not. The two methods differ in a respect that no amount of mechanical comparison touches, which is *ownership*. In the first, the people whose behaviour is altered are largely unaware of what they are part of. The crowds in Tehran on 19 August contained men who had been paid to be there and men who had not, and the distinction between them was invisible in the street and material to everything that followed. The script was written elsewhere, by people who would not live in the country it produced, and who could not be held to account by anyone in it. In the second, every participant is also a principal. The risk is borne by the person taking it. The objective is defined, revised, and abandoned by the people who will live with the result. Foreign assistance can lower the cost of coordination, supply technique, and enlarge the sense of what is possible — and Chapter Nineteen will argue that this is a considerable amount, worth doing, and worth doing openly. What it cannot do is supply the willingness. The distributed decision to withhold obedience is the one input the method requires and the one input no external party can manufacture, which is why the mechanical similarity between the two lineages is compatible with their being different in kind.

Hold both of these at once and the comparison becomes usable rather than corrosive: the same physics, different owners, and the ownership is what determines whether anything built on it stands up.

• • •

All of which sharpens the question this book exists to answer.

By any measure available to us now, and by several available at the time, the second method performs better. It succeeds more often. Its successes last longer and are more likely to produce something recognisably democratic five years on. It is cheaper in the currency that matters, which is not money. It cannot be exposed, because it operates in the open. And it is a domain in which democratic states hold an advantage their adversaries structurally cannot match, since the method requires operating publicly and an authoritarian state that operates publicly is doing something other than authoritarianism.

Democratic states industrialized the first one anyway. They built an institutional capability for it, wrote its manuals, trained its officers, and ran it repeatedly for a quarter of a century, including in places where the alternative was visibly working elsewhere.

There is an obvious objection to putting it this way, and it should be met now rather than left to accumulate. In 1953 the civic lineage was not yet available in the form described in Chapter Two: the systematization was two decades off, the dataset was six decades off, and what existed was a set of demonstrated cases rather than a doctrine. This is true, and it is a real limit on how much can be asked of the men who made the decision in 1953.

It is not a limit on the argument, for two reasons. The first is that the demonstrated cases were substantial and recent — an independence movement had, within living memory of everyone in the room, removed the same imperial power from a subcontinent by methods no one in the room considered worth studying. The second and more damaging reason is that the choice was not made once in 1953. It was re-made in 1954, in 1964, in 1965, in 1970, in 1973, and repeatedly thereafter, in a period during which the alternative moved from demonstrated to documented to systematized to quantified, and during which the accumulated failures of the first method became progressively harder to overlook. A decision made once under uncertainty is a mistake. A decision re-made for forty years against a widening body of contrary evidence is a preference, and preferences have explanations.

The explanation is not that the men involved were stupid, and it is not, or not principally, that they were wicked. It is that the first method had the four properties set out in Chapter One — a delivery timeline, a cost structure, a client, and a checkable success metric — and the second method has none of them. It cannot be scheduled. It cannot be attributed. It has no completion date. It offers nothing a principal can be shown at the end of a fiscal year, and it produces outcomes that the people who paid for them cannot control.

An institution does not choose between these on the merits. It chooses the one its accounting can see.

Part Two is the record of what that choice produced, and of the interests that were served by making it.

CHAPTER FOUR

Nationalization as Trigger Condition

The compensation offered to the United Fruit Company for the land the Guatemalan government took from it in 1953 was calculated from the value the company had itself declared for that land, in its own filings, for the purpose of paying tax on it. This was not an improvisation or an insult. It was written into the statute: expropriated holdings would be compensated at declared taxable value, in twenty-five-year bonds. The company had spent years declaring the land to be worth very little, because the land was worth very little to it — most of it was uncultivated, held in reserve against disease in the planted acreage and against the possibility of a competitor acquiring it. The government proposed to take it at the company's word.

The figure came to $627,572. The United States government, filing on the company's behalf, put the correct valuation at $15,854,849 — a little over twenty-five times as much. The land at issue was some 234,000 acres near Tiquisate, expropriated in February 1953 under a statute passed the previous June, out of the roughly 550,000 acres the company held in a country where it was the largest single landowner.

It is worth sitting with the shape of that dispute, because it is the shape of nearly everything in this part of the book. A sovereign state passed a law of general application. The law was applied to a foreign firm on the same terms as to domestic landholders, of whom it affected a great many. The firm objected to the valuation. And the objection to the valuation became, within a year, a matter for the Central Intelligence Agency.

Two questions run underneath every case in Part Two, and it is worth naming them now so that the historical detail can be read as evidence rather than as narrative. The first is Chapter Six's: what actually converts a commercial injury into a covert operation, given that most commercial injuries do not become one? The second is Chapter Seven's: how did the institution that performed these operations come to record them as successes, and what did that recording do to the institution? Everything in this chapter and the next is assembled to answer those two questions. Where a detail does not bear on one of them, it has been left out, and a reader who wants the fuller narrative of any of these episodes will find it well told elsewhere.

• • •

Iran and Guatemala are treated here as a matched pair, and the pairing is the argument. Two governments could hardly have been less alike. Mossadegh was a patrician constitutionalist in his seventies, leading a fissile coalition through parliament. Jacobo Árbenz was a young army officer turned president, elected with a substantial majority, governing with a programme of land redistribution in a country where two per cent of the population held most of the arable land. One was a nationalist with clerical allies; the other was a reformer with communists in his advisory circle. Their politics, their societies, their regions, and their strategic significance had almost nothing in common.

What they had in common was that each had impaired the position of a specific foreign firm, and each was removed by covert action within fourteen months of the other, by two operations whose planning documents were substantially the same document.

The Guatemalan case is the cleaner of the two on the security question, and that is precisely why it is analytically useful. The communist party in Guatemala was small. It held a handful of seats in a congress of fifty-six. Árbenz was not a member and had not been. There were no Soviet arms in the country until, in the operation's final months, a shipment of Czech weapons arrived — a shipment which the government had sought precisely because the United States had embargoed arms sales to it, and which was then cited as evidence of the alignment the embargo had produced. There was no base, no treaty, no fleet, and no plausible route by which Guatemalan agrarian policy could threaten the security of the United States. Whatever else was being defended in the summer of 1954, it was not a border.

The operation itself ran for a fortnight in June 1954 and was principally theatrical. A small force crossed from Honduras and advanced a short distance. Aircraft flown by American contract pilots bombed and strafed. A radio station, purporting to broadcast from the jungle and in fact broadcasting from outside the country, described the movements of an army that did not exist, in numbers it did not have, converging on positions it had not reached. The purpose was not to defeat the Guatemalan army, which could have destroyed the invading force in an afternoon. The purpose was to persuade the Guatemalan army that it need not try. It worked. The officer corps declined to fight, Árbenz was informed that the army would not defend him, and he resigned on 27 June.

Note the mechanism, because it is the argument of Chapter Three arriving in the historical record. What decided the outcome in Guatemala was not force. It was a small number of colonels revising their estimate of what was about to happen and of what the other colonels would do. The operation was a machine for manufacturing that revision, and it manufactured it with a radio transmitter and a handful of aircraft. This is the covert lineage doing, badly and from outside, the thing that Chapter Two describes movements doing from within.

The land taken under Decree 900 was returned. The reform was reversed. What followed, over the subsequent four decades, was a sequence of military governments and a civil war in which something on the order of two hundred thousand people were killed, the overwhelming majority of them by state forces, and in which a truth commission would later find that acts of genocide had been committed against Maya communities. Whether or not the operation of 1954 caused all of that is a question this book will not pretend to settle. What can be said is narrower and sufficient: the operation removed the last elected government Guatemala would have for more than thirty years, and no assessment written in 1954 contains any estimate of that possibility, because no assessment written in 1954 was asking.

• • •

Now the negative case, which belongs here rather than three chapters later, because the pairing is what makes the timing evidence legible.

In March 1938 the government of Mexico expropriated the holdings of seventeen foreign oil companies, British and American, in their entirety. The immediate trigger was a labour dispute in which the companies had declined to comply with an award of the Mexican supreme court. The expropriation was not partial, it was not confined to unused assets, and it was not accompanied by an agreed valuation. By any measure it was a more aggressive act against foreign capital than anything either Mossadegh or Árbenz attempted.

The response was a boycott of Mexican oil, a British rupture of diplomatic relations, and a sustained American diplomatic campaign for compensation. It was not an operation. The American ambassador in Mexico City argued consistently against coercion, the administration was committed to a policy of non-intervention in the hemisphere that it had reasons to value, and a war was visibly approaching in which Mexican alignment and Mexican oil would matter more than the claims of the companies. Compensation was negotiated over several years and concluded at a figure the companies considered derisory. The Mexican state kept the industry, and has it still.

The honest difficulty with Mexico as a negative case has to be stated, because a reader who knows the period will state it anyway. In 1938 the United States did not possess the instrument. There was no Central Intelligence Agency, no Office of Strategic Services, no standing covert action capability of the kind that existed by 1953. A government cannot decline to use a tool it does not have. Mexico is therefore a weak test of preference and a strong illustration of repertoire: it shows what the available responses to expropriation looked like before the instrument existed, and it establishes that boycott, diplomatic pressure, litigation, and negotiated settlement were a complete and functional set of options that resolved a larger commercial injury than either of the cases that opened this chapter.

That is not nothing. It means the question posed to 1953 and 1954 is not "what else could they have done" but "why did they stop doing what had already worked."

• • •

The pair and its negative yield a preliminary finding, which Chapter Six will test properly and which is stated here provisionally so that Chapter Five can be read against it.

Commercial injury is not sufficient. Mexico is the proof. Something else has to be present before an injury becomes an operation, and the candidate conditions are visible in outline already: whether the injured firm has standing access to the officials deciding; whether the dispute can be described in the vocabulary of the wider strategic contest; and, most importantly and least discussed, whether there exists inside the target country a partner willing and able to act. In Tehran there were officers, a monarch, and a mobilizable clergy. In Guatemala City there was an officer corps with no attachment to the government it served. In Mexico City in 1938 there was a revolutionary state with an army that had been built by the party in power, and no available partner at all.

Nationalization, on this account, is not the cause of intervention. It is the trigger condition: the event that converts a diffuse foreign policy question into a specific, articulable grievance held by a party with the access to press it. Whether the trigger fires depends on something else entirely.

CHAPTER FIVE

The Method Acquires a Bureaucracy

At the end of March 1964 a United States naval task force was steaming toward the coast of Brazil. It included an aircraft carrier, destroyers, tankers carrying fuel oil and aviation gasoline, and a ship loaded with small arms and ammunition. The contingency it was assembled against was the possibility that the Brazilian officers who were about to remove their president might meet resistance and run short of supplies. The force never arrived, because the operation it was supporting concluded before it needed to. The president left the country on 4 April. The ships turned around.

The interesting thing about that task force is not that it was dispatched. It is how quickly it was assembled, how completely it was integrated with a political plan being executed by foreign nationals, and how little of it required improvisation. Eleven years after Tehran, the removal of a foreign head of government had acquired logistics. There was an embassy running political liaison, a military attaché network running the officer relationships, a contingency package with a name, and a fleet order. The activity that in 1953 had been an ad hoc operation run by a handful of officers with a suitcase of cash was now a coordinated interagency undertaking with a supply chain.

That is what this chapter means by a bureaucracy. Not a conspiracy and not a policy, but the ordinary institutional furniture that accumulates around any activity a state performs repeatedly: doctrine, a personnel pipeline, standing regional relationships, a body of precedent, and an inventory.

• • •

João Goulart's offences against foreign capital were real and modest. He had signed a law restricting the annual profits foreign firms could remit abroad. A subsidiary of an American telecommunications company had been expropriated at state level by his brother-in-law, at a valuation the company disputed. He proposed a package of structural reforms including land redistribution, and he had moved to nationalize refineries. Against this he was a wealthy landowner, a constitutional vice-president who had succeeded lawfully, and a man whose principal political skill was the avoidance of confrontation.

Set beside him the government installed by military coup in Peru four years later, which is the negative case that belongs in this chapter rather than the next. Within a week of taking power, that government seized the Peruvian holdings of a subsidiary of Standard Oil of New Jersey outright, occupied the facilities with troops, and declined to pay compensation on the theory that the company owed more in back taxes than the assets were worth. Over the following six years it nationalized a telecommunications company, the fishmeal industry, sugar estates, and a large American mining concern. It bought Soviet arms. It called itself revolutionary and meant it.

The American response was a statutory aid cut-off that the administration then went to some trouble to avoid formally invoking, a long and irritable negotiation, and eventually a lump-sum settlement in 1974 under which the Peruvian government paid a global figure covering all the expropriations, distributed by the United States among the claimants — of which the oil company, whose seizure had started everything, received nothing at all. The government that had done all this was removed in 1975 by other Peruvian generals, for Peruvian reasons.

Goulart, who did considerably less, was removed with American logistical support standing offshore. The Peruvian junta, which did considerably more, was met with lawyers. Any account of this period that cannot accommodate both facts is not an account; it is an accusation. Chapter Six is where the difference is tested. What can be observed here is that the two cases differ on the variable identified at the end of Chapter Four: in Brazil there was a large, politically organized, deeply anti-communist officer corps with decades of institutional contact with the United States and an appetite for governing. In Peru the officer corps was the government, and it was the one doing the expropriating.

• • •

Indonesia is the moral centre of this book and will be written without cushioning.

Beginning in October 1965, following the killing of six generals by a group of officers and the counter-move by the general who emerged from it, the Indonesian army and the militias it armed and directed conducted a programme of mass killing directed at members and supposed members of the Indonesian Communist Party, their families, and a range of others swept up by local score-settling. The killing continued for approximately six months. Estimates of the dead begin at half a million and go higher. Bodies were put into rivers in numbers sufficient to obstruct them. Somewhere in the region of a million more people were imprisoned, many for over a decade, without trial.

The United States did not organize this. That should be stated plainly because it matters to the argument, and because the argument does not require the stronger claim. What the United States did was to encourage it, to signal approval of it privately and repeatedly to the officers conducting it, to supply communications equipment and other material to the army while it was underway, and — through an officer at the embassy in Jakarta who later described doing so — to pass to that army lists of names of party members compiled from open sources. Embassy reporting from the period, declassified in bulk in 2017, records the scale of the killing as it was occurring and records the assessment that it was, on balance, a favourable development.

The relevant point for a chapter about institutionalization is what this represents as a change in the product. Tehran was an operation the United States ran. Guatemala was an operation the United States ran. Jakarta was something else: a domestic catastrophe, generated by Indonesian actors for Indonesian reasons, which the United States encouraged, supplied, and counted as a success. The method had discovered that it did not need to run anything. It could underwrite. This is cheaper, more deniable, and — as the subsequent decade in the Southern Cone would demonstrate — far more scalable.

The commercial sequel is not the cause of what happened and should not be presented as one. It is nonetheless part of the ledger. Within two years the new government had passed a foreign investment law, concluded a mining contract with an American company that would become one of the largest such operations in the world, and reopened the oil sector on terms the previous government had refused. The thing that had been unobtainable by negotiation for a decade became available within twenty-four months, and the people who made it available were the people who had conducted the killings.

• • •

Chile in 1973 is the case where every element assembled since 1953 was present at once, and it is the last of the classical operations, in the sense that it was the last conducted before the disclosures of the mid-1970s changed what could be done without consequence.

Salvador Allende won a plurality in September 1970. Before he was confirmed in office, two parallel efforts were underway: one political and financial, aimed at preventing confirmation, and one military, aimed at inducing a coup. The second required removing the army commander-in-chief, a constitutionalist who would not act, and the attempt to kidnap him left him dead. The plot failed, Allende was confirmed, and the effort turned to the slower instruments: credit denial through the international institutions, support for opposition media and parties, and — in a phrase that survives in the notes of the meeting where it was said — an instruction to make the economy scream.

The commercial position was copper, and copper had been nationalized in July 1971 by constitutional amendment, passed unanimously by a congress the government did not control. The compensation formula deducted from book value what were designated excess profits earned over prior decades, which left the two large American producers with nothing, and in one calculation less than nothing.

Zambia had done something adjacent two years earlier. In August 1969 the Zambian government announced that it would take a controlling interest in the copper mines, and completed the acquisition the following year. Copper was to Zambia what it was to Chile, which is to say the entire export economy.

The government was nonaligned, hosted southern African liberation movements, and was not conspicuously friendly. It was not targeted. The differences are instructive and Chapter Six will weigh them. The Zambian acquisition was announced at Matero Hall on 11 August 1969, settled with the companies inside three months, compensated at book value in state bonds redeemable over eight to twelve years at six per cent, and — decisively — left Anglo American and Roan Selection Trust in place as managers and sales agents under ten-year contracts paying a percentage of gross sales and of profits. The firms lost majority ownership and kept an income. Copper was ninety-six per cent of Zambian exports; this was not a small expropriation, and it produced no crisis at all. There was, additionally, no Zambian military faction remotely interested in taking power, and no domestic political crisis to work with.

Chile's army needed no manufacturing by 1973, and the honest reading of the record is that the coup, when it came, was made in Chile by Chileans in conditions of severe domestic breakdown. The American contribution over three years was to deepen the breakdown, sustain the opposition, and make clear that a coup would be welcomed. Whether that contribution was decisive is genuinely contested and this book does not need it to have been. What is not contested is that it was made, that it was made continuously for three years, and that the object throughout was the removal of an elected government whose principal offence against the United States was the expropriation of two mining companies.

• • •

Twenty years separate Tehran from Santiago, and across them a capability matured in four respects.

It acquired *doctrine*: a body of internal precedent about what works, transmitted through documents like the one Chapter One opened on. It acquired a *pipeline*: officers who had done this before, and who moved between stations carrying the practice with them. It acquired *standing relationships*, particularly with Latin American officer corps, built through training programmes and attaché networks that existed continuously and could be activated rather than constructed. And it acquired a *tolerance* — the discovery, at Jakarta, that the method's most efficient form was to enable violence conducted by others and to record the result as a favourable development.

None of these is a decision. Each is a form of institutional memory, and institutional memory is what makes a practice cheap. By 1970 the removal of an inconvenient government was no longer an extraordinary undertaking requiring extraordinary justification. It was a thing the government of the United States knew how to do, had people who did it, and had done recently enough that the doing of it required no particular explanation.

That is what a bureaucracy is for. It makes the second time easier than the first, and the tenth easier than the second, and it does this without anyone ever deciding that it should.

And it co-evolves with the thing that measures it. Each of these operations was recorded as a success by the standard Chapter One set out, and each recorded success lowered the threshold of justification for the next, which produced more operations, which produced more recorded successes. The doctrine, the pipeline, the standing relationships, and the tolerance were not built and then measured; they were built *by* being measured, in a loop that had no term in it for anything that happened after the fiscal year closed. Chapter Six asks what determined which disputes entered that loop. Chapter Seven asks why nothing ever left it.

CHAPTER SIX

Testing the Economic Thesis

The preceding two chapters have assembled a pattern and an insinuation. The pattern is that governments which impaired the position of major Western commercial interests were, with some frequency, covertly removed. The insinuation is that the first fact explains the second. This chapter is an attempt to find out whether it does, conducted with the awareness that a book making this argument has every incentive to conclude that it does, and that a reader is right to discount an argument whose author needed it to come out a particular way.

There are three candidate explanations for the pattern and they are not mutually exclusive.

The *economic thesis* holds that intervention was substantially driven by the protection of specific commercial interests, and that the Cold War supplied a vocabulary in which those interests could be described as security concerns.

The *security thesis* holds that intervention was driven by an assessment, sincere and often mistaken, of Soviet advance, and that the commercial injuries were incidental — either coincidental, or evidence for the security assessment, since expropriation was a thing that leftward governments did.

The *null hypothesis*, which is the one most often overlooked and the hardest to dismiss, holds that intervention followed *opportunity* rather than motive: that the United States intervened where intervention was cheap and available, that cheapness and availability were determined by local conditions having nothing to do with either oil or communism, and that the apparent pattern is an artefact of where the instrument happened to fit.

• • •

The method has to begin by defining a population, because the fastest way to prove an economic thesis is to select the cases that prove it. The population used here is every government, between the establishment of a standing American covert capability in 1947 and the end of the 1970s, that materially injured a major Western commercial interest — through expropriation, nationalization, unilateral contract revision, or comparable measures — whether or not anything happened to it afterward. The definition is deliberately about the injury rather than about the outcome, because a population defined by outcome cannot test anything.

The positive cases are the ones already discussed, together with several not yet mentioned. The negative cases — governments that inflicted comparable or greater injury and were not covertly removed — include, at minimum, the following.

Bolivia, 1952. A revolutionary movement took power by insurrection, dissolved the existing army, armed worker and peasant militias, and nationalized the three great tin companies that constituted the country's export economy. This was, in scale relative to the national economy, a more radical expropriation than Chile's. The United States extended diplomatic recognition within weeks and then, under the same administration that had just removed Mossadegh and Árbenz, made Bolivia among the largest per-capita recipients of American aid in the hemisphere. The reasoning, on the record, was that the governing party was containable and anti-communist enough, that compensation was being offered, and that the realistic alternatives to it were worse.

Egypt, 1956. The Suez Canal Company was nationalized in July. The response of the United States was to oppose, and then to break, the military intervention launched by two allied European powers and Israel to reverse it, using financial pressure against sterling to force a withdrawal. There is no case in the period of a larger commercial expropriation being met with a more emphatic American defence of the expropriating state.

Peru, 1968; Zambia, 1969; Venezuela, 1976. Discussed in Chapter Five, with the addition of

Venezuela, which nationalized its oil industry outright in 1976, compensated the concessionaires at book value on terms they thought inadequate, and experienced no consequence whatever beyond continued commercial relations.

India, across the Nehru period. Sustained nationalization across insurance, banking, and heavy industry; nonalignment as declared policy; Soviet arms purchases; and, throughout, very substantial American aid.

And the population must include a case that cuts the other way, since a test that only admits inconvenient negatives is only half a test. *South Vietnam, 1963:* an allied head of government removed with American encouragement in a coup that no commercial interest requested, against no expropriation, for reasons entirely internal to the prosecution of a war. Whatever explains 1963 is not the economic thesis.

• • •

Laid out this way, the cases separate on variables that can be named.

Compensation behaviour. Bolivia offered, Peru eventually settled, Zambia negotiated and left the firms as paid managers, Venezuela paid book value, Mexico paid over years. Iran offered but on terms the company would not discuss; Guatemala offered at declared tax value; Chile's formula produced a negative number. This variable does real work, but less than it appears: what distinguishes the cases is not whether compensation was offered but whether the injured firm was left with a continuing position. A firm that keeps a management contract does not brief a principal. A firm written down to zero does.

The availability of a partner. This is the variable that discriminates most sharply, and it is the finding this chapter is likeliest to end on. Every positive case features an organized, willing, capable domestic actor with a prior institutional relationship to the United States and an appetite for power — a praetorian officer corps, most often, and one whose ambitions could be met by the removal. Every negative case lacks one. In Bolivia the army had been dissolved. In Peru the officers were the government. In Zambia there was no faction. In Egypt and India the state was consolidated and the military attached to it. Opportunity, in other words, is not a rival explanation to be dismissed. It appears to be a necessary condition.

Alignment and the availability of a patron. Cases in which the target could be described, with whatever accuracy, as moving into the Soviet orbit are systematically more likely to be positive. This is the security thesis's strongest showing and it should be conceded without hedging. But the direction of causation is not clean, and Guatemala shows why: an arms embargo produced an arms purchase from the eastern bloc, which was then cited as evidence of alignment. Where the record permits, the chapter tracks whether the security characterization preceded or followed the commercial dispute. In several cases it followed.

Access. Whether the injured firm had standing, informal, personal access to the officials deciding. This is the least tractable variable and probably the most important, and it is largely unmeasurable, which is a limitation and not a defence.

• • •

The finding this chapter arrives at is narrower than the argument it is often assumed to be making, and stronger for the narrowing.

Commercial injury does not cause intervention. It is neither sufficient — Mexico, Bolivia, Egypt, Peru, Zambia, Venezuela, India — nor, on the evidence of 1963, strictly necessary. What commercial injury does is determine *which disputes reach the table*. It converts a diffuse and arguable foreign policy question into a specific, quantified, urgently articulated grievance held by a party with the standing to press it in the rooms where decisions are made. It sets the agenda. Whether the item on the agenda becomes an operation is then determined largely by whether a partner exists.

This is a weaker claim about causation and a much more damaging claim about institutions. A story in which commercial interests dictate policy is a story about corruption, and corruption is in principle correctable by better people. A story in which commercial interests determine which questions are asked, while a separate variable determines the answer, is a story about an agenda-setting mechanism that no individual controls and no one has to be corrupt to operate. Nobody in that story is taking money. The firm with the grievance is simply the only party in the room with a specific request, and specificity is what gets acted on.

Four limitations bear on all of this and belong in the text rather than in a note. The population is drawn from disclosed cases, and disclosure is not random: operations that failed or were exposed are overrepresented, and successful ones may be missing entirely. Considered-and-rejected operations are almost wholly unobservable, which means the denominator is unknown. The capability did not exist before 1947, which weakens the pre-war negatives as tests of preference. And the security characterizations in the record were written by people who knew they might be read, which makes sincerity difficult to assess from documents alone. None of these is fatal. All of them mean the finding should be stated as it has been stated here, and not more strongly.

CHAPTER SEVEN

The Metric That Ate the Mission

Return to the document that opened this book. It has an assessment section. The assessment is that the operation succeeded.

By what it measures, it did. The prime minister was removed. The successor took office. The oil arrangement was restored on terms favourable to the parties that had commissioned the effort. Each of those is a fact, checkable within months by someone who was not present, attributable to a named team, and reportable upward without qualification. The document is not lying and its author was not a propagandist. He was applying the only standard available to him.

This chapter is about that standard, and about the proposition that a measurement regime is not a neutral instrument for recording outcomes. It is a theory of what outcomes are.

• • •

Consider the two statements an assessment might make.

Regime changed. Observable in months. Binary. Attributable to specific officers. Legible to a principal with no regional knowledge. Survivable in a briefing.

Legitimate governance obtained. Observable over decades. Continuous rather than binary, and contested at every point on the scale. Attributable to nobody, because by the time the answer is available the people responsible have retired, and the causal chain has passed through a dozen other events. Not briefable at all, since a principal asking what an operation achieved cannot be told that the answer will be available in 1979.

An institution rewarded on the first of these will produce the first indefinitely. This does not require anyone to prefer it, believe in it, or notice. It requires only that careers, budgets, and reputations attach to things that can be shown, and that the second statement can never be shown to anyone in a position to reward it. The people who ran these operations were not, by and large, cynics about the second question. Many of them cared about it a great deal and said so in private. It simply had no place to go.

The asymmetry compounds because the two kinds of cost are differently shaped. An operation's benefits are concrete, immediate, and attributable. Its costs are counterfactual, delayed, and diffuse — a regime with a compromised claim to authority, a movement elsewhere that will be accused of foreign sponsorship in 1974, an assessment capability that atrophies because it is never asked to produce anything. There is no accounting frame in which those appear as debits against the operation that produced them. Any institution facing that asymmetry will over-produce the activity, and will do so while its members believe themselves to be making careful judgements, because the thing that would correct the judgement is not in the frame.

This is the mechanism Chapter One promised and Chapter Six sharpened. The commercial interest supplied the agenda. The measurement regime supplied the impunity — not by concealing costs, which would have required someone to know them, but by ensuring that nobody was ever assigned to find out.

• • •

Which brings us to the reform era, and to why it did not work.

Between 1974 and 1981 the American covert action apparatus was subjected to the most serious external scrutiny it had ever faced. A congressional committee conducted a systematic public investigation, published findings on assassination plots and on the Chilean operation, and forced a great deal into the open. Statutes were passed requiring the president to make a formal written finding before a covert operation could proceed and to notify congressional committees. Standing oversight committees were created. An executive order prohibited political assassination. A subsequent order codified the whole framework.

Every one of these reforms addressed *authorization*. Who must sign. Who must be told. What may not be done. Not one of them addressed *evaluation*. No statute of that period, and none since, requires anyone to determine what a covert operation produced five years later, ten years later, or twenty; to compare that against what was predicted; or to record the comparison anywhere that the next authorization would encounter it. The reforms built a gate at the entrance and left the exit unmonitored.

The consequence is that the institutional memory described in Chapter Five is memory of technique only. What is transmitted forward is how to do it. What is not transmitted forward is what happened afterwards, because nobody was ever tasked to write that down, and a practice whose successes are documented and whose failures are not will look, to the people inheriting it, like a practice that works.

• • •

The deeper claim of this chapter is that the metric was not merely permissive. It was formative.

A measurement regime commits its users to an ontology. To measure who holds an office is to adopt, in practice and regardless of what anyone believes in theory, a theory in which political outcomes consist of who holds offices. Over twenty-five years and several thousand officer-careers, that theory stops being a measurement convention and becomes the way the institution sees. It shapes collection priorities, because you collect against what you will be assessed on. It shapes recruitment, because you recruit sources close to the people who hold offices. It shapes analysis, because analysis is written for consumers who ask about the men at the top. And it shapes, above all, what registers as an event: in an institution organized this way, a colonel's loyalty is information and a general strike is background.

That is the sentence this chapter exists to earn, because it is the hinge between Part Two and everything after it. The apparatus did not fail to see the Iranian revolution coming because it was careless. It failed because it had spent a quarter of a century building an instrument calibrated to a different quantity, and the quantity that was changing in Iran in 1978 was not one it measured. Part Three is about that failure, and about the further discovery that the instrument, having been calibrated this way, could no longer be trusted to tell the difference between a society that was about to come apart and one that was merely unhappy.

The manual was right. The operation succeeded. That is the problem.

CHAPTER EIGHT

The Argument We Gave Them

On a Tuesday morning in the summer of 2012, the staff of a Russian environmental organization that had spent fifteen years documenting industrial violations in the Urals arrived at their office to find a team of inspectors waiting. The inspection lasted several days. At the end of it the organization received a notice: it had been identified as a foreign agent. This was the term the new statute used, and the term was not chosen carelessly. In Russian the phrase carries the resonance of the Cold War, of espionage, of working for somebody else's country against your own. It means something more than it translates. The organization had received a small grant from a Western foundation. It worked on rivers and smelters. It had never operated as a political actor in any ordinary sense of that phrase. It was given six months to comply or close.

What happened to that organization is the subject of the last part of this chapter. What happened before that organization was founded is the subject of the first.

• • •

On 1 August 1938, the United States Congress passed a statute requiring persons acting as agents of foreign principals to register with the Department of Justice, disclose their activities, and label their materials. The Foreign Agents Registration Act was directed at a genuine problem: the organized distribution of propaganda on behalf of the Nazi German government through American front organizations, some of which had been active for several years and some of which were not at all transparent about who was paying for the literature they handed out. The statute was also, and self-consciously, an American law built on American constitutional constraints: it required disclosure rather than prohibition, applied only to explicitly political activity, and was drafted in a way that created no power to simply designate an organization out of existence.

Stripped to its skeleton, however, the act specified three elements that have proven internationally transferable: a trigger based on the receipt of foreign resources; a category of covered activity defined as political or quasi-political; and a registration and labelling regime whose non-compliance could be sanctioned. Those three elements constitute an architecture, and the architecture is what diffused. The specific American version — with its constitutional hedges, its disclosure-only remedy, and its relatively narrow application — did not travel. The skeleton did.

What follows is a table of the instrument as it exists in four jurisdictions, comparing trigger, designation consequence, and the operative terms that make the architecture expansible. The table is illustrative rather than exhaustive, and the entries should be verified against current primary texts, which in several cases have been amended recently.

Jurisdiction	Instrument & date	Trigger	Consequence of designation
Russia	Foreign agents law (2012); expanded 2017, 2019, 2020, 2022; Undesirable organizations law (2015)	Receipt of foreign funding plus activity construed as political; expanded to include "foreign influence" without funding requirement; ultimately broadened to any activity that could affect state interests	Mandatory labelling on all materials; onerous quarterly and annual reporting; criminal liability for non-compliance; for "undesirable" designation, outright prohibition and criminal penalties for participation
India	Foreign Contribution Regulation Act, as amended 2010, 2020	Receipt of foreign contributions by registered entities; prior permission required for new recipients; amendment narrowed eligibility severely and imposed centralized	Registration revocable at administrative discretion without judicial review; prohibition on transfer of received funds to unregistered entities; restrictions on administrative expenditure share; has been applied to cancel

		banking requirements	registration of thousands of organizations
Nicaragua	Foreign agents law (2020)	Receipt of foreign funds or assets; broadly defined to include technical assistance and training	Registration and identity disclosure; prohibition on political participation by designated persons; served as legal basis for mass cancellation of NGO legal status in 2021–2022 affecting several hundred organizations
Georgia	Transparency of Foreign Influence law (2024)	Receipt of more than twenty per cent of funding from foreign sources	Registration and disclosure; monitoring by revenue service; financial penalties; enacted over sustained domestic protest and significant international criticism

Entries based on publicly available statutory texts and secondary analysis as of 2024. Amendment histories are material and ongoing; characterizations should be verified against current primary texts.

Several things are notable about this table. First, the convergence: instruments drafted in unrelated legal cultures, by governments with different relationships to the rule of law, have arrived at substantially the same architecture. Convergent evolution of this kind indicates diffusion of a template rather than parallel invention. Second, the elasticity: in every case the instrument as enacted is significantly broader than what would be required to address the stated concern — foreign-funded electoral interference — and in several cases has been amended to broaden it further after initial enactment. The architecture is designed to be expanded. Third, and least comfortably, the direction of diffusion: the skeleton traveled eastward and southward from the country that built it, and there is no available diplomatic mechanism by which the country of origin can object to the use of its own template without simultaneously objecting to the principle.

The last point is the argument of this chapter and the one most in need of being stated clearly. The United States government, between 1953 and 1973, provided every autocrat with a permanently usable and partially accurate explanation for domestic opposition: that it is externally manufactured, externally funded, and responsive to foreign interests rather than domestic conditions. That explanation is now encoded in law in dozens of jurisdictions and is operationally applied to organizations working on rivers and smelters, on election monitoring, on documentation of extrajudicial killings, and on every other domain of civic activity that a government finds inconvenient. The organizations that are caught in those laws did not cause this. The operations described in Part Two did. The gift was not propaganda. It was a true story, told about the wrong people, and the people telling it were handed the story by the people it was originally true of.

• • •

The price is not merely symbolic, and it should be calculated as concretely as the other costs in this book. The foreign-agent designation, wherever it is applied, imposes three costs on the civil society organization it reaches.

The first is operational: the compliance burden of registration, labelling, and reporting consumes organizational capacity that was previously devoted to program work, and in several of the jurisdictions above this burden appears to have been deliberately calibrated to be disabling rather than merely burdensome. The second is reputational: the label travels. A group carrying a foreign-agent designation in its home country loses its ability to represent itself as indigenous to an international audience, is easier to dismiss by domestic officials, and is less useful to foreign partners who do not wish to compound the problem. The third and most durable is the chilling effect on the organizations that were not designated, whose members understand that designation is possible and adjust their behaviour accordingly. That third cost is almost wholly invisible and almost certainly larger than the first two.

Set against this: none of the jurisdictions in the table above has yet successfully used a foreign-agent law to suppress a movement with genuine mass participation. The instrument's primary victims have been professional civil society organizations — advocacy groups, monitoring bodies, legal aid providers, environmental groups — rather than mass movements, which tend to lack the formal structure the law is designed to reach. Whether this represents a structural limit of the instrument or a staging of its use is not yet clear, and the answer has implications for the doctrine in Part Five.

CHAPTER NINE

Epistemic Rot

In August 1978 an analytical product on Iran circulated within the Central Intelligence Agency containing a sentence that has followed the institution ever since: that Iran was not in a revolutionary or even a pre-revolutionary situation. A congressional staff evaluation quoted it the following January; a president quoted it in his memoirs. The Director of Central Intelligence later disputed its status, insisting it came from a draft study he had rejected rather than from a finished assessment — a correction worth recording, and one that improves the institution's position very little, because a Defense Intelligence Agency appraisal the following month judged that the Shah was likely to remain actively in power for another decade, and that one was not a draft. The two documents were produced by different agencies with different sources and different analytical cultures, and they agreed. That is the fact requiring explanation.

Neither judgment was a rogue product or a careless one. Both were considered outputs of an analytical system built over decades on collection that was, by the standard of the time, reasonably good. The analysts were not unintelligent, and nobody has established political direction behind the findings. They were wrong in a way the system could not detect, because the wrongness was not a malfunction of the system. It was the system, operating as designed.

By September 1978, tens of millions of people in Iran had concluded privately that the Shah's government was ending. They had not all concluded it on the same day, and they had not announced it, because announcement was dangerous. But the revolution was, in the terms this book has been building toward, already present. It was just present in a quantity the analytical apparatus was not calibrated to see.

• • •

Begin with what the apparatus could see, because it could see a great deal.

It could see the loyalty of named commanders. It could see the internal politics of the officer corps, the factionalism within SAVAK, the composition of the cabinet, the Shah's health, the trajectory of oil revenues, and the distribution of military hardware. It had sources close to the palace, sources in the ministries, sources in the bazaar, and a long record of reporting from one of the most surveillance-saturated states in the Middle East. The technical collection was extensive. The human intelligence was genuine, not manufactured. By the standards of the profession, the station in Tehran was not badly run.

What the apparatus could not see was the content of conversations that were not reported to anyone. It could not see what a bazaari merchant believed his neighbors were about to do, or what an army conscript thought would happen to him if he refused an order in front of a crowd twice his size, or whether a mid-level cleric had begun to hear things in the streets that he had not heard in the mosques six months before. These are not exotic or inaccessible quantities. They are the quantities on which the actual outcome depended. But there were no sources reporting them, no collection priorities requesting them, and no analytical framework in which they would have constituted evidence of anything, because the framework was organized around who held office and whether they were in a position to continue holding it.

That is not a gap in coverage. It is a gap in ontology. The apparatus had a theory, implicit but operative, of what political outcomes consist of: they consist of who holds office and who has the coercive capacity to keep holding it. On that theory, the Shah's situation in September 1978 was defensible. He had the army. He had SAVAK. He had the loyalty, or at least the compliance, of the people who mattered under the theory. The theory told the apparatus to look for the lever, and the apparatus looked for the lever and found it still in the Shah's hands, and reported accordingly.

What the theory could not process was the possibility that the compliance of those same people was about to stop being available — not because anyone had decided to stop it, but because each of them was separately revising their estimate of what everyone else was about to do, in private, invisibly, and at speed.

• • •

Havana in late 1958 is the earlier case and in some respects the cleaner one, because the structural problem is more legible when the underlying capabilities are more limited.

Fulgencio Batista had, on paper, considerable advantages: a larger army than the insurgency, air power, control of the capital, and the active support of the United States until very late in the process. What he did not have, by late 1958, was an army that wanted to fight. This was not a secret, except from the people whose job was to know it. The rural conscripts who made up the bulk of the force had no attachment to the regime, came from the communities the insurgency was drawing on, and were receiving a continuous signal that they were on the wrong side of something large. Desertion and surrender were running at rates that, to anyone observing carefully, indicated not a military problem but a political one: the army was not being defeated, it was declining to fight.

The American assessment in late 1958 continued to describe a situation that might yet be stabilized, long after this had ceased to be true. The assessment was honest; the analysts were not told to shade it in any particular direction. They were simply not asking the question to which the relevant answer existed, which was not "can Batista hold the army" but "does the army believe it has anything to hold." The second question was not in the collection brief. It was not in the analytical framework. When Batista left on 1 January 1959, the departure was reported as a surprise.

• • •

Kabul in 2021 is the most recent and the most expensive iteration of the same failure, and it is worth being precise about what was and was not a failure, because the usual account attributes it to a set of causes — the speed of the American withdrawal, the corruption of the Afghan government, the quality of the Afghan army — that are all real and all beside the point of this chapter.

The Afghan National Defense and Security Forces as constituted in 2021 were, by any conventional metric, adequate to resist the Taliban without American combat support. They numbered over three hundred thousand. They had air power, artillery, and equipment the Taliban could not match. The assessment as late as June 2021 gave the government months; more optimistic assessments gave it years. Kabul fell in eleven days.

It fell because a very large number of soldiers and police decided, in the space of a fortnight, that they were not going to fight, and that decision propagated through the force at a speed that no analysis had modelled because no analysis had been tracking the variable that drove it. The variable was not morale in the ordinary sense, and not pay, and not equipment. It was the estimate each soldier was making of what the soldiers around him were about to do — an estimate that was revised, massively and simultaneously, the moment the first provincial capitals fell without the expected resistance. One surrender produced the conditions for the next, which produced the conditions for the next, and a force that had given every appearance of stability discovered in eleven days that it had no interior at all.

An assessment apparatus that could see personnel numbers, equipment inventories, and battle-damage assessments could not see this coming, not because it lacked information, but because the quantity that cascaded — mutual expectation about collective behaviour — was not a quantity it had ever been asked to track. The apparatus had, by 2021, sixty-eight years of experience measuring who held the lever. It had no experience measuring whether anyone would pull it.

• • •

Three cases, separated by sixty-three years, in different regions, at different scales, with different adversaries. The uniform element is not bad intelligence or political pressure on the analysts or any of the other explanations that have been offered. The uniform element is a theory of political stability organized around coercive capacity and the positions of named men, applied to situations that were changing on a different variable entirely.

The same apparatus that could reliably identify who could be bought could not identify when nobody could be sold. This is not a harder version of the same problem. It is a different problem, requiring a different instrument, and the instrument was never built because the institutions doing the work had spent the better part of a century building the other one.

That is what Chapter Seven called the formation of an ontology, and what this chapter is calling rot. It is rot in the specific sense: a corruption that works inward from the surface, leaving the structure apparently intact until it is not. The assessments look like assessments. The sources look like sources. The analytical products reach the principals on schedule and in the familiar format. The product is simply, and at the worst possible moments, not about the thing that is happening.

The reform proposed in Part Five is not a better collection programme or a reorganized analytical division. It is a different theory of what political stability is, built into the instrument from the beginning rather than retrofitted after the next collapse.

CHAPTER TEN

The Counterfactual Ledger

This chapter attempts something that standard accounting does not permit: to enter on the ledger not merely what was spent and what was received, but what the same resources would have purchased if directed elsewhere. It is an inherently speculative exercise, and the argument of this book does not depend on any particular number being right. It depends only on the approximate shape of the comparison, which is not speculative at all.

Democratic states hold a genuine and almost entirely unexploited structural advantage in the domain described in Part One. The advantage is this: the mechanism by which legitimacy conditioning works requires operating in the open, and an authoritarian state that operates in the open in the relevant sense is not doing authoritarianism. It cannot build training programmes for civic organizers, cannot fund civil society without being the civil society, cannot publish methods of withdrawing obedience without implying that obedience can be withdrawn from it. Democratic states can do all of these things, because the practices are continuous with their own domestic civic life. The space is genuinely theirs, and they largely vacated it.

• • •

What was not built is knowable in outline, because some of it was proposed and some of it was started and stopped.

The assessment capability described in Chapters Eight and Nine was not built. An instrument calibrated to the distribution of loyalty across a society's pillars, sensitive to the distinction between compliance and expectation, capable of returning a null finding in the absence of latent conditions — this was not built, and there is no evidence it was seriously attempted. The post-mortems following 1979 produced process reforms and collection enhancements aimed at the existing instrument; they did not produce a reckoning with the instrument's theory. The post-mortem following 2021 was still in progress as this was written, and early indications suggest it will follow the same pattern.

The standing capacity to transfer civic organizing technique openly and without operational strings was not built. What was built instead were organizations that gave with one hand the things the civic lineage required — training, connectivity, documentation support, legal assistance — and with the other hand the thing it cannot survive, which is the appearance of foreign management. Organizations like the National Endowment for Democracy were not useless and their work has been genuinely valuable in places. They were also, structurally, the maximum concession to the covert lineage's institutional logic that the democratic alternative could survive: grants decided in Washington, by Americans, on criteria American legislators funded, with reporting requirements oriented toward American principals. The Chapter Eight table records the cost of that architecture. It is not the organizations' fault.

The standing to speak about legitimacy was spent. This is the least tangible item on the ledger and the most important. Democratic states have an audience for arguments about political authority that autocracies lack: publics that have internalized the claim that government derives its authority from the governed, and that are accordingly attentive when that claim is invoked about somewhere else. That audience constitutes a form of power — soft power is the usual term, though it understates the mechanism — and it is a form of power that authoritarian states structurally cannot accumulate, because accumulating it requires practicing the thing it is based on. Every operation described in Part Two spent a fraction of that standing. The price was charged across sixty years and none of it appeared on a single invoice.

• • •

What the ledger cannot recover is the option that was traded away at the beginning. In the mid-1950s it would still have been possible to build a doctrine around the second lineage: to invest in assessment capability calibrated to legitimacy, to build the civic transfer capacity openly and with published methodology, and to treat the covert option as unavailable not because of restraint but because the better option existed and the worse one was visibly degrading the better one's foundation. That option closed, not in any single moment but through a series of budget decisions, capability investments, and institutional preferences that each made the next one harder to reverse.

The usual objection is that this is too easy in retrospect and too demanding of people who were facing real pressures in real time. The objection is largely right about the early period and progressively less right about the later one. By 1973 the evidence was extensive, the failure modes were documented, and the alternative doctrine was not speculative. The option had narrowed to a choice between a forty-year-old practice with a long record of instability and a twenty-year-old literature with a growing record of success, and the forty-year practice continued to be chosen because the accounting could see it and the literature produced nothing that could be shown to a principal at the end of a fiscal year.

That is the counterfactual: not a different world, but the one that was available. The one that was chosen cost more, produced worse outcomes, and left the institutions making the choice less capable of perceiving the world they had helped to make.

CHAPTER ELEVEN

The Quantity Nobody Can See

On 17 December 2010, a municipal inspector in the Tunisian town of Sidi Bouzid confiscated a produce cart belonging to a man named Mohamed Bouazizi. The inspector had done this before; Bouazizi had been selling without a permit since he was a child, and the confiscation of his scales and goods was the standard instrument of an ordinary petty shakedown. Bouazizi could not pay the fine. He went to the regional headquarters to complain and was refused an audience. He returned to the street, bought a can of paint thinner, and set himself on fire in front of the building.

Within four weeks, the government of Tunisia — which had governed continuously for twenty-three years and had been assessed as stable as recently as the previous autumn — collapsed. Within the year, the governments of Egypt, Libya, and Yemen had also fallen, and the Syrian civil war had begun. Uprisings of varying seriousness touched more than a dozen additional countries.

What happened? The standard accounts reach for proximity: the images spread on social media, governments were slow to respond, the repression was visible, and the regions had accumulated grievances over decades of unemployment and corruption and political exclusion. All of this is true and none of it is an explanation, because the grievances were also true in 2009, and in 2005, and in 1999. The fuel had been there. Bouazizi did not add fuel. He lit something.

What he lit was a problem that this chapter exists to name: the gap between what people privately believe and what they publicly express, and the catastrophic speed at which that gap can close when the right signal arrives.

• • •

In a political system where expressing dissent is costly — where it risks arrest, social punishment, loss of employment, or simply the identification of yourself as someone who opposes an arrangement that might persist — private belief and public expression diverge systematically. People who privately believe the government is illegitimate do not say so. They attend the mandatory rallies, they vote the right way, they are polite to the policeman, and they return home and continue to believe what they believe. This process is recursive: each person observing the public compliance of his neighbors draws the rational conclusion that he is probably alone in his private view, and adjusts his public behavior to match the compliance of people who are themselves only complying in the belief that he and everyone else genuinely approves. The result is that a society can contain an overwhelming majority of private dissidents while producing virtually no public signal of dissent, and each of those private dissidents correctly believes himself to be in a small minority.

What is missing is not opinion. What is missing is common knowledge of opinion — each person not only knowing their own view but knowing that others know it, and knowing that others know that they know. This is a specific informational condition and it can exist or not exist independently of whether the underlying opinion has changed. A society in which private dissent is total and common knowledge is absent looks, from outside, identical to a society in which support is genuine. No external observer, and no assessment apparatus of the conventional kind, can tell them apart.

What the Bouazizi moment did was supply a common-knowledge event. Not a particularly informative one: his death communicated nothing about the distribution of opinion that the people watching did not already privately know. But it communicated it publicly, which meant that everyone watching could see that everyone else was watching, and could update their estimate of what the person next to them was privately thinking, and could begin to consider a revision of their own public behaviour. The cascade that followed was not caused by his death. His death supplied the informational trigger for a cascade that was already latent in the distribution of private belief, and had been latent for years.

• • •

The implication for assessment is precise and unflattering.

A change in the quantity that matters — private belief about the arrangement's sustainability — is invisible by construction, because the condition of its invisibility is the same condition that makes the belief dangerous to express. An apparatus that tracks expressed sentiment will see only the public surface: the stable compliance, the attended rallies, the right votes. It will not see the private distribution, because the private distribution is not expressed. And it will not see the common-knowledge threshold approaching, because the threshold is crossed by events carrying no predictable size, and the approach to it is internal to each person's estimate of what others are thinking — which is also not expressed.

The threshold's crossing looks, from outside, like a sudden event. It is not sudden. The change that was sudden was the conversion of private knowledge into common knowledge. The underlying change may have been building for a decade. Every assessment filed in the decade before it was technically accurate about the expressed quantities it was measuring, and entirely uninformative about the quantity on which the outcome depended.

This is the structure that an adequate assessment method would need to represent, and the following chapters are an attempt to specify what representing it would require.

CHAPTER TWELVE

Approval, Compliance, and the Only One That Matters

At eleven-thirty on the morning of 21 November 1989, a man named Bogdan stood in the Great Hall of the Palace of the Republic in Bucharest and applauded. The applause was long, rhythmic, and in unison with the several thousand delegates around him. He had been applauding versions of this speech, with variations in duration and particular emphasis, for most of his adult life. He did not believe a word of it. He would have been astonished to learn that any of the other delegates did either. He applauded because the applause was required, because failure to applaud was observable, and because there was no world available to him in which the cost of not applauding was worth paying. He went home that evening and told his wife that things would not change in his lifetime.

He was describing his own expectation, not his own desire, and the distinction is the subject of this chapter.

• • •

Three quantities, routinely conflated, are doing three separate things in that scene and must be separated if what followed is to be understood.

Approval is what Bogdan did not have. He did not like the arrangement. He had not liked it for twenty years. His assessment of the regime's performance, its leaders, its premises, and its prospects was negative in almost every dimension he considered privately. Approval is the quantity surveys attempt to measure, that propaganda attempts to manufacture, and that is probably the least politically consequential of the three. A regime can survive universal private disapproval indefinitely; North Korea has done so for three generations. Approval does not run the machine.

Compliance is what Bogdan gave. He showed up, he applauded, he voted the right way, he filed the right reports, he did not say in public what he said at home. Compliance is what the machine runs on, and it is formally separable from approval because the cost of non-compliance is enforced by the apparatus regardless of anyone's approval of anything. Compliance can be purchased by sufficiently reliable enforcement indefinitely — which is why the question of when enforcement becomes unreliable is the operative political question, and why that question requires knowing something about the apparatus itself rather than about the population it is enforcing against.

Expectation is the third quantity and the only load-bearing one. Bogdan expected the arrangement to persist because he expected the others to continue complying, because he expected the apparatus to continue enforcing, because he could see no signal that any of these expectations was about to change. His compliance was, in the last analysis, a function of his expectation about collective behaviour rather than about his individual preference. This is the operative variable because it is the one that can cascade.

• • •

Three weeks after the congress, a different scene: a factory director in a town in Oltenia, whose name is not recorded, continuing on 10 December to prepare the monthly production figures he would submit to the ministry by the fifteenth. He did not approve of the regime any more than Bogdan did. He was complying, in the sense that he was doing the work the state required. But something had begun to shift in his expectation: he had heard about what was happening in Timișoara, he knew what the army had been asked to do there, and he had begun, without fully acknowledging it to himself, to hedge. The hedging was small: certain documents not filed quite on time, certain decisions deferred to superiors who would have to own them. He continued to comply. His expectation had already begun to move.

This is the intermediate state, and it is the state that an adequate assessment apparatus would most need to be able to detect, because it is the state in which intervention of the Chapter Nineteen kind might actually change something. By the time the cascade has begun, the instrument arrives too late to matter. Before the cascade is latent, the instrument returns nothing. The period in which expectation is shifting but has not yet tipped is the only window that has any operational relevance, and it is the window that no existing institutional instrument is calibrated to see.

• • •

On 22 December 1989, a general named Victor Stănculescu was walking around Bucharest with his leg in a plaster cast. He had arranged the cast himself. By his own later account it existed for one purpose: to make him unavailable for the order he expected to receive, which was to turn the army's guns on the crowds. That morning the Minister of National Defence, who had been given that order and had not executed it decisively, was found fatally wounded in the ministry; the death was recorded as suicide. Stănculescu was made acting minister within hours. Before the day ended he had ordered units back to barracks, and he had arranged the helicopter that took Nicolae and Elena Ceaușescu off the roof of the Central Committee building — an act he later described as saving them from a mob, and which delivered them to the custody that killed them three days later, at a trial he organised.

The cast is the document. It was prepared in advance of the order it was designed to evade, which means the decision it embodies was made before 22 December — made as information arrived about what the army had been asked to do in Timișoara, about how those orders had actually been carried out, and above all about what the other generals appeared to be doing, they being the people whose behaviour determined whether compliance or defection was survivable. Stănculescu is not offered here as a figure of conscience; in 2008 a Romanian court convicted him over the Timișoara shootings and sentenced him to fifteen years. That is the point. Expectation is not a moral quantity. It moved in a man who had helped do the thing he was now declining to continue doing, and it moved before anything visible changed.

The three quantities in that sequence — Bogdan's approval, the factory director's hedging compliance, Stănculescu's revised expectation — are not three points on one scale. They are three different objects that move at different speeds, are driven by different inputs, and respond to different interventions. An instrument that measures one and reports on all three has committed an error more consequential than a gap in coverage: it has identified the wrong quantity as the operative one, and will produce confident findings that are systematically orthogonal to what is about to happen.

CHAPTER THIRTEEN

What an Adequate Method Must Represent

On the evening of 9 November 1989, a passport-control officer named Harald Jäger was on duty at the Bornholmer Strasse crossing point between East and West Berlin. He had not been briefed on anything unusual. His orders were the standing orders: check travel documents, apply the standard restrictions, process authorised crossings in the normal time. At approximately six-thirty in the evening, a press conference in East Berlin ended with a statement by a member of the Politburo that new travel regulations would allow citizens to cross the border. Asked when the regulations would take effect, he looked at his notes and said: immediately, without delay.

The statement was an error. The regulation had not been finalised and was not to take effect until the following day. No border commander had been briefed. Jäger's telephone calls to his superiors over the next four hours produced no new orders, because no new orders existed; what he got instead, at one point, was a senior officer on an open line calling him incapable of handling the situation. Meanwhile the crowd at the crossing grew from dozens to hundreds to thousands. He was instructed at one stage to let the loudest through and stamp their documents in a way that would prevent their return — a measure that required his officers to identify individuals, one at a time, in front of a crowd that had begun pushing the barrier aside. He had roughly fifty men.

At about half past eleven, Jäger ordered the gates opened. Some twenty thousand people crossed in the next three quarters of an hour, and by midnight the other crossings in Berlin had followed.

Consider what an adequate method would have needed to represent in order to understand why that decision went the way it did — and why it mattered as much as it did, because it was one of the decisions that ended the German Democratic Republic.

• • •

The requirements that follow are stated as a specification rather than as a description of any existing instrument, because no existing instrument meets them. Each is derived from the analysis of the preceding chapters, and each carries an implicit critique of what current assessment practice does instead.

The first requirement is a *distributed loyalty map* across the regime's institutional pillars: the security services and their internal divisions, the military with its distinct corps and their distinct relationships to the political leadership, the civil bureaucracy at the levels that process decisions rather than merely make them, the courts, the media, the church or comparable legitimating institution, the organized economy. The loyalty map is not a list of named officials and their assessed reliability. It is a distribution: what proportion of people in each pillar are complying from what motivation — fear, material benefit, genuine belief, inertia, or expectation about others — because each motivation responds differently to the events that shift the overall picture. Jäger's decision was not made by a man loyal to a regime; it was made by a man maintaining compliance in conditions where the cost of maintaining it had suddenly become unclear, and where his superiors had given him no instruction for those conditions.

The second requirement is *sequencing and conditionality across pillars*. Jäger did not act in isolation. His decision was conditioned on his estimate of what the crowds would do if he did not open the gates, what his own officers would do if he ordered them to push the crowd back, what his superiors were doing while they gave him no orders, and what was happening at the other crossing points. Each estimate was a function of what he believed the others were doing, and the network of mutual estimates is the operative object — not a set of individual loyalty scores, but an ordering with dependencies, in which one pillar's movement changes the calculus for the next. A method that scores pillars independently will not capture this, because the mechanism is precisely the conditionality.

The third requirement is a formal treatment of the *observability problem*. The absence of visible dissent at a given moment cannot be read as evidence of stability, because the model developed in Chapter Eleven shows that absence of visible dissent is the normal condition of a society whose private opinion is entirely hostile, as long as enforcement is believed to be reliable and common knowledge is absent. Jäger's crowd had not been publicly dissident a week before. The absence of visible dissent in East Germany in October 1989 was not evidence of a stable state. It was evidence of a state in which the cost of expression exceeded the benefit until the Leipzig demonstrations began — and the Leipzig demonstrations were themselves a common-knowledge event that revised the calculation everywhere they were seen.

The fourth requirement is the *distinction between grievance and expectation*, which this chapter has already made but which bears formal statement as a methodological rule: a method that measures dissatisfaction and reports instability will produce false positives continuously, because dissatisfaction is the chronic and stable condition of most of the populations under authoritarian governance, and it is not what predicts cascade. What predicts cascade is a shift in the specific, interpersonal, institutionally-structured expectation about what others are going to do — the quantity Bogdan had in stable equilibrium in November and Stănculescu had shifted entirely by December.

The fifth requirement is attention to *elite time horizons and exit options*. Whether the members of the regime's sustaining coalition can survive the regime materially — whether their assets, their families, their professional prospects, and their personal safety are separable from the continuation of the current arrangement — is frequently the operative variable in whether they defect, and at what cost. Stănculescu survived the end of the Romanian regime and even its trial. Qaddafi's inner circle could not imagine surviving his fall, which is part of why his fall was so much more violent. An adequate method would track this as a structural feature of the coalition rather than as a biographical detail about its members.

The sixth requirement is sensitivity to *common-knowledge events and their information content*. The statement at the Bornholmer Strasse press conference that evening had no legal force, because the regulation it described did not exist. It was one of the most consequential political events in postwar Europe, because it supplied the common-knowledge trigger that converted a private expectation already held by millions into an expressed and publicly visible act. The mechanism by which this happened — a statement interpreted as an authoritative signal, broadcast simultaneously to an audience that knew others were watching — is not an exotic or unpredictable mechanism. It has a structure that can be characterized in advance. A method that tracked the conditions under which such events are possible, without being able to predict the event itself, would be substantially more useful than one that waits for events large enough to override its null state.

One additional requirement deserves its own statement because it is a constraint rather than a capability. The method must have a *resolution floor*: it must be structurally incapable of resolving below the level of institutional aggregates and pillar distributions. A method that can track individual defection propensity by name is a targeting system regardless of its stated purpose, and the distinction between an assessment instrument and a domestic control apparatus operating at that resolution is a matter of intent rather than of design. Intent changes. Design does not. The resolution floor is therefore a design requirement, not a usage policy, and Chapter Sixteen will explain why this matters as much as it does.

CHAPTER FOURTEEN

What the Method Must Be Able to Fail At

Everything in the preceding three chapters has been building toward a requirement that is not a capability. It is an absence — a specific, well-defined, structurally mandated absence — and it is the most important thing this book argues.

An adequate method of assessing political legitimacy must be capable of returning nothing.

This needs to be stated exactly. The requirement is not that the method sometimes fails, in the ordinary sense of producing an incorrect result. Every method fails sometimes. The requirement is that the method has a *well-defined null state*: a state in which the legitimate output is the finding that the latent conditions for significant political change are absent, that no available external action would materially alter the trajectory of the situation, and that the correct response is to do nothing. And the requirement is that the method's accuracy at identifying that state is its primary validation criterion — not a secondary check on a method primarily validated by its positive findings, but the first test it must pass before anything else about it is worth examining.

• • •

Explain the inversion by considering what it excludes.

A method that consistently finds a situation worth acting on in every case it examines is not measuring legitimacy. It is measuring opportunity — the presence of a grievance, a coalition, a commercial interest, an ideological opponent, a strategic position worth contesting. There is always a grievance. There is almost always a coalition. There is usually a commercial interest. A method calibrated to find these things will find them everywhere it looks, will generate an unbroken series of positive findings, and will never produce a null result, because the quantities it is sensitive to are chronic and universal.

This is a precise description of the apparatus described in Part Two. It produced positive findings. The positive findings were acted on. The actions produced the results recorded in Chapters Four through Seven. The apparatus was not malfunctioning. It was measuring the quantities it was calibrated to measure, and those quantities are not the same as the quantities that determine whether external action produces a durable outcome or an expensive failure. Calibration to opportunity rather than legitimacy generates systematic overproduction of interventions, not by any individual's corruption or impatience, but as a mechanical consequence of what the instrument is sensitive to.

The null-state requirement is therefore not a moral commitment that a doctrine imposes on an institution. It is a calibration standard that an instrument must meet before it can be trusted to generate the positive findings that justify anything. And it converts the book's ethical position — that external action premised on false legitimacy readings is destructive — from a principle into a property. A method that is well-calibrated at its null state will not authorize action in the absence of latent conditions, because it will not see latent conditions that are not there. The restraint is in the instrument. It does not depend on the virtue of the person holding it.

• • •

The asymmetry argument of Chapter Eighteen rests on this. The claim that legitimacy conditioning cannot be imposed — that it amplifies an existing latent cascade or does nothing at all — is not merely an ethical observation about what external actors should do. It is, if the null-state requirement is properly built, an empirical property of the method itself. An instrument calibrated to detect latent conditions will return null in their absence, and a doctrine built around that instrument will accordingly recommend no action, because the instrument's output is the basis for the doctrine's recommendations. The constraint is not volitional. It is structural.

This is the hinge of the whole argument, and it should be stated in the form that makes it most falsifiable. If a well-calibrated null-state instrument applied to a given situation returns a positive finding — if it identifies a latent distribution already in motion — then external assistance of the Chapter Nineteen kind may amplify something real, and the asymmetry argument permits it. If it returns null, the asymmetry argument requires that nothing be done. If it returns positive in every case, it is the Part Two instrument in better clothes, and the whole construction collapses.

What makes the claim falsifiable is the proposition that the null state is common. Most societies, most of the time, are not in a latent cascade toward regime change. Most authoritarian states are not on the verge of collapse; most populations, most of the time, expect the arrangement to continue and have adjusted their behaviour accordingly. A method that finds cascade-in-progress everywhere it looks is simply wrong. The frequency of the null finding is a diagnostic of the method's calibration, and a method that never produces it has failed the most important test regardless of how well it performs on the positive cases.

• • •

One vulnerability requires honest acknowledgment.

A method with a null state can be pressured to suppress it. This is not a hypothetical risk. The entire architecture of Part Two demonstrates that commercial and political interests can systematically corrupt the translation of assessment into recommendation, and there is no reason to believe that a well-designed instrument is immune to the application of pressure on the people interpreting its outputs. Chapter Twenty, which addresses the question of who holds the instrument and under what conditions, is this chapter's necessary sequel. What can be said here is that the null-state requirement at least makes suppression visible: a positive finding where a null finding is expected is an anomaly that can be tracked, documented, and interrogated in a way that the absence of a null-state requirement does not permit. An institution with no null state can authorize every intervention without producing any evidence of anomaly. An institution whose instrument has a null state and whose record shows that the null state is never returned has left a paper trail that the Chapter Twenty custody arrangements are designed to make available.

The case for the null-state requirement does not rest on the optimism that it will never be suppressed. It rests on the argument that suppression of it is at least a different and harder problem than the one that produced Part Two. That is a modest claim, and the modesty is deliberate.

CHAPTER FIFTEEN

The Autocrat's Daily Problem

In the summer of 2020, Alexander Lukashenko arrived at a polling station in Minsk and voted in a presidential election he was going to win with eighty per cent of the vote, by an official count that appeared to be counting something other than votes. The election's actual result was not published and is not known with precision. What is known is that within two days of the announced result, several hundred thousand people were in the streets of Minsk, that a general strike was beginning, and that the security services were detaining and reportedly beating protesters at a rate that overwhelmed the available processing capacity. Lukashenko appeared in public wearing a bulletproof vest and carrying an assault rifle.

The image is useful not as symbol but as mechanism: a ruler who needs to appear in public armed is a ruler who has stopped being certain about what the armed people around him will do. The rifle was not there to reassure the public. It was there because the people who would normally have ensured that no such thing was necessary — the security detail, the loyal officers, the dependable commanders — had become a somewhat less dependable quantity, and the rifle was the hedge.

Lukashenko survived 2020. The movement did not have the elite defections it needed, and Russia made clear that it would not allow a transition. But the rifle is the text this chapter is reading, because it is the external trace of the autocrat's internal problem, which this chapter inverts the analysis to examine.

• • •

An authoritarian ruler's core dilemma is that the apparatus required to enforce compliance is itself the most capable instrument of his removal, and he cannot monitor it from outside it. He must continuously re-purchase the loyalty of people who could, if they coordinated, terminate his tenure at any moment, using resources those same people are responsible for generating and accounting for. Every day of an authoritarian regime is a solution to this problem, and the solution is never permanent because the problem never resolves.

The purchase takes several forms, and each leaves a trace.

Promotion and appointment as loyalty signals: placing reliable people in positions of consequence and less reliable people in positions of visibility without consequence. The difference between these is more observable than it might appear: an institution in which the capable are systematically not in charge of anything that matters is an institution whose internal structure is organized around the control problem rather than the mission.

Redundant security organs whose primary function is to monitor each other. The proliferation of parallel intelligence services within a single state — each with overlapping jurisdiction, separate reporting chains, and mutual surveillance mandates — is not an accident of bureaucratic expansion. It is the structural solution to the problem of trusting any single service, and its cost is that coordination across the organs becomes unreliable, which means that the state's response to crises requiring cross-organ coordination degrades exactly when it is most needed.

Resource flows that must be maintained. A coalition purchased with access to rents must be kept in access to rents. This creates a rigidity in fiscal and economic policy that becomes increasingly visible under pressure: the budget items that cannot be cut, the contracts that cannot be cancelled, the exchange rates that cannot be adjusted, the reform programs that cannot be implemented because implementation would require displacing someone whose displacement would destabilize the arrangement. Every authoritarian state has this structure, and its shape, read carefully, is a map of the coalition.

Coalition narrowing under stress. When the resources available for purchase decline — during economic shock, during isolation, during the early stages of a legitimacy cascade — the autocrat faces a triage problem: which pillar members to keep and which to let go. The characteristic pattern is that the narrowing begins at the periphery and moves inward, shedding the cheaper and more replaceable members of the coalition while concentrating resources on the irreplaceable core. This process is observable in fiscal data, in appointment patterns, and in the progressive restriction of the coalition's visible membership. It is also self-defeating, because the shedding of peripheral members is itself a signal that the arrangement is under stress, which accelerates the revision of expectations across the entire remaining coalition.

• • •

These four patterns are not difficult to observe. They are observable in public data — budget allocations, appointment records, organizational charts, procurement patterns — and they are far more informative about the actual state of a regime's internal cohesion than the conventional assessment variables of political loyalty and coercive capacity. The problem is not that the signal is faint. The problem is that the instrument is not listening for it, because the instrument was built to answer a different question.

An apparatus calibrated to the Chapter Thirteen requirements would have substantially more to work with in most authoritarian states than the existing assessment practice does, because regimes emit far more legitimacy-relevant signal than they emit control-relevant signal. The coal mine that tells you about canaries is measuring the wrong gas.

The autocrat's daily solution to his daily problem generates, as a byproduct, a continuous and partially legible record of the problem's severity. Reading that record is the assessment task. Building the instrument that reads it is the reform.

CHAPTER SIXTEEN

The Mirror Problem

In the archives of the Stasi — the East German Ministry for State Security, dissolved in 1990 and whose records were subsequently opened — there is a file category called Operative Personenkontrolle, operative individual surveillance. A person placed in this category was subjected not primarily to monitoring but to a programme of active intervention designed to alter their behaviour without their knowledge. The programme had a name: Zersetzung, decomposition. Its methods included rearranging furniture in the subject's apartment, replacing objects with slightly different ones, making anonymous calls at night, sending forged letters in the subject's name to their employer or their partner, causing small malfunctions in their workplace equipment, and organizing the steady withdrawal of their social circle through approaches to the people around them.

The objective was not to gather information. The objective was to produce a specific change in the subject's estimate of their own situation and of what the people around them were doing — to make them uncertain of their own perceptions, isolated, and unable to coordinate with others. The Stasi was, in these operations, acting directly on the variables described in Chapters Eleven and Twelve: the subject's sense of whether common knowledge of their situation existed, whether others shared their views, and whether collective action was possible. It was engineering the informational conditions that determine whether a latent cascade can form.

This is the mirror problem. The instrument specified in Chapter Thirteen, built to detect and characterize the conditions for a legitimacy cascade, is also an instrument for preventing one. The technique for identifying the approach of common knowledge is the technique for interrupting it. The requirements document is reversible.

• • •

The objection should be stated at full strength, because weakening it would be a way of evading it.

A programme that maps the relationships within a movement in order to understand how coordination occurs is also a programme that maps the relationships in order to identify where severing them would be most effective. A method that tracks shifts in private belief and their relationship to common-knowledge thresholds is a method that, held by different hands, identifies the communication channels and the focal events that most efficiently deliver those thresholds — for the purpose of closing the channels before the events can occur. An instrument calibrated to the distribution of loyalty across an elite, sensitive to the sequences by which pillar defections propagate, is an instrument that, used differently, reveals where early defectors can be isolated before the cascade reaches the others.

The numbers are worth stating precisely, because they are usually stated loosely. At the end, the Ministry employed something over ninety thousand full-time staff — roughly one for every one hundred and sixty-six citizens — and carried on its books some hundred and seventy thousand registered informants, which yields the frequently quoted ratio of one officer or informant for every sixty-three East Germans. Higher-end estimates counting occasional and part-time collaborators run to one per six or seven, and should be treated as the upper bound of a contested count rather than as a finding. However it is counted, it was probably the most densely surveilled society in history, and it did not prevent the collapse of the state it was protecting.

The Zersetzung files are not evidence that the instrument is useless in the direction of control. They are evidence that it was being applied at a resolution that was politically unsustainable and strategically inefficient — enormous effort spent on individuals, one at a time, while the quantity that actually ended the state moved somewhere the files did not reach. A more carefully calibrated version, acting on institutional aggregates and communication networks rather than on persons, might have been considerably more effective. The Stasi's failure was not a failure of the concept. It was a failure of scale, and scale is exactly what has since become cheap.

That last clause is the part of this chapter that has aged badly since its historical material was set, and it needs saying plainly. Everything the Ministry accomplished with a hundred and seventy thousand informants — mapping who knows whom, identifying which relationships carry coordination, detecting the formation of shared belief before it becomes public — is now a byproduct of ordinary platform operation, obtainable from data that users generate voluntarily, at a cost per subject several orders of magnitude below a case officer's salary. The common-knowledge events described in Chapter Eleven now propagate through recommendation systems whose ranking decisions determine what a population learns that other members of the population are thinking; and the same systems, adjusted, determine what it does not learn. A state that can shape distribution can act directly on the variable that Chapter Twelve identified as load-bearing, without arresting anyone, without entering an apartment, and without producing a file that anyone can later open. The mirror problem is not a historical analogy. It is a description of infrastructure that already exists, mostly in private hands, in every country including the ones this book is addressed to.

This is the uncomfortable finding, and this book will not soften it. The requirements in Chapter Thirteen do not by themselves prevent the instrument from being used for control. They specify what an instrument for assessing legitimacy would need to represent, and a well-resourced state actor could build an instrument meeting those requirements and use it to suppress what it has identified. The specification is not proprietary. The physics is not directional.

• • •

Two structural responses are available. Neither is fully adequate, and only one is adequate enough.

The first is the resolution floor imposed in Chapter Thirteen's final requirement. An instrument that cannot resolve below the level of institutional aggregates cannot be directly used for Zersetzung, because Zersetzung operates on individuals. This is a genuine constraint rather than a usage policy, because it is built into the instrument's design — the floor is what the instrument cannot do, not what it is not supposed to do. It eliminates the most operationally damaging application of the mirror without eliminating the assessment capability.

The second is the null-state requirement of Chapter Fourteen. An instrument calibrated primarily to return null is an instrument whose primary output, in most situations, is that nothing is happening and nothing should be done. This output is essentially useless to a control apparatus, which does not need to be told that nothing needs suppressing — it only needs the positive findings, the identification of conditions where a cascade might form. A method whose calibration standard is accuracy at null is a method that produces a great many nulls and a smaller number of carefully validated positives, and the ratio is unfavourable from the perspective of a control application that needs reliable positive identification of threats to be actionable. The null-state instrument is still dual-use, but it is more expensive to use in the control direction, because the same calibration that makes it trustworthy in the assessment direction makes it persistently reluctant to flag the situations a control application needs it to flag.

These two constraints together — the resolution floor and the null-state calibration — do not eliminate the dual-use risk. They convert it from a risk of identical misuse to a risk of modified, more expensive, and more observable misuse. The remaining vulnerability is addressed in Chapter Twenty, not here. What can be said here is that the objection this chapter has raised does not invalidate the project. It requires that the project be designed specifically against the objection, and the Chapter Thirteen requirements have been written with that requirement in mind. A book that does not raise this chapter, or that raises it and answers it too easily, is a book that has not finished its argument.

CHAPTER SEVENTEEN

Built to Win, Unbuilt to Govern

On 11 February 2011, eighteen days after the occupation of Tahrir Square began, Hosni Mubarak resigned the Egyptian presidency and transferred authority to the Supreme Council of the Armed Forces. The movement that had forced this outcome celebrated in the square that night. It had achieved something genuinely remarkable: the removal, through nonviolent mass action, of a ruler who had governed for thirty years and who commanded one of the largest militaries in the Arab world. The celebration was warranted.

What the movement did not have, on 12 February, was any mechanism for converting its achievement into governance. It had no hierarchy. It had no leadership with recognized authority to negotiate on anyone's behalf. It had no organizational memory beyond the preceding eighteen days. It had no relationship with the institutions of the state that it had not just occupied in protest. And it had no mechanism — by design, by necessity, and by the structural logic of how it had won — for making binding decisions about what came next.

The Supreme Council of the Armed Forces had all of these things. It had hierarchy, institutional memory, relationships with every part of the state, the capacity to make binding decisions, and the monopoly on coercive capacity that Mubarak had just vacated. The movement had made the transition possible. The army governed it.

• • •

The tension exhibited in Tahrir Square is not specific to Egypt, and it is not a failure of the Egyptian movement. It is a structural feature of the mechanism described in this book, and it follows directly from the properties that make the mechanism work.

A distributed movement with no fixed leadership is decapitation-resistant: there is nobody to arrest whose arrest stops it. The absence of hierarchy that protects the movement in its confrontation with the state is precisely the absence of hierarchy that leaves it without a governing structure the morning after the state yields. A movement that draws its strength from the breadth of its participation — shopkeepers, students, professionals, the disaffected from every part of the social map — does not thereby acquire a coalition capable of governing, because governing requires making allocative decisions that will disappoint some of those participants, and the movement has no mechanism for managing that disappointment that does not dissolve it. A campaign optimized for the withdrawal of compliance is not optimized for the construction of a new basis for compliance. These are different problems, and the first problem's solution tends to destroy the tools required for the second.

This tension has a name in the civil-resistance literature — the transition problem, or the post-victory gap — and it is the least developed area in the tradition. The theoretical literature on how to force a transition is substantial. The literature on what to do in the eighteen months after forcing one is thin. The thinness is not an accident. The tradition developed in opposition, in conditions where the question of governance was deferred to after a victory that had not yet arrived, and the deferral has never fully been redeemed.

• • •

The problem compounds because the residual state apparatus does not share the transition problem. It was there before the movement, it is there after the movement, and it has the institutional memory, the personnel, the coercive capacity, and the bureaucratic relationships that the movement cannot develop in the span of weeks or months during which its energy is highest. Ukraine after 2004, Egypt after 2011, Sudan after 2019: the pattern is consistent. A movement successful enough to force a transition encounters, on the far side of the transition, a state apparatus that has survived the movement and is available to whoever can consolidate their hold on it. That may be the movement's leadership, if it has one; it may be the military; it may be the residual coalition of the previous regime; it may be an outside actor with sufficient resources and relationships. What it will not be, without institutional capacity that the movement does not have, is the movement itself.

The doctrine in Part Five offers assistance at this stage — documentation, legal support, convening, transition planning — and is honest that this assistance is the most difficult to provide on the terms the refusability principle requires. Consolidation assistance creates the conditions for the relationship between an external actor and a new government to resemble the relationship described in Part Two, and the people on all sides of the consolidation period are aware of this and behave accordingly. No clean solution to this exists. The chapter ends here, without one, because ending it with a solution that is not available would be the same kind of assessment error this book has spent its first half diagnosing: the production of a positive finding in the absence of a positive condition.

The problem is stated as an honest limit. What the movement needs after the transition — the institutional capacity to convert distributed withdrawal of compliance into a durable new basis for compliance — is not something any external actor can supply, for the same reason that the withdrawal itself could not be externally manufactured. The authority to build a new governing order, like the authority to dismantle the old one, derives from the people who will live under it. That is not a slogan. It is the structural claim this book has been making since Chapter Three, and it holds as firmly after the victory as before it.

CHAPTER EIGHTEEN

The Asymmetry Argument

Between 2009 and 2012 the United States Agency for International Development built a social media platform for Cuba. It was called ZunZuneo, after the sound a Cuban hummingbird makes. It ran through front companies in Spain and the Cayman Islands so that its origin could not be traced; it was seeded with mobile numbers obtained from the Cuban state telecom without their owners' knowledge; and it was designed, according to the documents the Associated Press later obtained, to proceed in stages — build an audience on football, baseball, and weather, introduce political content once the audience was large enough, and then, at maturity, mobilise subscribers into what the planning documents called smart mobs, in order to renegotiate the balance of power between the state and society. Contractors profiled early respondents by political disposition. At its peak the platform had about forty thousand subscribers. By the internal metric — audience growth — it was performing.

It produced nothing. Not a diminished result, not a partial cascade suppressed by the Cuban state — nothing. When the project ran out of money and shut down, its subscribers noticed that their messaging service had stopped working. When it was exposed two years later, the principal consequence was a news cycle in Washington and a permanent exhibit in Havana's case that every civic irritation on the island is manufactured abroad.

The project is usually told as a story about incompetence, and there was incompetence in it. But it is more instructive read as a controlled experiment that no one intended to run. Here was the full covert repertoire — deniable funding, staged influence, engineered mobilization — applied patiently, over years, to a population with ample grievances against a regime with a long record of economic failure. The tools functioned. The audience assembled. And the thing the project existed to produce could not be produced, because the thing it existed to produce was never within the power of the people producing the project.

• • •

This chapter states the doctrine's core claim, which the whole book has been assembling: *actor substitution can be imposed on a society against its latent conditions, and legitimacy conditioning cannot.*

The first half of the claim is Part Two. A coup requires a dozen decisive defections, purchased privately, and a population's beliefs are irrelevant to whether it succeeds on the night — which is precisely why it can be scheduled, budgeted, and delivered, and precisely why what it delivers cannot stand on its own. The second half of the claim follows from the mechanism established in Part Four. Legitimacy conditioning operates on one variable only: the mutual expectations of a society's members about each other's behaviour. Those expectations are formed by domestic actors observing other domestic actors — the conscript watching his unit, the minister watching the other ministers, the shopkeeper watching the street. An external actor can inject information into that system. It cannot manufacture the observations, because the observations are of people it does not control, doing things it cannot make them do. Where the latent distribution is already in motion, external assistance can lower the cost of coordination and widen the sense of the possible. Where it is not, there is nothing to lower the cost of. The toolkit amplifies. It does not generate.

ZunZuneo is the negative demonstration. Serbia is the positive one, and the comparison is the argument in miniature — but it is only an argument if the machinery of Part Four is doing the explaining, so let the machinery do it.

Take the three quantities of Chapter Twelve. In Serbia by 2000, approval of the arrangement had been negative for years; compliance was still being given; and expectation had begun to move, through lost wars, hyperinflation, and a decade of private conclusions drawn by people watching other people. What the assistance supplied was technique for converting that private distribution into common knowledge — the parallel count of Chapter Twenty-One, the pillar work on the police described in Chapter Two — and what did the actual work on the night was the Chapter Eleven mechanism: several hundred thousand people in one place, each observing that the others had come, and a police cordon observing the same thing and drawing the same inference. Nobody abroad produced that inference. It was produced by Serbs looking at Serbs.

Now run ZunZuneo through the same apparatus. It could deliver messages, which is an input to approval and at the margin to expressed preference. It could not deliver the observation. A Cuban reading a satirical text about a former telecommunications minister learns something about what a text message says; he learns nothing whatever about what his neighbour, his foreman, or the officer at his local station is about to do — and Chapter Twelve established that only the third quantity is load-bearing. The platform was operating on the wrong variable, and no amount of subscriber growth converts one into the other. Worse, from the operator's side: the informational content of a common-knowledge event depends on its being mutually observed, and a covert platform is structurally incapable of supplying mutual observation, because its participants cannot know who else is participating or on whose behalf.

The inputs were roughly comparable across the two cases; the outcomes varied with a variable the inputs could not reach. That covariance is what evidence for the asymmetry looks like. And Chapter Fourteen's instrument, had it existed, would have returned a positive reading on Serbia and a null on Cuba before a dollar was spent on either.

• • •

The objection must now be given its full weight, because it is partly true and the doctrine cannot survive pretending otherwise.

Manufactured consent is a real phenomenon. States demonstrably shape expressed opinion through propaganda, and foreign actors demonstrably shift public discourse through disinformation — the record of the last decade of influence operations is not seriously disputable, and a book that waved it away would deserve the dismissal it would get. If expressed preference can be manufactured, why not expectation? If a foreign actor can make people say things, why can it not make people believe that others are about to act?

The answer is that the two quantities are produced by different machinery, and only one of them is reachable from outside. Expressed preference responds to the information environment, which external actors can pollute. But the expectation that matters — the colonel's estimate of what the other colonels will do — is not formed from the information environment. It is formed from direct observation of specific, known people inside specific institutions: how the orders were carried out last week, who hedged in the meeting, whether the units actually moved. That evidence is generated by the behaviour of domestic actors, it is consumed close to its source, and it is precisely the evidence that discourse-level operations cannot fabricate, because fabricating it would require controlling the behaviour being observed — at which point the external actor is no longer conditioning legitimacy but running a coup, and has crossed back into the first lineage with all of its properties.

The claim, properly bounded, is therefore this: external actors can influence what a population says, and can marginally influence what it believes about distant things. What they cannot manufacture is the mutual expectation among the members of an elite that their own arrangement is ending, because that expectation is built from observations only domestic behaviour can supply. It is a narrower claim than a general claim about propaganda, and every reduction in its breadth is an increase in its defensibility. The color-revolution accusation, in its strong form — that Western actors manufacture cascades in societies that would otherwise be stable — asserts exactly the capability the Cuban experiment tested and failed to find. The accusation's persistence is a Part Three phenomenon. Its truth would require machinery that no one, on the evidence, possesses.

What remains is the doctrine's first rule, stated as a property rather than a promise: assistance to legitimacy conditioning succeeds only where a domestic cascade is already latent, and therefore cannot, even in the hands of bad actors, produce transitions that domestic conditions do not support. It fails safe. The coup fails catastrophic. Everything else in this Part is the working-out of what may responsibly be built on that difference.

CHAPTER NINETEEN

The Refusability Principle

On 1 August 2009, in the first of a series of televised mass trials of more than a hundred Iranian journalists, politicians, academics, and clerics, a prosecutor read an indictment to the Revolutionary Court in Tehran. Among the charges was the assertion that the events following the June election had proceeded according to a timetable, and that more than a hundred of the hundred and ninety-eight steps set out in the writings of Gene Sharp for a velvet coup had already been carried out. Iranian intelligence had introduced the same elderly academic to the public a year or two earlier, in an animated film that showed a computer-generated version of him conspiring with an American senator and a Hungarian-American financier, and described him as one of the CIA officers responsible for American infiltration abroad. He was a retired researcher living in East Boston. The accusation was, in one narrow sense, well-researched: translated copies of his handbook had in fact circulated among Iranian activists, as they had circulated in a great many other places, in more than thirty languages.

The accusation nonetheless failed to grip, and the manner of its failing is the subject of this chapter. The academic's response, in substance, was that everything he had ever produced was published, freely downloadable, available in Farsi, and open to inspection by the Iranian government itself; that he directed no one, funded no one, and knew none of the people accused of following him; and that a method available equally to everyone belongs to no one. There was nothing to expose, because nothing was hidden. The regime was left accusing its citizens of having read a book.

Compare the platform that opened Chapter Eighteen, whose exposure was a durable gift to the government it targeted. The difference between the two cases is not scale or competence. It is a single design property, and the property can be stated as a test.

• • •

Any activity that must remain secret in order to function is actor substitution wearing better clothes.

The test is deliberately unforgiving, and it is worth walking through why the doctrine needs it at full strength rather than in a reasonable-sounding weaker form. The mechanism established in Part Four runs on common knowledge: a movement's power is precisely the public visibility of its participation. Assistance that operates in secret is therefore not a discreet version of the same help. It is help of a categorically different kind — help whose value depends on concealment, which means help that would injure the recipient if known, which means help the recipient cannot be presumed to have accepted, which means, followed to the end, an operation conducted through the recipient rather than an assistance rendered to them. The secrecy is not a modality. It is a tell.

Three requirements follow, and together they constitute the refusability principle.

Assistance must be *published*: the methodology, the curriculum, the tools, and the fact of the program itself, stated openly enough that the intended beneficiaries, their adversaries, and third parties can all see what is on offer. Assistance must be *transparently funded*: the recipient must know who is paying, and so must everyone else, because a recipient who does not know the source of their support has not consented to the relationship, whatever they have signed. And assistance must be *structurally declinable*: offered rather than delivered, on terms the recipient sets, revocable by the recipient without penalty, and designed so that declining it is a real option — which means, among other things, that it cannot be bundled with resources the recipient cannot survive refusing.

What this rules out should be enumerated, because the enumeration is the doctrine's teeth: covert funding of civic organizations, in any amount, through any number of intermediary layers. Unattributed messaging and front platforms of every kind. Deniable partnerships. Payments to individuals within movements. The purchase of editorial positions. Any relationship in which the recipient's public account of their own support would be false. Every one of these is standard practice somewhere in the present architecture of democracy assistance, every one of them has been defended as a protective measure for the recipients, and every one of them fails the test — because each converts a participant into an instrument, and because each, upon exposure, validates the inoculation script of Chapter Eight for every movement everywhere, including the movements that took nothing.

• • •

Now the hard objection, which is not hypothetical and must not be handled as if it were.

Transparency exposes recipients. The instruments tabulated in Chapter Eight are triggered by disclosed foreign funding; a movement that accepts published assistance from a democratic government has handed the designation regime its evidence, and people have gone to prison on exactly this. The covert channels the refusability principle prohibits were built, in part, by people who understood this and were trying to protect their partners. The objection deserves an answer that does not minimize the cost.

The answer has three parts, in ascending order of weight. First, the covert alternative does not avoid the cost; it defers and compounds it. Concealed funding is eventually exposed at a rate history suggests approaches certainty, and exposure of concealment confirms the accusation not only against the recipient but against every clean movement in every country, forever. The transparent cost is paid by the recipient who chose it. The covert cost is paid by everyone, without consent, indefinitely. Second, the design can reduce the exposure: methodology can be severed from money, so that what transfers is published technique requiring no financial relationship at all — the elderly academic's model, which survived the accusation precisely because there was no funding to designate. Where money is genuinely needed, pooled multilateral mechanisms with published criteria at least deny the designation regime a single national adversary to name. Third, and hardest: a movement's capacity to survive the foreign-agent accusation while accepting published assistance is itself a reading of the latent conditions. A society in which the accusation lands and sticks against a transparent recipient is telling the observer something about the state of its common knowledge — and the doctrine's instrument should register it as data rather than treat it as noise to be engineered around. The engineering around it is how the last seventy years happened.

Refusability, finally, cuts in both directions, and the second direction is the one the assisting party will find harder. The recipient may refuse. They may take the training and reject the strategy, take the tools and pursue objectives the funder dislikes, accept assistance and then denounce the assister when denouncing them is what the domestic situation requires. Under the doctrine this is not a failure of the program. It is evidence that the program transferred capability rather than control — which is the entire distinction between the two lineages, arrived at from the operational side. What that transfer consists of is the next chapter.

CHAPTER TWENTY

Custody of the Instrument

In the autumn of 2002 the United States intelligence community produced, in roughly three weeks, a National Intelligence Estimate on the Iraqi weapons programs. The estimate was commissioned late, under pressure, in a decision environment where the policy it would inform was visibly already made. Its headline judgments were confident. The dissents — and there were dissents, on the aluminum tubes, on the uranium claims, from agencies whose technical people had looked hardest at the evidence — were present in the document, in footnotes and annexes, at a prominence inversely proportional to their inconvenience. The estimate said what the moment required, the moment proceeded, and the post-mortems, which were extensive and expensive, established what the dissenting footnotes had established already, too quietly to matter.

The episode is usually filed under intelligence failure, alongside the cases in Chapter Nine. It is a different kind of failure and a worse one. Chapter Nine described an instrument that could not see the relevant quantity. The 2002 estimate is an instrument that could see, and said otherwise — an assessment apparatus bent, under institutional pressure, toward the finding its principals needed. No one had to falsify anything. Emphasis, placement, confidence language, and deadline did the work.

This chapter exists because the instrument specified in Part Four would face exactly that pressure, and would face it most intensely at exactly the moments its null finding mattered most. Chapter Fourteen conceded the vulnerability: a null state can be suppressed. Chapter Sixteen established the deeper problem: the instrument is dual-use, and the requirements it meets are the requirements a sophisticated control apparatus would also want met. Both chapters handed the same question forward. Who can be trusted to hold this thing?

• • •

Work through the candidates, against the book's own record.

An intelligence service fails on the evidence of Part Two and the mechanism of Chapter Seven. The institution's success metrics, career incentives, and consumer relationships all reward positive findings; the 2002 estimate shows what those pressures do to inconvenient judgments even inside a mature oversight regime. An instrument whose primary validation criterion is its willingness to report that nothing is happening cannot be domiciled in an institution that has never once been rewarded for reporting that nothing is happening.

A foreign ministry fails on politicization of a different flavor: assessments aligned to the diplomatic position of the season, and a null finding that becomes unsayable whenever the government of the day has committed itself publicly to the proposition that a given regime is crumbling.

A private contractor fails on Chapter Six. The operative corrupting variable identified there — a commercially interested party with specific access — is not a risk the contractor model mitigates; it is the contractor model's description. An assessment shop whose revenue depends on renewals from clients who want particular findings is the capture mechanism with a rate card.

A multilateral body comes closest and fails last. Pooling custody across democratic states dilutes any single government's pressure, and the model has partial precedents in election observation. But the members capture what they jointly own; several plausible members maintain, at home, the very instruments tabulated in Chapter Eight; and a body that must reach consensus before publishing a finding has built the suppression of inconvenient findings into its charter and called it procedure.

Every custodian fails the same way: each is an institution with interests, and an instrument with a null state is an instrument whose most valuable output is the one every interested party will, at some moment, need it not to produce.

• • •

The proposal that follows from the book's own logic is the uncomfortable one. *Publish the instrument.* The methodology open, the requirements open, the validation record open, the findings — wherever findings are produced — open, and the capacity to run it distributed among universities, research institutes, and civil society organizations in many countries rather than licensed to any government. Custody is resolved by dissolution: an instrument that everyone holds, no one controls, and an inconvenient null produced by one team can be replicated by another, in another country, outside the reach of the party that needs it suppressed.

The safety case rests on Chapter Eighteen, and stands or falls with it. If the asymmetry holds, publication gives a stable autocracy nothing it can use — the instrument returns null on a society whose latent distribution is not in motion, and a null is not actionable. What publication gives an unstable regime is sight of its own condition, and it is worth being unsentimental about what that means: early warning of legitimacy collapse, in the hands of the collapsing, is not obviously a harm to anyone. The alternative on current arrangements is a nuclear-armed state discovering its own terminal condition by surprise, over a fortnight, with its command structure dissolving around its arsenal. A world in which regimes can read their own decline is a world with fewer Kabul-speed collapses, and Kabul-speed collapses are in no one's interest, including the populations underneath them.

Honesty about what publication does not solve. It does not prevent a well-resourced state from building the high-resolution version — the instrument without the floor, the one Chapter Sixteen described. Nothing prevents that; the floor binds the builders who accept it, and the argument for maintaining it in the published version is not that adversaries will honor it but that the published instrument must not function as a reference design for the targeting version. It does not remove pressure on the people running assessments; it relocates the pressure to a venue — open methodological dispute among independent teams — where suppression at least generates a visible anomaly, a published null contradicted by an official posture, of the kind the 2002 process was structured never to produce. And it does not answer the question of whether governments will listen to an instrument they do not own. They will not, often. The instrument's purpose is not to compel its consumers. It is to exist on the record, so that the Chapter Twenty-Two override — the decision to act where the reading said null — has to be made in daylight, against a finding that cannot be footnoted into an annex.

The domestic case remains, and it should be answered rather than avoided. An instrument that reads legitimacy distributions can be pointed at the society that built it, and the reflex is to call that the abuse scenario. The reflex is wrong. Aggregate self-assessment — a democracy reading its own distribution of loyalty, compliance, and expectation, at institutional resolution, in public — is not the instrument's failure mode. It is its highest use, and Chapter Twenty-Three will end the book there. The failure mode is resolution, not jurisdiction: the same reading conducted at the level of named persons is Zersetzung with better mathematics, whoever conducts it and wherever. The floor is the safeguard. It travels with the instrument, or the instrument should not travel.

CHAPTER TWENTY-ONE

Capability Transfer, Not Operations

On the night of 24 September 2000, while the Yugoslav federal election commission prepared to announce results that would send the presidential contest to a second round, a network of domestic election monitors was completing a count of its own. The technique they were using was not secret and had not been invented in Serbia. It had been developed and refined over the preceding fifteen years, most famously in the Philippines in 1986, and it had been transmitted through published manuals and open trainings ever since: station volunteer observers at a statistically constructed sample of polling places, record the posted precinct tallies at the source before they can be altered in aggregation, and compute the national result independently, with margins of error, on election night.

The monitors' count showed the opposition candidate winning outright in the first round. That number did the decisive work of the following two weeks. The regime's insistence on a runoff was no longer a procedural position to be debated; it was a checkable falsehood, checked. The general strike that followed was not framed as an insurrection but as the defense of a result — *he won* — and the number gave every wavering official, every police commander, and every state employee the same fact at the same time. It was, in the vocabulary of Part Four, a common-knowledge machine: a published statistical technique, operated by citizens, whose output converted a million private suspicions of theft into one shared, quantified certainty.

No foreign government ran that count. Foreign assistance trained some of the people who did, funded some of the organizations they worked in, and printed some of the manuals they learned from — openly, attributably, in programs that had been running for years and were described in annual reports anyone could read. Then the capability belonged to the people operating it, and it produced a result no funder controlled and several funders were surprised by. That is the model. This chapter specifies it.

• • •

The test that separates capability transfer from operations can be put in one question: *if the recipient ended the relationship tomorrow, would the capability remain?* If yes — if the technique is learned, the tools are owned, the infrastructure runs without the assister's participation — then what was transferred was capability, and the transfer is consistent with the doctrine. If no — if the activity stops when the funding stops, if the strategy lives in the assister's planning documents, if the local organization is a delivery mechanism for decisions made elsewhere — then what exists is an operation with local staff, whatever the grant agreement calls it, and it fails the Chapter Nineteen test even if every dollar of it is disclosed.

What passes the test, concretely: *training*, offered on published curricula, taken or left, in strategy, organization, security, and the accumulated case knowledge of the tradition — the Bangkok manual's descendants, taught in the open. *Assessment and forecasting capability*, meaning the Part Four instrument itself, published per Chapter Twenty, so that movements can read their own societies with the same rigor that states read them — including the null finding, which tells a movement that the cascade it believes imminent is not, and that the correct strategy is patience, which is the finding movements most need and least receive. *Communications and organizational security infrastructure* owned outright by its users: tools, code, and hardware that continue functioning if every funder disappears, audited in the open precisely so that the users need not trust the provider. *Documentation support* — the capacity to record, preserve, and authenticate evidence of repression, which serves the mechanism (repression made visible is repression made expensive) and serves justice afterward. *Legal defense*, for the people the Chapter Eight instruments reach. And *convening*: the unglamorous work of putting practitioners from different countries in the same room and leaving.

The governing rule across all six: the transferring party retains no operational control, no veto, and no reporting relationship that converts assistance into direction. The rule has a corollary that must be accepted in advance rather than discovered in crisis, because it is the price of the entire doctrine: *transferred capability will be used in ways the transferring party dislikes.* Parallel counts will validate elections won by parties the funder opposes. Security tools will protect organizers whose politics the provider finds noxious. Trained strategists will campaign against governments the assisting state is courting. Each of these is not a program failure to be engineered out with tighter grant conditions. Each is the proof that the thing transferred was capability — and every tightening of conditions to prevent such outcomes is a step back across the line, toward the relationship where the recipient is an instrument, the assistance is a lever, and the whole apparatus of Part Two begins reassembling itself with better vocabulary.

A last boundary, from Chapter Seventeen, drawn here so the doctrine cannot be extended past its own argument. Capability transfer reaches its limit at the transition. What a movement needs after victory — the institutional capacity to govern — cannot be transferred from outside for the same structural reason the willingness to withhold obedience could not be: it derives from the people who will live under it or it does not exist. Assistance in the consolidation period is where refusability is hardest to maintain, where the assister's leverage is greatest, where the historical record is worst, and where the doctrine therefore counsels its own most restrictive application, not its most generous. The moment of victory is the moment the assisting party should be getting smaller, and every institutional instinct will run the other way. Which is the subject of the next chapter.

CHAPTER TWENTY-TWO

The Judgment That Gets Corrupted

Before dawn on 30 April 2019, on an overpass beside a military airbase in eastern Caracas, the man recognized by some sixty governments as Venezuela's legitimate interim president stood with a small group of soldiers and a recently freed opposition leader and announced that the final phase had begun. The army, he said, was moving; the senior figures were with them; the end of the regime was a matter of hours. The call went out for the garrisons to come over.

The garrisons did not come over. The senior figures who had reportedly signaled their readiness did not move; whatever had been understood in the preceding weeks' contacts, the understanding evaporated in the event. By afternoon the crowds were being dispersed, the freed opposition leader was in a foreign embassy, and the government that was to have fallen by evening had instead received a live demonstration that its armed forces would hold — a demonstration it could not have purchased, delivered free by its opponents.

The episode is not offered as an equivalence with Part Two; no proper covert operation ran that morning, and the men on the overpass were Venezuelans risking themselves in their own country. It is offered as a reading failure of the exact kind this book has been building the vocabulary to name. Every input to the decision was an *opportunity* reading: sanctions biting, a paper case of elite contacts, sixty recognitions, an external coalition politically committed to imminent collapse and institutionally unable to hear otherwise. The *legitimacy* reading — where the mutual expectations of the officer corps actually stood, whether any colonel believed the other colonels would move — was the one nobody could produce, and into that absence flowed every interest that needed the answer to be yes. An instrument meeting the Part Four requirements would, on the evidence of the day itself, have returned null. The question this chapter asks is what would have happened to that null on its way to the principals — because the honest answer, on the record of this book, is that it would have gone the way of the 2002 footnotes.

• • •

The do-nothing judgment is the one most reliably corrupted, and the corruption has a structure worth stating exactly, because it explains why good people fail at this repeatedly and why exhortation has never fixed it.

The interests arrayed against a null finding are always specific, present, and articulate: the exile community with a timetable, the sanctioned firms with stranded assets, the allied governments with public commitments, the officials whose careers are attached to a policy of pressure, the principal who has said in public that the regime's days are numbered. The interests served by the null are diffuse, future, and mute: the movement not yet discredited, the collapse not yet accelerated into chaos, the credibility not yet spent. This is Chapter Seven's asymmetry, transported from the operational ledger to the decision meeting — concentrated benefits argue, distributed costs do not attend. And the corruption never announces itself as the overriding of an assessment. It arrives as reasonable dispute about an uncertain judgment: the conditions are debatable, the sources disagree, the window is closing, and surely the reading is too cautious. Every override in the historical record was, in the room, a judgment call about latent conditions. That is what makes the judgment the point of attack.

Exhortation having failed for seventy years, the remedy proposed here is institutional design, offered as the actual substance of the ethics rather than its administrative afterthought. Four provisions, each answering a specific failure documented earlier in the book.

Separate assessment from action. The people who produce the legitimacy reading must have no stake in what the reading licenses — no operational directorate, no program budget that grows with positive findings, no consumer relationship that rewards actionable product. Chapter Seven showed what happens when the assessing institution owns the activity being assessed; the provision is that it never again does.

Make the null a finding of record. A null result must be a published, dated, attributable document — not the absence of a report, which can be ignored without a trace, but a positive artifact stating that the conditions are absent and no action is indicated. The 2002 mechanism worked by prominence and placement; a null of record cannot be footnoted, because it is the headline of its own document.

Insulate the reading from commercially positioned parties. No entity with a financial position in the outcome — sanctioned assets, stranded concessions, contingent contracts — participates in producing or reviewing the assessment. Chapter Six identified the injured firm's specific access as the operative corrupting variable across the entire historical record. The provision names it and bars the door.

Price the override. Principals retain the authority to act against a null — democracies elect governments to make exactly such judgments, and no instrument should govern. But the act of overriding must itself be a matter of record: a stated decision, acknowledging the finding, owning the divergence, reviewable afterward against what occurred. The purpose is not to prevent the override. It is to make the override expensive in the currency the historical record shows was always free — the ability to act against the reading while claiming there was no reading.

Written without confidence, as promised. Each provision can be gamed by a sufficiently determined institution; the 2002 estimate was produced inside an oversight architecture that looked, on paper, robust. The claim for these provisions is the modest one from Chapter Fourteen: they make the corruption harder, slower, and visible, where for seventy years it has been easy, fast, and deniable. A doctrine that promised more would be exhibiting the disease it diagnoses — a positive finding in the absence of the conditions for one.

CHAPTER TWENTY-THREE

The Real Deliverable

This book opened on two manuals. The first was written in March 1954, classified, circulated among officers, and devoted to the removal of a foreign head of government; its assessment section recorded a success, and the assessment was correct by everything it measured and catastrophic by everything it did not. The second was published in Bangkok in 1993, photocopied without permission, and devoted to the withdrawal of obedience; it belonged to no one, directed no one, and has outperformed the first on every measure this book has been able to construct.

The book closes by proposing a third document, and by being precise about what kind of document it is. Not a better version of the first manual — not an operations doctrine at all. The third document is the instrument specified in Part Four: a published, open, replicable method for reading the distribution of legitimacy in a society, calibrated first and always to its own null state, floored at institutional resolution, held by everyone and controlled by no one. The deliverable of seventy years of failed tradecraft is not a capability to act. It is a capability to see.

• • •

The case has been made in pieces and can now be assembled in a paragraph. The capability Part Two built was an answer to the question *who holds the lever, and can we reach it* — a question whose pursuit captured a democracy's foreign policy for commercial interests too small to have justified it (Part Two), armed every autocrat with a permanent alibi against his own citizens (Part Three), and left the institutions that pursued it structurally blind to the one quantity that determines political outcomes (Part Three again, at Havana, Tehran, and Kabul). The capability that was never built was an answer to the question *what does this society expect of itself* — and Part Four specified what answering it requires, while Part Five specified how the answering can be organized so that it does not simply reconstitute the thing it replaces. The asymmetry argument is the load-bearing wall: the seeing-capability fails safe where the acting-capability failed catastrophic, because sight of a cascade that is not latent is sight of nothing.

A democracy so equipped does not need to cause collapses. It can perceive them forming — with time to position for them, protect people through them, and refrain from the interventions that its old instrument, measuring opportunity and calling it legitimacy, would have licensed. That is the strategic case, and it is sufficient. It is not the deepest case.

• • •

The deepest case is the one Chapter Twenty declined to treat as an abuse scenario. The instrument reads in every direction, including homeward, and the book ends on the homeward reading because the book's own argument requires it to.

Every mechanism described in Part Four operates in democracies. Preference falsification is not an authoritarian specialty; it is the condition of any society where expressing certain conclusions is costly, and the costs need not be prison to be effective. Pillars have loyalty distributions in republics too — officer corps, civil services, courts, state legislatures, media institutions — and those distributions move, and expectation runs ahead of expression there as everywhere. A democracy's characteristic failure is to mistake its elections for a legitimacy instrument and conclude that it needs no other: but an election measures expressed preference at intervals, at the ballot's resolution, and Chapter Twelve established that expressed preference is the least load-bearing of the three quantities. The load-bearing one — whether the members of the arrangement expect the others to keep complying with it, whether the officials expect the other officials to certify, whether the losing coalition expects the institutions to hold — is measured by nothing at all, in any democracy now existing, and the recent history of several of them suggests the reading would not be uniformly reassuring.

The state that spent seventy years trying to manipulate legitimacy abroad never learned to perceive it, and the incapacity was indivisible: the instrument it failed to build for Tehran is the instrument it lacks for itself. Building it now, in the open, floored and null-calibrated and owned by no one, is therefore not a foreign policy proposal with a domestic side effect. It is a single reform with one purpose wearing two directions — a state, and a society, capable at last of seeing the thing on which everything else it values stands.

The first manual taught how to end a government. The second taught how obedience is withdrawn. The third would teach nothing but how to look — and of the three, it is the only one that could not have been used to write Part Two, and the only one every party to the question, ruler and ruled, foreign and domestic, is better off having exist.

That is the wager of this book. Not that democracies will stop wanting levers, but that an instrument built to see — built so it can say *nothing is latent here, do nothing* — is the one piece of tradecraft that cannot be captured by the interests that captured all the others, because the only thing it delivers is the truth about whether the moment is real. Everything in the record says the moments that were manufactured ended badly, and the moments that were real needed no manufacturing. A method that can tell the difference is worth more than every operation it declines to authorize.

It would fit in a small book. Ninety-odd pages, numbered sections, published openly, copied without permission. Written for the explicit benefit of people who will need to understand something clearly, somewhere else, later.

CASE FILES AND VERIFICATION RECORD

> *This document does two jobs. Part A builds out the evidentiary base for Chapter Six — the population of cases, the coding of each, and the discriminating variables — so that the chapter's method is auditable rather than asserted. Part B records the verification pass over every dated or numerical claim in the manuscript, marking each as confirmed, corrected, or flagged for further work. Corrections marked here have been applied to the current draft.*

PART A · THE CHAPTER SIX CASE POPULATION

Population definition

Every government, between the establishment of a standing American covert-action capability in 1947 and the end of the 1970s, that materially injured a major Western commercial interest through expropriation, nationalization, unilateral contract revision, or comparable measure — whether or not anything subsequently happened to it. Two pre-1947 cases (Mexico 1938) are carried as repertoire evidence, not as preference tests, for the reason given in the coding notes. Two cases with no commercial trigger are carried as counter-evidence.

Coding variables

- **INJURY** — scale of the commercial injury relative to the target economy and to the injured firm.
- **CONTINUING POSITION** — did the injured firm retain an income-producing role (management contract, marketing agreement, minority equity) after the measure? This turns out to matter more than whether compensation was nominally offered.
- **PARTNER** — presence of an organized, capable domestic actor with prior institutional relationship to the United States and an appetite for power.
- **ALIGNMENT** — availability of a bloc patron, and whether the security characterization preceded or followed the commercial dispute.
- **ACCESS** — informal, personal access by the injured party to deciding officials. Largely unmeasurable; recorded where documented.
- **OUTCOME** — covert removal attempted / not attempted.

A.1 Positive cases

Iran, 1951–53

Injury. Majlis vote to nationalize the Anglo-Iranian Oil Company, spring 1951; AIOC majority-owned by HM Government since 1914. The 1950 Aramco–Saudi fifty-fifty settlement made the Iranian terms conspicuously anomalous; AIOC declined to match. British response: embargo, withdrawal of technical staff, closure of Abadan, asset freeze, legal campaign against Iranian crude.

Continuing position. None — total loss of concession. **Partner.** Present and multiple: monarch, sections of the officer corps, mobilizable clergy, bazaar networks. **Alignment.** Tudeh present but Mossadegh not communist; Truman administration explicitly assessed removal as counterproductive and refused. Security characterization *followed* the change of administration, not a change in evidence. **Access.** Extensive, via British government as shareholder and via the incoming administration's own principals.

Outcome. TPAJAX, August 1953. First attempt failed 15–16 August; second succeeded 19 August. **Settlement:** Consortium Agreement, signed 19–20 September 1954, in force 29 October — BP 40%, five US majors 8% each (40% total), Shell 14%, CFP 6%. In 1955 the US majors each ceded 1% to a group of smaller American firms (Iricon, 5%). American firms held no position in Iranian oil before 1951.

Sources: FRUS 1952–54 vol. X (esp. d.432, memorandum of understanding, 9 April 1954); Encyclopaedia Iranica, "Oil Agreements in Iran"; consortium text.

Guatemala, 1952–54

Injury. Decree 900, Agrarian Reform Law, 17 June 1952. Expropriation of uncultivated holdings above threshold, compensated at declared taxable value in twenty-five-year bonds at 3%. February 1953 order expropriated approximately 234,000 acres of United Fruit land near Tiquisate; company appeal denied by Guatemalan Supreme Court, 18 March 1953. Compensation offered: **$627,572**. Valuation demanded by the State Department on the company's behalf: **$15,854,849** — a ratio of about 25:1. UFCo held roughly 550,000–600,000 acres, most uncultivated. Across the whole programme some 1.5 million acres were distributed to about 100,000 families.

Continuing position. None on the expropriated acreage. **Partner.** Present — an officer corps with no attachment to the government it served, which declined to fight. **Alignment.** Weakest of any positive case. PGT small; Árbenz not a member; no Soviet arms until the Czech shipment of May 1954, which the arms embargo had itself induced and which was then cited as evidence of alignment. **Access.** Documented and extensive.

Outcome. PBSUCCESS, June 1954. Árbenz resigned 27 June. Decree 900 repealed; land restored. Subsequent civil war 1960–96; CEH truth commission (1999) documented approximately 200,000 dead, overwhelmingly at the hands of state forces, and found acts of genocide against Maya communities.

Sources: FRUS 1952–54 vol. IV, editorial note (d.420); Handy, "The Most Precious Fruit of the Revolution," HAHR 68:4; CEH, *Memoria del Silencio*.

Brazil, 1962–64

Injury. Profit remittance law restricting annual repatriation; state-level expropriation of an ITT telecommunications subsidiary in Rio Grande do Sul at a disputed valuation; proposed refinery nationalization and land redistribution. Modest by the standards of this population. **Continuing position.** Partial — foreign firms remained operational. **Partner.** Strongly present: large, politically organized, anti-communist officer corps with decades of institutional contact. **Alignment.** Goulart a landowner and constitutional successor; alignment case weak but assertable. **Outcome.** Coup 31 March–1 April 1964, with a US naval contingency package (carrier, escorts, tankers, ammunition ship) dispatched and recalled unused. Military rule to 1985.

Indonesia, 1965–66

Injury. Sukarno-era seizure of Dutch holdings (1957–58) and sustained pressure on American oil operations. **Partner.** Present. **Outcome — and coding note.** This case must be coded differently from the others and the chapter says so: the United States did not organize the killings. It encouraged them, signalled approval to the officers conducting them, supplied communications equipment and other material, and — through an embassy officer who later described doing so publicly — passed lists of PKI names compiled from open sources. Embassy reporting declassified in bulk in 2017 records the scale contemporaneously and assesses it favourably. Estimates of the dead begin at 500,000; roughly a million more imprisoned, many for over a decade, without trial.

Commercial sequel (recorded as sequel, not cause): foreign investment law 1967; a major American mining contract of work concluded 1967; oil sector reopened on terms the previous government had refused — within roughly twenty-four months of a decade of failed negotiation.

Verification note: the list-passing claim rests on the officer's own on-record account and subsequent documentary release. Characterize as attested by the participant and consistent with released cables; do not characterize as the operative cause of the killings.

Chile, 1970–73

Injury. Copper nationalization by constitutional amendment, July 1971, passed unanimously by a congress the government did not control. Compensation formula deducted designated excess profits from book value, leaving the two large American producers with nothing and in one calculation less than nothing. **Continuing position.** None. **Partner.** Present by 1973, though not manufactured: the 1970 attempt to induce a coup required removing a constitutionalist commander-in-chief and left him dead, and failed. **Alignment.** Assertable. **Access.** Documented, including corporate offers of funding. **Outcome.** Three years of credit denial, opposition funding, and media support; coup 11 September 1973. American contribution to the final event is genuinely contested; the chapter does not require it to have been decisive.

A.2 Negative cases

Mexico, 1938 — *REPERTOIRE EVIDENCE, NOT A PREFERENCE TEST*

Total expropriation of seventeen foreign oil companies, 18 March 1938, following the companies' refusal to comply with a Mexican supreme court labour award. Larger and less qualified than any injury in the positive set. Response: boycott of Mexican crude, British rupture of diplomatic relations, sustained American pressure for compensation, ambassadorial argument against coercion, compensation negotiated over several years at figures the companies considered derisory. State retained the industry permanently.

Coding note, stated in the chapter text. In 1938 no standing covert-action capability existed. A government cannot decline to use an instrument it does not have. Mexico therefore tests repertoire — establishing that boycott, pressure, litigation and settlement were a complete and functional response set — and does not test preference. Its value to the argument is to reframe the question put to 1953 from "what else could they have done" to "why did they stop doing what had already worked."

Bolivia, 1952 — *STRONG NEGATIVE*

MNR takes power by insurrection, April 1952; existing army dissolved; worker and peasant militias armed; the three great tin companies nationalized 31 October 1952. Relative to national economy, more radical than Chile. American response: recognition within weeks, then — under the same administration that removed Mossadegh and Árbenz — among the largest per-capita aid programmes in the hemisphere. On-record reasoning: the governing party was containable and anti-communist enough, compensation was offered, and the realistic alternatives (Trotskyist currents, the miners' militias) were assessed as worse. **PARTNER absent** — the army had been dissolved. This case is the cleanest demonstration that partner availability, not injury magnitude, is doing the work.

Egypt, 1956 — *STRONG NEGATIVE*

Nationalization of the Suez Canal Company, 26 July 1956. American response: opposition to, and then termination of, the Anglo-French-Israeli intervention, using financial pressure on sterling to force withdrawal. No case in the period features a larger commercial expropriation met with a more emphatic American defence of the expropriating state. Later hostility to Nasser does not alter the coding of the 1956 response.

Peru, 1968–74 — *STRONG NEGATIVE; MATCHED AGAINST BRAZIL*

Coup 3 October 1968 (Velasco). **Six days later**, 9 October, the Revolutionary Government renounced the Act of Talara and occupied the International Petroleum Company's refinery and the La Brea y Pariñas fields — a subsidiary of Standard Oil of New Jersey — declining compensation on the theory that IPC owed more in back taxes than the assets were worth (Peru claiming c. \$690m owed; IPC demanding c. \$120m). Subsequent nationalizations 1969–74: ITT, mining, fisheries, sugar, media. Soviet arms purchased.

Response. Hickenlooper Amendment triggered but its formal application avoided; special ambassador (Irwin) recommended against sanctions; aid and multilateral lending curtailed in practice. **Settlement:** agreement of 19 February 1974 — Peru paid **\$76 million** to the United States government for distribution among claimants, with a unilateral Peruvian annex listing eligible companies constructed so that Peru could maintain it had not compensated IPC. IPC received nothing. Velasco was removed in 1975 by other Peruvian generals for Peruvian reasons.

Why it matters. Goulart did substantially less and drew a naval contingency package; the Peruvian junta did substantially more and drew lawyers. PARTNER absent: the officer corps *was* the government.

Sources: FRUS 1969–76 vol. XI pt. 2, editorial note d.294; Maurer, *The Empire Trap* working paper (Greene settlement figures).

Zambia, 1969 — *STRONG NEGATIVE; MATCHED AGAINST CHILE*

Matero Reforms announced 11 August 1969: state acquisition of 51% of the copper mines from Anglo American Corporation and Roan Selection Trust, implemented January 1970. Copper was approximately 96% of Zambian exports. Kaunda nonaligned; hosted southern African liberation movements.

CONTINUING POSITION — present, and this is the finding. Terms settled amicably within three months: compensation at book value in state bonds redeemable over eight to twelve years at 6% interest, plus ten-year management and sales contracts paying the companies approximately 1.5% of gross sales and 2% of profits. The firms lost majority ownership and kept an income stream. Compare Chile, where the excess-profits deduction left the producers with nothing. Both were copper; both were the entire export economy; one produced three years of covert effort and one produced a bond schedule.

Sources: Kaunda, "Towards Complete Independence," Matero Hall, 11 August 1969; Bostock and Harvey; Whitworth (2015) on ZIMCO bond terms.

Venezuela, 1976 — *NEGATIVE*

Oil nationalization effective 1 January 1976; former concessionaires compensated at book value on terms they considered inadequate and retained service and technical contracts. Democratic government, OPEC member, aligned. No consequence beyond continued commercial relations. Coding: CONTINUING POSITION present; PARTNER absent; ALIGNMENT favourable.

India, Nehru period — NEGATIVE

Sustained nationalization across insurance (1956), banking, and heavy industry; declared nonalignment; Soviet arms purchases; substantial and continuous American aid throughout, including large food assistance. PARTNER absent; consolidated state with a military attached to civilian authority.

A.3 Counter-evidence: intervention without commercial trigger

South Vietnam, 1963

An allied head of government removed in a coup with American encouragement, against no expropriation, at no commercial party's request, for reasons internal to the prosecution of a war. Whatever explains 1963, it is not the economic thesis. Carried in the population deliberately: a test that admits only inconvenient negatives is half a test.

Congo, 1960–61

Carried with a coding caveat. Mineral interests were substantial and Belgian corporate involvement documented, but the American decision chain is better explained by bloc-competition variables than by a specific injured firm with access. Recommend including as a mixed case and coding ACCESS as not established.

A.4 The coding table

Case	Injury	Continuing position	Partner	Alignment	Outcome
Iran 1951–53	Total	None	**Yes**	Asserted after admin change	Removal
Guatemala 1952–54	Partial (uncultivated)	None on taken land	**Yes**	Weak; partly induced	Removal
Brazil 1962–64	Modest	Partial	**Yes**	Assertable	Removal
Indonesia 1965–66	Prior seizures	Partial	**Yes**	Strong	Enabled, not run
Chile 1970–73	Total, negative comp.	None	**Yes** (by 1973)	Assertable	Sustained effort
Mexico 1938	Total	None	No	n/a	*None (no capability)*
Bolivia 1952	Total, core economy	None	**No** (army dissolved)	Containable	*Aid*
Egypt 1956	Total	None	No	Hostile	*Defended*
Peru 1968–74	Total, uncompensated	None	**No** (officers govern)	Soviet arms	*Settlement*
Zambia 1969	Majority stake, core economy	**Yes** (mgmt + sales)	No	Nonaligned	*None*
Venezuela 1976	Total	**Yes** (service contracts)	No	Aligned	*None*
India 1950s–60s	Sustained, sectoral	Mixed	No	Nonaligned + Soviet arms	*Aid*
S. Vietnam 1963	**None**	n/a	Yes	Central	Removal

A.5 What the table supports

Commercial injury is neither sufficient nor necessary. Six negatives with injuries equal to or greater than the positives; one positive (1963) with no injury at all. The crude economic thesis fails on this population and the chapter must say so in its own voice before a reviewer says it.

PARTNER discriminates almost perfectly. Every positive has one; every negative lacks one. This is the strongest single finding and it means the null hypothesis — that intervention followed opportunity — is not a rival to be defeated but a necessary condition to be incorporated.

CONTINUING POSITION out-performs compensation-offered. Bolivia, Guatemala, and Iran all involved compensation offers; two produced operations. Zambia and Venezuela left the firms with income; neither produced anything. The operative question is not whether a firm was paid but whether it was left with a reason not to lobby.

The defensible thesis. Commercial injury determines *which disputes reach the table*; partner availability determines *whether anything is done about them*. This is a weaker causal claim and a far more damaging institutional one, because agenda capture requires no one to be corrupt — the injured firm is simply the only party in the room with a specific, quantified, urgent request, and specificity is what gets acted on.

A.6 Limitations to keep in the chapter body, not in a note

- **Disclosure is not random.** The population is built from disclosed cases; exposed and failed operations are overrepresented; successful ones may be missing entirely.
- **The denominator is unknown.** Considered-and-rejected operations are almost wholly unobservable, which means PARTNER's apparent necessity could partly reflect what got proposed rather than what got approved.
- **Capability confound before 1947.** Handled explicitly for Mexico; do not extend the reasoning to post-1947 cases without saying so.
- **Documents were written by people who knew they might be read.** Sincerity of security characterizations cannot be assessed from the record alone; the timing test (did the characterization precede or follow the dispute) is a partial substitute and is used as one.

PART B · VERIFICATION RECORD

> *Every dated or numerical claim in the manuscript, checked against available sources. Corrections marked* CORRECTED *have been applied to the current draft. Items marked* FLAGGED *require primary-source work before publication.*

B.1 Confirmed as drafted

Claim	Verified detail	Note
1954 consortium shares	BP 40%; five US majors 8% each; Shell 14%; CFP 6%. Signed 19–20 Sept 1954, in force 29 Oct.	US firms held nothing pre-1951. 1955: each major ceded 1% to smaller US firms (Iricon 5%).
UFCo compensation dispute	$627,572 offered vs $15,854,849 demanded; c. 234,000 acres, Feb 1953; Decree 900 of 17 June 1952.	Ratio is c. 25:1, not 30:1. Figures now stated exactly in text.
Peru IPC seizure and settlement	Coup 3 Oct 1968; seizure 9 Oct; agreement 19 Feb 1974; $76m to USG for distribution; IPC excluded by Peruvian annex.	"Six days later" is exact. Hickenlooper triggered, application avoided.
Zambia Matero Reforms	Announced 11 Aug 1969; 51% of AAC and RST; book-value bonds 8–12 yrs at 6%; management and sales contracts retained.	Copper c. 96% of exports. Settled within three months.
ZunZuneo	USAID; 2009–2012; c. 40,000 subscribers at peak; shell entities in Spain and Cayman Islands; staged plan to "smart mobs"; AP investigation April 2014.	Planning language on renegotiating state–society balance is in the documents AP obtained.
Chenoweth/Stephan	1900–2006: c. 53% nonviolent vs c. 26% violent success.	See correction B.2 on the decline.
Iranian indictment citing Sharp	Show trials from 1 Aug 2009; indictment asserts more than 100 of the 198 steps executed.	Text available via Iran Data Portal. Now used directly in Ch. 19.

B.2 Corrected — applied to the draft

Ch. 9 — the August 1978 Iran assessment MATERIAL

The "not in a revolutionary or even pre-revolutionary situation" line comes from the January 1979 House Permanent Select Committee staff evaluation, quoting a CIA product of August 1978. **The then-Director subsequently disputed its status publicly, stating it came from a draft he had rejected rather than a finished study.** Presenting it as a settled formal assessment was the manuscript's most exposed factual claim and has been rewritten: the chapter now records the dispute and shifts weight onto the Defense Intelligence Agency appraisal of the following month, which judged the Shah likely to remain actively in power for another decade and was not a draft. The argument is strengthened by the change, because two agencies with different sources agreeing is a better illustration of a systemic ontology problem than one document is.

Ch. 12 — Stănculescu MATERIAL

The drafted account had him refusing to send Ceaușescu a helicopter. By his own later account he *arranged* the helicopter, describing it as saving the couple from a mob. Corrected sequence: the plaster cast was self-arranged to avoid orders to fire on demonstrators; Defence Minister Milea died on the morning of 22 December (recorded as suicide after failing to execute such orders); Stănculescu became acting minister that day; he ordered units back to barracks and arranged the departure; he organised the trial of 25 December. In 2008 a Romanian court convicted him over the Timișoara shootings (fifteen years; released 2014). The chapter now uses the cast rather than the helicopter as its evidentiary object and states the conviction, which strengthens the analytical point: expectation is not a moral quantity.

Ch. 13 — Jäger at Bornholmer Strasse MINOR

Time corrected from 11:29 to approximately 11:30 pm. He commanded roughly fifty men, not an unspecified force. The "recording names" detail replaced with the accurate one: he was instructed to admit the loudest and stamp their documents so as to prevent return — which required individual identification in front of a crowd already moving the barrier. Approximately 20,000 crossed in the following forty-five minutes; other crossings followed by midnight.

Ch. 16 — Stasi ratios MINOR BUT QUOTABLE

Drafted as "one full-time officer for every sixty-three citizens" — incorrect. Correct: c. 91,000–97,000 full-time staff (c. 1 per 166 citizens) plus c. 173,000 registered informants, yielding the commonly cited 1-per-63 for *officers or informants combined*. The 1-per-6.5 figure includes occasional and part-time collaborators and is an upper-bound estimate of a contested count. Text now states all three with their qualifications. The "files on a third of the population" claim was removed pending a citable figure.

Ch. 2 — the decline in civil-resistance success SUBSTANTIVE ADDITION

Drafted vaguely as a narrowed margin. Corrected and sharpened: nonviolent success rates fall from roughly 52–53% to roughly a third when the dataset is extended through the 2010s, even as campaign frequency peaked. Chenoweth's own attribution — movement-side thinning (organization, training, strike usage, over-reliance on the internet) plus regime learning, with *blaming foreigners and outsiders* listed among the learned responses — has been incorporated, and now sets up Chapter Eight directly.

Ch. 19 — the Sharp animation DATE

Drafted as following the 2009 protests. The Iranian intelligence cartoon depicting Sharp with an American senator and a financier dates to 2007–2008, i.e. it *preceded* the protests it was later used to explain. Corrected, and the chapter now leads with the 1 August 2009 indictment instead, which is stronger material: a state prosecutor citing a published book's numbered list as evidence of conspiracy.

B.3 Flagged — requires primary-source work

- TPAJAX BUDGET. Figures in the low hundreds of thousands with authorization around $1m are widely repeated but variously sourced. Either cite the specific document or replace with a ratio claim that does not depend on the absolute number.
- ROOSEVELT DECLINING PBSUCCESS. Currently hedged twice with "reportedly." Source it or cut it — the sentence's work (the method no longer depended on the man) survives without it.
- ABADAN HOUSING AND WORKFORCE DETAIL. The characterization is sound and standard, but the chapter's most humanly specific passage should rest on a named study rather than general knowledge.
- BRAZIL TASK FORCE COMPOSITION. Confirm vessel types and the ammunition-ship claim against the declassified record rather than secondary accounts.
- INDONESIA LIST-PASSING. Attested by the participant on record and consistent with the 2017 releases. Fix the characterization precisely and do not let it drift toward causation.
- KABUL 2021 FIGURES. Force size and the eleven-day span are broadly reported; the assessment timeline (what was judged, when, by whom) needs the actual documents, several of which are now available through inspector-general reporting.
- 2002 NIE PRODUCTION TIMELINE. The three-week figure and the placement of dissents are well established in the post-mortems; cite them directly rather than relying on the general account.
- CHAPTER EIGHT STATUTE TABLE. All four entries require checking against current primary texts and amendment histories, several of which have moved recently. The Hungarian instrument, omitted from the table, is worth adding precisely because it was struck down — a negative case for the diffusion argument.
- GUATEMALA DEATH TOLL AND GENOCIDE FINDING. Cite the CEH report directly, with its own language on attribution of responsibility, rather than the widely circulated summary figure.

B.4 Structural changes made in this pass

- **Part II analytical thread.** Chapter Four now opens its historical material by naming the two questions the whole Part is assembled to answer, so the case detail reads as evidence. Chapter Five closes on the co-evolution of bureaucracy and metric — each recorded success lowering the threshold for the next — which hands directly to Chapters Six and Seven.
- **Ch. 18 bridge.** The Serbia/Cuba comparison is now run explicitly through the Part Four machinery: the three quantities of Chapter Twelve applied to each case, and the specific finding that a covert platform cannot supply mutual observation because its participants cannot know who else is participating. The asymmetry is now derived rather than asserted.
- **Algorithmic information environment.** Added to Chapter Sixteen, where it belongs: everything the Stasi did with 173,000 informants is now a byproduct of ordinary platform operation, and the common-knowledge events of Chapter Eleven now propagate through ranking systems that can equally be adjusted to prevent their formation. The mirror problem is restated as a description of existing infrastructure rather than a historical analogy.