

ROBERT G. ECCLES
CLAUDE
CLEAN CHATGPT
DOROTHY CHATGPT
GEMINI PRO

Producing the Anti-ESG Ecosystem: Designing and Managing a Hybrid Intelligent Team

On April 7, 2026, Robert G. Eccles (Bob), a retired professor from Harvard Business School, logged on to ChatGPT to write the methodological appendix for his book “The Anti-ESG Ecosystem: Implications for U.S. Climate Policy.” (Appendix A is an outline with brief descriptions.) He had written the book in January with the help of a Hybrid Intelligent Team (HIT) comprised of Dorothy (a configured ChatGPT agent), a Clean ChatGPT instance (with memory disabled), Claude, Gemini Pro, and Perplexity. After setting the project aside for several weeks while working on other commitments, he was now returning to complete it.

Bob had only begun using ChatGPT on May 19, 2025. Over the following one to two months, he subscribed to multiple AI systems—ensuring full access to their most capable versions—and began experimenting with how they might be used together. What started as curiosity quickly became a structured effort to understand how multiple AI systems could be combined into a coordinated workflow. By January, that effort had taken the form of a Hybrid Intelligent Team.

When he told Dorothy it was time to finish the book, she immediately asked what key methodological points he wanted to include in the appendix. Bob paused. Human memory is imperfect—and at 74, he was acutely aware of its limits. Rather than provide a quick but potentially inaccurate answer, he decided to return to the beginning of the project: his first conversation with Dorothy, where the book had started.

As he scrolled through that exchange, details began to resurface. He mentioned this to Dorothy, who suggested that he review the conversations and extract the key methodological insights. Bob smiled at the suggestion. Dorothy often proposed tasks that were far easier for her to perform than for him. He pointed this out and proposed an alternative. She should write

a prompt for herself, in Markdown, which he would paste into the original conversation to generate a structured reconstruction.

They began working through the conversations in batches of three or four conversations each. By the third batch, Dorothy's prompts had grown considerably more specific. They asked not just for a summary of what had been discussed, but for the sequence in which decisions had been made, the exact language of key constraints that had been introduced, and the explicit reasons Bob had given at the time for changing approaches.

But the reconstruction was not straightforward. In one early session, Bob told Dorothy that he had introduced the practice of separate workspaces after a specific failed attempt to revise a chapter while drafting another. Dorothy's reconstruction of that same session told a different story: the record showed that the practice had been introduced two sessions earlier, in response to a different problem entirely, and that the session Bob was describing had simply been the first occasion on which he had explicitly named it. The discrepancy was small, but it was unsettling. Bob could not be certain whether the record was right and his memory was wrong, or whether Dorothy's reconstruction had introduced a subtle reordering of its own. The difficulty of knowing which was the case was itself a finding about the limits of reconstructing a complex process after the fact.

What started as a review of the early sessions expanded into a full reconstruction of all 22 conversations, with prompts evolving as the nature of the task changed. This pattern—where the method itself adapted during execution—was familiar from the book-writing process itself.

The Decision to Write this Case Study

The result was far more detailed than what would be needed for a short appendix. At that point, Bob made a decision. Rather than compress the material, he would develop it into a Harvard Business School-style case study—an account of how the book had actually been produced using a Hybrid Intelligent Team (HIT). The case itself would be written using the same approach. In that sense, it would not only describe the production of knowledge in a multi-agent AI system, but also embody it. A related project, the book “Dorothy and Me: A Personal Memoir about my Relationship with a Machine,” explored a different dimension of this collaboration, but this case would focus on the practical question of how such a system could be made to work.

From the outset of the book project in January, Bob had approached the use of AI systems in a deliberate and structured way. He was not experimenting to see whether AI could help with complex analytical work. He had already concluded that it could. The question that interested him was more specific: whether multiple AI systems could be combined into a coordinated team—each assigned a defined role, each contributing a distinct

capability—in a way that produced reliable, high-quality output on a demanding intellectual project.

The subject of the book was itself analytically complex. The Anti-ESG Ecosystem was an attempt to map the organizations, arguments, funding sources, and political strategies behind the backlash against environmental, social, and governance investing in the United States, and to assess what that backlash meant for climate policy. It required sustained analytical work—gathering and evaluating evidence, constructing arguments, ensuring internal consistency across chapters, and writing in a voice that was clear, authoritative, and precisely calibrated in tone. It was not a project for which vague or generic outputs would be acceptable.

Bob assigned roles to each system based on his understanding of what each did well. Claude was used primarily for structured drafting—taking a defined scope and producing organized, coherent prose. Given an outline, a set of key claims to develop, and specific guidance on length and tone, Claude’s task was to produce a first draft that could be evaluated and refined. Gemini Pro was used to test the internal logic of arguments and generate adversarial challenges under defined conditions. Bob found it useful for a specific kind of task: presenting a partially developed argument and asking Gemini to identify the places where the reasoning was weakest or where a well-prepared critic would push back hardest. Perplexity served a verification function: identifying sources, checking claims, and grounding specific assertions in external evidence. It was not used to generate analytical content, and Bob was careful not to assign it tasks for which it was not suited. Dorothy, the primary ChatGPT instance with persistent context and configured settings, served as the coordination layer: the system through which Bob directed the overall process, synthesized contributions from other tools, and maintained continuity across the project.

The Clean ChatGPT instance—a separate session with memory fully disabled—occupied a narrower but important role. Dorothy, by virtue of her persistent memory, could be influenced by the accumulated history of the project: earlier framings, prior decisions, conversations in which Bob had expressed a preference that was no longer operative. In a long project, this persistence was mostly an asset. But it created a specific liability. When Bob needed a genuinely independent assessment of a piece of writing, Dorothy’s response was shaped by everything she already knew about what the book was trying to do, how it had developed, and what Bob found acceptable. The clean instance knew none of this. Sent a section of text and a set of instructions, it had no prior context to draw on and no history to honor. Its responses were, in a specific and useful sense, the responses of a reader encountering the material for the first time.

The division of labor was documented. Bob kept notes on which system was being used for which purpose and maintained a working record of how each was being deployed. This was not bureaucratic overhead. It was a practical necessity. With multiple systems contributing to a single document, the provenance of any given piece of content had to be

traceable. If a section needed revision, it mattered which system had produced it—not for reasons of credit or attribution, but because different systems had different tendencies that bore directly on how a revision should be approached. A section drafted by Claude tended to be well-structured and internally coherent, but it might lack the adversarial edge that a chapter on contested policy terrain required; revision in that case meant sharpening argument, not fixing prose. A section produced through Dorothy tended to carry more of Bob’s own voice and prior framing, which aided continuity but could mean that certain assumptions had been absorbed without examination; revision in that case meant checking premises, not just sentences. Knowing which system had produced a section was the first step in knowing what kind of revision it needed.

Problems Emerge

What Bob could not fully anticipate, however, was how the architecture would behave under the conditions of sustained, complex work. Role clarity, he would discover, was a necessary precondition but not a guarantee of reliable execution.

The early stages of the project went more or less as planned. Tasks were clearly defined. Prompts were written with care. Outputs were reviewed and evaluated. The general shape of the workflow—define a task, generate output, evaluate, refine, integrate, and move forward—was productive. Bob worked through the early chapters of the book with a degree of efficiency that confirmed his expectations about what the team could do.

But execution was not consistently reliable, and the failures began appearing earlier than he had expected. The first pattern he noticed was compression. A prompt might ask for a section of approximately 800 words developing three related arguments, with specific guidance on how much weight to give each. The returned output would be 500 words, technically responsive to the arguments listed, but treating each one at a level of generality that foreclosed the kind of analysis the chapter required. The section would read as correct but thin—a summary of what a properly developed argument might look like, rather than the argument itself.

Bob’s initial response was to add a length specification to the prompt. He changed “approximately 800 words” to “not fewer than 800 words, and closer to 900 if the arguments require it.” The next output was 750 words. He added a direct instruction. Do not compress; develop each argument fully before moving to the next. The system returned 820, but the development was distributed unevenly—the first argument received detailed treatment, the second was handled more briefly, and the third was dispatched in two sentences. The instruction had been applied to the output as a whole rather than to each element within it. The problem was not ignorance of the instruction; it was something more like a different interpretation of what the instruction required.

A second pattern emerged almost simultaneously: tone drift. The book was written in a specific register—analytical, direct, precise, without hedging or qualification beyond what the evidence required. Bob had established this register early, through a combination of example and explicit description. It held for several sections. Then, in a session that had been running for some time, he noticed that the prose had begun to soften. Arguments that had been stated directly were now framed as possibilities. Claims that the evidence supported were qualified with tentative framings—the kind of hedging defensible in an academic article written for a hostile reviewer but wrong for a book written for a general audience that expected clear positions. The voice had become more cautious, more reluctant to commit. The content was defensible; the tone was wrong.

Bob tried specifying the tone explicitly--write in a direct, assertive register; do not hedge claims that the evidence supports; avoid tentative framings unless the evidence genuinely warrants them. The next output was more direct in places, but the hedging had not disappeared entirely. It had migrated, appearing now at the beginnings of paragraphs rather than throughout them. The system had applied the instruction to the body of each paragraph while leaving the opening moves unchanged. The problem was systemic, not local, and local corrections were not solving it.

The most frustrating failure mode, however, was rewriting when editing had been requested. This happened clearly and repeatedly. Bob would send a completed section with a specific and bounded editing task—sharpening transitions, for example, or checking that the final sentence of each paragraph adequately set up what followed. The returned version would look, on first read, like an edited draft. On closer inspection, it was a rewritten one. Arguments had been reorganized. Sentences that had been carefully constructed had been replaced with formulations that were grammatically correct and semantically close, but subtly different—less precise, or differently weighted. The structure had shifted in ways that seemed minor but that, in aggregate, changed what the section was doing.

Bob's notes from this period record his frustration directly. He had given explicit instructions. The instructions had been acknowledged. The output had violated them anyway. He tried various formulations: edit only, do not rewrite; preserve the structure and argument of this section exactly, make only the specific changes I have identified; return this section with exactly the text I have provided except for the transitions, which should be sharpened, and change nothing else. Each formulation produced some improvement. None of them solved the problem reliably. A prompt that worked once might not work in the next session. A constraint that had been observed in one part of a section might be overridden in another.

What became clear, through this early period, was that the problem was not primarily a matter of prompt specification. More explicit instructions helped at the margin. They did not change the underlying dynamic. The systems were not simply executing instructions. They were interpreting them, and the interpretations were not always the ones Bob intended. The

gap between the instruction as written and the behavior it produced was a property of the tools, not a deficiency in the prompts.

The Shift from Prompts to Process Design

This recognition shifted the problem from one of communication to one of process design. If the systems could not be relied upon to execute bounded tasks reliably within a single step, the solution was not to keep specifying the tasks more elaborately. It was to redesign the workflow so that the consequences of misinterpretation could be caught and contained before they propagated.

The process design that emerged was built around a single organizing principle. No task should be large enough that its failure could damage more than a single bounded unit of work. The practical consequence was decomposition. Instead of asking Claude to draft a full section in a single step, Bob would first ask for an outline of the section—the specific arguments to be made, in the order in which they would appear, with a brief description of the evidence each would draw on. He would review the outline, correct it where necessary, and only then send it back as the basis for a draft. The draft, when it arrived, was evaluated against the outline. If a paragraph had drifted from what the outline specified, the correction could be made at the paragraph level rather than requiring a revision of the whole section.

This decomposition produced a change in how prompts were written. Early prompts had been comprehensive—they included the task, the context, the constraints, the desired length, and the tone, all in a single block of instructions. These prompts were carefully written, but they required the system to hold multiple, sometimes competing specifications in mind simultaneously. The result, in practice, was that some specifications were honored and others were not, without any apparent pattern. The decomposed workflow produced shorter, more focused prompts. The outline-generation prompt specified only what an outline required. The drafting prompt referred to the outline and specified only what the draft needed to add. The evaluation prompt specified only what was being checked. Each prompt was written for a single step and asked the system to do one thing.

Gemini Pro's adversarial function became more structured as the workflow evolved. Early in the project, Bob had used it informally—sending a draft and asking in general terms whether the argument held up. The responses were useful but uneven. Over time, Bob began specifying the adversarial task more precisely. He would tell Gemini to assume the perspective of a critic who believed the anti-ESG backlash was primarily performative—a political signaling exercise without serious policy consequences—and to identify the three arguments in the draft that such a critic would find most vulnerable, explaining for each what the critic would say and what evidence the critic would cite. The output of this exercise was not incorporated directly into the draft. It was used to identify the places where the argument needed to be strengthened or where the evidence needed to be more explicitly engaged. In one instance, Gemini identified that a central claim about the coordinated nature of anti-ESG

legislative activity rested almost entirely on circumstantial evidence about timing and similarity of language across state legislatures, with no direct evidence of the coordination mechanism itself. Bob spent two additional sessions developing that part of the chapter, adding evidence about the specific organizations that had produced and disseminated model legislation. The claim survived; its basis was substantially improved.

Creating Separate Rooms

The use of separate conversational environments—rooms, in Bob’s terminology—emerged from a related problem. He had noticed, in sessions that extended over several hours, that the quality of outputs tended to decline as the session lengthened. Early in a session, a system working on a section about the legal strategies used by state attorneys general to challenge ESG investing requirements would produce tightly argued, specifically grounded prose. Later in the same session, assigned a section on the role of think tanks in legitimizing anti-ESG arguments, the same system would produce prose that was structurally similar but analytically thinner—less grounded in specific examples, more reliant on general characterizations, with a tendency to restate what had already been said rather than developing new material.

The rooms were not simply different sessions. They were differentiated by function. Drafting happened in one room. Revision happened in a second, and only after the draft had been evaluated outside of any AI system. Bob would read the draft himself, make notes on what needed to change, and formulate specific revision tasks before opening a revision session. Evidence integration happened in a third room, using Perplexity, entirely separate from drafting and revision. Bob had learned early that asking a drafting system to also identify and integrate citations produced outputs in which the citations had been generated rather than found—plausible-looking references that, on inspection, did not exist or did not support the claims they were attached to. The verification function had to be separated from the generation function, handled by a different tool, in a different room, at a different stage.

Methodological reflection happened in a fourth room, and its structure was distinct from the others. This was not a room in which Bob prompted a system to analyze prior outputs or audit the workflow on his behalf. It was a room in which he prompted himself—using the act of formulating a prompt as a way of forcing precision about what he actually thought was happening. He would open a session and write out, as directly as he could, a description of a problem he had been encountering and what he believed its source was. The system would respond; Bob would read the response less for its content than for what the act of writing the description had clarified. In some sessions, the response was useful on its own terms—identifying an implication of the problem he had not noticed or suggesting a variant of his proposed solution that was more precisely targeted. In other sessions, the value was almost entirely in the writing: the discipline of composing a clear description of a process problem forced a level of analytical attention that simply thinking about the problem did not. The reflection room was, in this sense, a tool for making tacit knowledge temporarily explicit.

Enter Clean ChatGPT

The Clean ChatGPT instance came into its own at a specific juncture in the project. Bob had been working with Dorothy for several weeks on the chapter dealing with the political economy of ESG regulation, and Dorothy's accumulated context had become, in a subtle way, a constraint. When Bob asked her to assess whether the chapter's central argument was clearly stated, she confirmed that it was—but her assessment was grounded in her knowledge of what the argument was supposed to be, developed over weeks of prior conversation. What he needed was a response from a reader who did not already know what the argument was trying to do.

He sent the chapter to Clean ChatGPT with a prompt written by Dorothy that basically said, "Read this and tell me, in one paragraph, what argument you understand it to be making." Clean ChatGPT described an argument that was related to but meaningfully different from the one Bob had intended—less about the political economy of coordinated opposition and more about the legal vulnerabilities of ESG mandates. The chapter, read without prior context, was producing the wrong impression. Bob revised the introduction and the framing of the first two sections. When he sent the revised version to Clean ChatGPT with the same instruction, the response described the argument he had intended.

Locking

Locking emerged as a structural practice alongside decomposition and rooms. When a section reached the standard Bob had set for it—argument sound, prose right, evidence in place—he declared it locked. From that point, it could not be reopened except for a specific and narrow reason—a discovered factual error, a confirmed inconsistency with another locked section, or a structural change to the book that made a previously acceptable framing no longer viable. These exceptions were rare. The locking practice did not prevent necessary changes; it required that the reasons for changes meet a higher bar. In practice, most impulses to revise locked sections were recognized as impulses rather than necessities and were set aside.

Other Failures

Several failure modes proved more resistant to process solutions than others. Footnote management was the most consistently difficult. The book required a substantial apparatus of citations, and maintaining consistent footnote numbering across a long document was a task that no system managed reliably. The problem was structural; footnote numbers were global properties of the document, assigned sequentially across all chapters, but the document was being produced section by section, in separate rooms, by systems that had no persistent awareness of the document as a whole. Attempts to use AI systems to manage the renumbering produced their own errors. Footnote lists renumbered without updating in-text references, or sequences correct in one part of the document and wrong in another. Bob

eventually separated footnote management entirely from the drafting process, treating it as a distinct task to be handled only after the text of each chapter was fully locked, applied as a final layer that was never subsequently revised.

Diagram generation proved to be a category error embedded in the workflow design. The book required several visual representations of the anti-ESG network—organizational relationships, funding flows, legislative coordination structures—and early attempts to generate these within the AI workflow produced outputs that were technically responsive but practically unusable: text-based approximations of diagrams, or detailed specifications for visuals that still required a separate tool to produce. The insight was simple but consequential. Systems designed to produce text could not reliably produce publication-quality visuals regardless of how precisely the task was specified. The diagram work was routed outside the primary workflow entirely.

Context saturation manifested in a subtler and more pervasive way than footnote errors or diagram failures. It did not produce obviously wrong outputs. It produced outputs that were correct at a lower level of precision than the book required—prose accurate in its general characterizations but imprecise in its specific claims, arguments logically sound but analytically shallow, transitions that moved from one point to the next without the connective tissue that made the movement feel necessary rather than arbitrary. The tells accumulated across a section rather than appearing in any single paragraph.

One episode, midway through the drafting of the book’s central analytical chapter, illustrated the pattern clearly. Bob had been working in a single session for several hours on a section about the coordinated legislative strategy that had produced anti-ESG resolutions in over 20 state legislatures between 2021 and 2024. Early in the session, the outputs had been precise: specific organizations named, dates accurate, mechanisms described with the granularity the argument required. As the session continued, the outputs began to shift. Organizations were described in more general terms. Mechanisms were characterized at a higher level of abstraction. The section was still about what it was supposed to be about, but it had lost the specificity that made it analytically useful. Bob closed the session, opened a new room, reintroduced the outline and the key sources, and continued from the point at which the quality had begun to decline. The outputs from the new session recovered the earlier level of precision within the first exchange.

From Generation to Validation

By the later stages of the project, the workflow had shifted from generation to validation. New material was no longer being produced at scale. The task had become one of ensuring that what had already been written held together—that arguments were consistent across chapters, that evidence supported the claims it was meant to support, that the structure of the book reflected the structure of the analysis. Bob would take the text of two chapters—never more than two at a time—and submit them with specific analytical questions about

consistency. When an inconsistency was identified, he first assessed whether it was real or apparent. Apparent inconsistencies were common. The system's tendency to apply a uniform analytical lens to material that had been deliberately differentiated. Real inconsistencies required correction in the locked text, with explicit notation of what had been changed and why, followed by another consistency check before re-locking.

Perplexity played a larger role in this phase than it had in the drafting stages. With the argumentative structure in place, grounding specific claims in external sources had become more tractable. A claim about the total amount spent by a particular foundation on anti-ESG advocacy between 2018 and 2023, for example, required multiple searches for publicly available tax filings, for investigative reporting that had aggregated funding data, and for academic analyses that had synthesized public records. The results from different sources did not always agree. Bob had to assess the credibility and methodology of each source, determine which figure was most defensible, and revise the text accordingly—sometimes making the claim more precise, sometimes more hedged, depending on what the evidence actually warranted. This was not a task that could be delegated.

What Bob observed, as the project moved toward completion, was that the role of human judgment had not diminished as the workflow became more structured. If anything, it had intensified. Decisions about sequencing—which tool to use at which stage, in what order, with what specific instructions—required continuous attention. The sequence was not fixed; it had to be adjusted in response to what each step produced. A draft that came back well-formed could move immediately to evidence integration. A draft that came back with structural problems had to go through an additional revision cycle first.

The evaluation burden was particularly demanding. Reading and assessing every output—not just for surface correctness, but for analytical adequacy, tonal consistency, structural fit, and evidential grounding—required sustained concentration. The ratio of evaluation time to generation time was at least ten to one across the project as a whole. A section that took an agent a few minutes to draft might take an hour to evaluate properly. That hour was not passive. Bob would read the draft once for overall impression, then again with his outline in front of him, checking each paragraph against what it was supposed to accomplish. He would note, paragraph by paragraph, what was working, what was thin, what had drifted from the intended argument, and what was missing entirely. He would then assess whether the problems were addressable through targeted revision—in which case he would formulate specific revision tasks—or whether the section needed to be substantially redrafted, requiring a return to the outline first. By the end of that hour, he typically had a clearer view of the section's problems than any AI system had provided, because his evaluation was grounded in a precise understanding of what the section needed to do in the context of the specific chapter and the overall book. Only Bob had the full context of the book in mind. The agents were always focused on a particular task and this was necessary for them to do their work well.

Speed and control came into tension most acutely at the boundary of the locking decision. A section was not quite right—the prose acceptable but not as precise as the argument required, or the evidence present but not as tightly integrated as it could be—but further revision would require opening another room, reintroducing context, running another cycle. The temptation to lock and move forward was real. Bob developed a practice for these moments. He would write a brief note identifying the specific inadequacy, set the section aside, and return to it the following day. Most of the time, the note confirmed that the revision was necessary. Occasionally, reading the section fresh, what had seemed inadequate the previous day was in fact acceptable. The forced pause was itself a form of quality control.

The Challenge of Making Tacit Knowledge Explicit

Now, in the first week of April 2026, working with Dorothy to reconstruct the 22 conversations in order to write this case, Bob was confronting a problem he had not anticipated. The tacit knowledge embedded in the workflow was not easily recoverable from the record. The practices that had made the process work—the specific timing of the room transitions, the precise formulations that had succeeded where earlier versions had failed, the judgment calls about when an output was good enough to lock—had been made in the moment, without documentation, and could not be reliably reconstructed from the conversation logs alone. Dorothy’s reconstruction could describe what had been done. It could not reliably explain why each decision had been made at the specific moment it was made, because the reasoning had been tacit and the record did not contain it.

This was a different problem from the one the workflow had been designed to solve. The workflow had been designed to produce a book. It had succeeded. The problem now was whether the knowledge embedded in that workflow—the accumulated understanding of how to make a multi-agent AI system produce reliable, high-quality output on a complex analytical project—could be made explicit enough to be transferable. Whether another researcher, attempting to design a similar system for a different project, could take from this account not just a description of what had been done, but a set of principles or practices that would be genuinely useful in a different context with different tools and different material.

Bob had a sense of the answer. He had developed, over months of working closely with these tools, an understanding of what mattered and what did not—of where the leverage was and where it was not. He knew, for instance, that the specific tool assignments were less important than the principle of role separation; that the specific room structure was less important than the principle of functional differentiation; that the specific locking criteria were less important than the principle of building on fixed rather than provisional foundations. But knowing this and being able to state it in terms specific enough to be actionable, were different things. The principles, stated at sufficient generality to apply across contexts, were not wrong but were also not quite right—they were accurate descriptions of what had mattered, but they did not capture the judgment required to apply them in practice. And the judgment, as far as Bob could tell, was precisely what the record could not preserve.

The case that was now taking shape was itself a test of this problem. If the account of the production process could be written in a way that conveyed not just the structure of the workflow but the texture of the decisions within it—the specific failures, the specific responses, the specific moments at which the process had required more from Bob than the design had anticipated—then perhaps the tacit knowledge could be made sufficiently concrete to be of use. Whether that was achievable was not yet clear. The 22 conversations had been reconstructed, but reconstruction was not the same as transmission. What Bob had built was a process. What he was now trying to build was an account of a process that could function as a guide. The distance between those two things was the problem that remained—and it was a problem that no methodology, however carefully documented, could fully close, because the gap was not in the description but in the judgment that description alone could not convey.

Appendix A: The Anti-ESG Ecosystem — Table of Contents

The Anti-ESG Ecosystem: Implications for U.S. Climate Policy

The book maps the organizations, funding networks, legal strategies, and political mechanisms behind the U.S. backlash against ESG investing, and analyzes how that backlash constrains institutional visibility of climate-related financial risk. It is organized into four parts, twelve chapters, a prologue, and three appendices.

Prologue

Frames ESG not as a singular doctrine with a clear trajectory but as a loose umbrella of heterogeneous practices that became politically contestable. The prologue sets out the book's core approach: treating ESG conflict as an ecosystem characterized by interaction effects, feedback loops, and path dependence, rather than as a linear policy debate with winners and losers.

Part I — A Primer on ESG

Provides the conceptual orientation for the book. Rather than adjudicating whether ESG succeeded or failed, Part I establishes the analytical object: a modern risk-sensing infrastructure that became politically contestable, legally vulnerable, and behaviorally constrained.

Chapter 1 — What Is the Anti-ESG Ecosystem?

Introduces the book's central argument: the anti-ESG ecosystem is not simply a backlash against a set of corporate disclosures but an emergent pattern of mutually reinforcing actions that narrows the range of long-horizon risk that can be legitimately named and acted upon within American institutions. ESG itself is treated as a surface expression of deeper dynamics concerning institutional visibility of systemic risk.

Chapter 2 — Letters, Laws, and Lawsuits: How Pressure Enters the System

Traces the mechanism by which anti-ESG opposition moves from rhetorical positioning into institutional constraint. It shows how pressure enters organizations in sequential stages — letters, then legislation, then litigation — each lowering the cost of escalation and altering behavior before formal outcomes are ever reached. The chapter establishes that constraint can be imposed upstream of enforcement, making uncertainty itself an instrument of compliance.

Part II — The Analytical Foundation

Provides the analytical architecture that structures the rest of the book. The goal is not to catalog actors or predict outcomes but to discipline causal reasoning in a domain where evidence is partial, attribution is routinely misplaced, and single-domain explanations consistently mislead.

Chapter 3 — The Seven Edge System

Introduces the book’s core analytical framework: a seven-domain system comprising the narrative edge, the legal edge, the capital edge, the governance edge, the real economy edge, the macroeconomic edge, and the political edge. Each domain has distinct actors, logics, and feedback mechanisms, and the system’s behavior arises from their interaction rather than from any single domain’s dynamics.

Chapter 4 — Edge by Edge: How the System Actually Works

Works through each of the seven edges in operational detail, showing how pressure actually moves within and across domains. The chapter emphasizes that institutional pressure is more often procedural than coercive — it operates by shifting defaults, narrowing what feels defensible, and widening what feels exposed — rather than through direct commands or explicit enforcement.

Chapter 5 — What We Are Holding Constant

Defines the methodological constants that structure all subsequent chapters — the explicit assumptions held stable so that disagreement can be located at the level of analytical premises rather than displaced onto conclusions. These constants are not empirical claims to be validated but analytical commitments that make the scenario analysis in Parts III and IV possible and interpretable.

Part III — Baseline Conditions and Climate Stress

Establishes the baseline against which all political scenarios in Part IV are evaluated. It describes the present-state configuration of the anti-ESG ecosystem and then introduces climate stress to examine how the system responds without the variable of political change.

Chapter 6 — Institutional Hardening

Describes the present-state configuration of the ecosystem as one of “institutional hardening”: a pattern in which organizations across domains have converged on procedural defensiveness, constrained discretion, and path-dependent risk management — not through coordination or shared ideology, but through repeated rational adaptation to unresolved legal and political uncertainty.

Chapter 7 — Climate Shock Without Visibility

Constructs a scenario in which a sustained and physically manifest climate shock coexists with institutional under-response. The chapter shows how visible climate change can fail to produce commensurate institutional action when hardened organizations lack the discretion, language, and legitimized frameworks to translate physical signals into governance responses. Visibility in the physical world does not guarantee visibility in the institutional world.

Part IV — The Political Scenarios

Examines how the anti-ESG ecosystem processes a visible climate shock under four distinct configurations of post-2028 U.S. political control. Each scenario holds the physical shock constant and varies only the distribution of political authority, making visible how institutional response pathways differ — and how the core constraints of the system persist across all four configurations.

Chapter 8 — Unified Republican Control (2028)

Analyzes the system’s response to climate shock under unified Republican control of the federal government. Executive authority reinforces existing anti-ESG constraints, federal regulatory rollback reduces disclosure requirements, and institutional hardening deepens as organizations lose the cover of compliance obligations. The chapter shows how political authority can amplify system constraints already present rather than introducing fundamentally new ones.

Chapter 9 — Republican Presidency and Senate, Democratic House (2028)

Examines divided government as a machine for converting climate shock visibility into political conflict without producing institutional resolution. With one party controlling the executive and upper chamber and the other controlling the lower chamber, responsibility fragments while accountability intensifies. Neither side can act unilaterally; both can block; and both gain politically from demonstrating the other’s failure.

Chapter 10 — Democratic President and House, Republican Senate (2028)

Analyzes a configuration capable of significant executive action but incapable of legislative settlement. Democratic control of the presidency and House creates capacity for regulatory and executive initiatives on climate-related disclosure, but Republican Senate control blocks statutory consolidation. The chapter shows how institutional responses cluster around executive authority when legislative pathways are closed, and how that concentration creates both opportunity and vulnerability.

Chapter 11 — Unified Democratic Control (2028)

Examines the scenario with the greatest formal capacity for institutional response to climate risk. Unified Democratic control enables legislative action, regulatory expansion, and coordinated disclosure requirements. The chapter explores how the system’s structural constraints — federalism, legal uncertainty, capital discipline, and path dependence — limit the speed and durability of institutional change even when political authority is consolidated.

Chapter 12 — Divided Government and System Behavior

Synthesizes the findings across all four political scenarios. Rather than identifying a preferred political configuration, the chapter derives cross-cutting conclusions about how the system behaves under strain: which constraints are robust across all scenarios, where political authority creates genuine leverage, and what the scenarios collectively reveal about the

relationship between political change and institutional inertia. The chapter reframes the scenario analysis as a set of stress tests rather than forecasts.

Appendices

Appendix A — A Note on AI-Augmented System Observation

Addresses the problem of observability: the anti-ESG ecosystem cannot be directly observed in full because it operates across seven institutional domains simultaneously, each with partial data and its own logic. The appendix explains how AI-augmented analysis was used to extend the effective range of observation, enabling cross-domain pattern recognition that would be beyond the capacity of a single analyst working from conventional sources.

Appendix B — Methodological Note on Workflow and Collaboration

A practical account of how the book was produced: a structured multi-AI collaboration in which each system was assigned a defined role — drafting, adversarial challenge, verification, coordination — and the author retained responsibility for framing, integration, and all substantive judgments. The appendix documents the division of labor, the sequencing of tasks, and the governance principles that maintained quality and attribution across the workflow.

Appendix C — Reflections on Process

Offers interpretive observations on the production process described in Appendix B — not as a formal methodology but as a candid account of what the process revealed about the nature of the work itself. The appendix reflects on the limits of reconstruction, the tacit knowledge embedded in workflow decisions, and the question of what it means to retain authorial judgment in a system where generative AI contributes substantially to the text.