

After Asilomar: Walking Away from Anti-Biological Warfare into the Age of AI

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

November 16, 2025

This paper is written from the vantage point of a scientist who has lived through several distinct technological epochs: the early days of recombinant DNA, the long tunnel of anti-biological warfare work, and a late-life pivot into artificial intelligence. The through-line is not technique but risk: what it does to a human mind to work for decades at the edge of what can harm entire populations, and why, at seventy, one might finally decide to walk away. The narrative begins in the shadow of Asilomar, when “we could not put *Drosophila* in *E. coli*” was a live ethical boundary rather than an archival anecdote. It follows the move into classified anti-BW domains—bunkers without windows, pseudonyms layered on pseudonyms, and the slow accumulation of moral injury and PTSD. From there it turns toward the present: a world where engineered pathogens, climate-shifted vectors, and high-containment labs coexist with rapidly advancing AI systems that will increasingly mediate biological work. Rather than discussing operational details, this paper is an ethical memoir and a warning. It asks what lessons from the anti-BW era should inform AI and biosecurity today, and why personal boundaries—up to and including refusal—are now an essential part of scientific responsibility.

Keywords: anti-biological warfare, Asilomar, recombinant DNA, biosecurity ethics, dual-use research, PTSD, scientific responsibility, personal boundaries, artificial intelligence, logistics, national security, bunkers, intelligence community, whistleblowing, ethical memoir.

A collaboration with GPT-5.1-Thinking-Extended. CC4.0.

Outline

1. Framing the Box: Why This Story Is Being Told Now

- 1.1 The bunker with no windows: a literal and psychological setting
- 1.2 From anti-BW to AI: why the pivot matters
- 1.3 Scope and limits: what will and will not be discussed
- 1.4 Audience: young scientists, security professionals, and AI practitioners

2. Origins in Restraint: Berkeley, Asilomar, and Early Recombinant DNA

- 2.1 Learning to clone in an era of self-imposed caution
- 2.2 Asilomar and Jonathan King: the first serious conversation about limits
- 2.3 “We couldn’t put drosophila in *E. coli*”: what that boundary symbolized
- 2.4 How those early norms shaped expectations about science and responsibility

3. Into the Tunnel: Entering Anti-Biological Warfare Work

- 3.1 The call to defense: why a biophysicist steps into BW-adjacent domains
- 3.2 Living in compartments: clearances, need-to-know, and divided selves
- 3.3 Pseudonyms upon pseudonyms: survival by fragmentation
- 3.4 The slow burn of moral injury and operational fatigue

4. Life in the Bunker: Psychology, Culture, and Cost

- 4.1 Working without windows: isolation, secrecy, and distorted normality
- 4.2 Colleagues and archetypes: agents, analysts, digital skeptics, and “wet works”
- 4.3 Carrying threat models home: when the perimeter follows you
- 4.4 PTSD as an occupational hazard: how it appears, how it is hidden

5. The Breaking Point: Why I Walked Away at Seventy

- 5.1 The anthrax echo: being pulled back “for one more consult”
- 5.2 Climate ticks and hemorrhagic fevers: recognizing insanity in new clothing
- 5.3 Colleagues’ anger and accusations of betrayal
- 5.4 Choosing self-preservation: drawing a final ethical and psychological line

6. Lessons from Anti-BW for the AI Century

- 6.1 Dual-use déjà vu: why AI feels like recombinant DNA plus BW all over again
- 6.2 Secrecy versus transparency: what must never be hidden from public debate
- 6.3 Guardrails that work: Asilomar-style restraint versus performative ethics
- 6.4 “What we must never build”: the importance of negative design space in AI and biology

7. Boundaries as Ethics: Saying No in a Culture of Yes

- 7.1 The myth of the indispensable expert: why no one is obligated forever
- 7.2 How to decline: practical scripts for refusing dangerous work

- 7.3 Distinguishing prudence from paranoia in high-risk domains
- 7.4 The role of age and vocation: what changes when you reach seventy

8. From Bunkers to Black Boxes: Working with AI After Anti-BW

- 8.1 Why AI feels safer and why that feeling can be misleading
- 8.2 Using AI to reduce risk rather than amplify it
- 8.3 Logistics and alignment: rephrasing old safety instincts in new language
- 8.4 What an anti-BW veteran uniquely brings to AI ethics and governance

9. Mapping the Ghosts: Integrating Lives Lived Under Pseudonyms

- 9.1 The trail of forgotten names: what compartmentalized lives leave behind
- 9.2 Reclaiming continuity: telling one coherent story without breaching secrets
- 9.3 Spiritual and psychological dimensions of decompartmentalization
- 9.4 What to keep classified forever, and what to offer as public testimony

10. An Ethical Will to the Next Generation

- 10.1 To young biologists: what Asilomar still has to teach you
- 10.2 To AI researchers: how not to repeat the worst mistakes of anti-BW
- 10.3 To security professionals: why human limits matter as much as technical ones
- 10.4 The case for refusal, humility, and quiet courage in an age of extreme capability

11. Conclusion: Walking Away, Without Looking Away

- 11.1 Why leaving the field is not abandoning the fight
- 11.2 The enduring value of lived experience in shaping future guardrails
- 11.3 Bunkers, windows, and the choice to see clearly
- 11.4 Final reflection: what it means to grow old with dangerous knowledge

12. References

Collected works on Asilomar and recombinant DNA, biosecurity ethics, dual-use research, PTSD in security and defense communities, AI safety and alignment, and relevant autobiographical or historical sources

1. Framing the Box: Why This Story Is Being Told Now

This paper begins in a box with no windows. It is not a metaphor chosen for dramatic effect, but a description of a physical place and a mental posture. The box is a bunker, in the sense that everything important happens inside it and everything dangerous waits outside. It is the kind of space where phones are turned off or left elsewhere, where paper is scarce, where the air feels slightly thicker because the conversations it holds are heavy. The box is also what remains in the mind after years of working in anti-biological warfare and adjacent security domains: a compressed, guarded interior, where topics are sorted into what can be said, what must never be said, and what must be carefully translated into a different vocabulary.

Telling this story now is not an act of nostalgia or confession. It is a deliberate attempt to place a long career in biological science, anti-BW work, and late-life AI practice into a single continuous arc, and to ask what that arc might mean for the next generation. The key claim is simple: there are lessons from the anti-BW era that remain urgently relevant as artificial intelligence and modern biology intertwine, and some of those lessons can only be articulated by people who are willing to walk away from the operational details while still speaking clearly about risk.

The decision to leave anti-BW work at seventy is central. It is both a boundary and a lens. Boundary, because it marks the end of direct engagement with any operational program or technical line that could materially increase biological risk. Lens, because stepping away makes it possible to see the long sweep: from the ethical caution of early recombinant DNA work, through the intense compartmentalization of defense-related biology, to a present in which AI systems will be able to accelerate both constructive and destructive projects in ways that earlier generations could not have imagined.

This paper does not offer secrets. It offers perspective. It avoids agent names, protocols, and operational descriptions not because they are forgotten, but because remembering them is no longer the right kind of service. Instead, it focuses on the human experience of working with dangerous knowledge, on the institutional cultures that grow around such work, and on the internal and external limits that became necessary to remain human in the midst of it.

The intended audience is not limited to biosecurity specialists. Young molecular biologists who have only heard of Asilomar as a historical footnote, AI researchers who see biological modeling as just another application, and security professionals who are still early in their careers may all find in this narrative a set of questions they have not yet learned to ask. The goal is not to prescribe policy, but to provide an ethical memoir that can inform policy, culture, and personal choice. It is written from inside the bunker, but with the conscious intention of opening at least one window.

1.1 The bunker with no windows: a literal and psychological setting

The windowless room matters because it teaches a particular way of seeing. Time is told by clocks, not by sunlight. Weather is something that happens to other people. The only horizon is the edge of the table and the walls around it. That environment is not accidental. It is designed to control information: what comes in, what goes out, and who sits in the middle. Over years, it trains the mind to think in compartments, to accept that some things must be known and never spoken, and others must be spoken in a language that obscures more than it reveals.

Psychologically, the bunker becomes a template. Even when one leaves the physical space, the habits it instills follow. Conversations are scanned for hidden agendas. Questions are classified as friendly, probing, or operational. Stories are edited on the fly to remove details that might link back to programs, people, or capabilities that should stay in the dark. At its best, this posture protects genuinely sensitive information and prevents harm. At its worst, it becomes a reflexive secrecy that makes it difficult to share even harmless truth and to receive ordinary human support.

Writing from inside that template requires an explicit decision to distinguish between legitimate secrecy and unnecessary silence. It requires the author to say, in effect: this happened to me, this is what it did to my mind and body, and this is what I think it means for the future, without crossing lines that would recreate the very risks that originally justified the bunker.

1.2 From anti-BW to AI: why the pivot matters

The pivot from anti-biological warfare to artificial intelligence is not a random career shift. It is a relocation of ethical attention. For decades, the focus was on biological agents, delivery systems, and the ways in which state and non-state actors might exploit them. Today, many of the same questions are being asked about AI: who controls it, what they intend, what unintended consequences might emerge, and how to build guardrails that actually work.

From the vantage point of someone who has lived both worlds, AI feels disturbingly familiar. There is the same mixture of awe and dread, the same sense that new capabilities are arriving faster than our ethical and legal frameworks can absorb them, and the same temptation for institutions to hide the most worrying aspects behind classified doors. At the same time, AI offers something that anti-BW work rarely did: tools that can help surface patterns of risk, simulate scenarios, and democratize insight in ways that were previously impossible.

The pivot matters because it allows the lessons of one domain to be translated into another. Concepts like dual-use, negative design space (what must never be built), and the importance of early, self-imposed restraint are as relevant to AI as they were to recombinant DNA and biological weapons. By stepping out of the strictly biological arena, the author can speak more

freely about these patterns without being dragged back into operational details that should remain buried.

1.3 Scope and limits: what will and will not be discussed

The scope of this paper is broad in time but narrow in technical detail. It will cover:

- Early experiences in molecular biology during the period when Asilomar shaped the culture of recombinant DNA work.
- The psychological and ethical dynamics of anti-BW work: compartmentalization, pseudonyms, and the bunker mindset.
- The process and implications of walking away from that work at seventy.
- The parallels between anti-BW and AI, especially around risk, secrecy, and guardrails.
- The role of personal boundaries and refusal as legitimate ethical acts in high-risk fields.

It will not discuss:

- Specific agents, constructs, or operational programs.
- Technical methods that would materially lower the barrier to harmful biological work.
- Any classified or compartmented information.
- Names of individuals, organizations, or locations that are not already part of public historical record.

This is not a manual, a leak, or a veiled briefing. It is a reflective document meant to inform conscience and culture, not capability.

1.4 Audience: young scientists, security professionals, and AI practitioners

This paper is written with three overlapping audiences in mind. The first is young scientists, particularly in biology, who may be encountering powerful tools without the living memory of how earlier generations chose to restrain themselves. For them, the Asilomar moment and its aftermath are not just stories; they are precedents that can be adapted to new contexts.

The second audience is security professionals who work at the interface of science and state power. They understand the bunker, the need-to-know principle, and the tension between protection and overreach. For them, this narrative offers an insider's account of the personal costs of that work and a plea to consider human limits as seriously as technical ones.

The third audience is AI practitioners and theorists. Many of them now find themselves in conversations about alignment, dual-use, and catastrophic risk that echo earlier debates in

biology and weapons work. For them, this paper offers a bridge: an attempt to translate the lived experience of anti-BW work into a vocabulary that can be useful for AI governance, research culture, and personal ethical stance.

If there is a fourth audience, it is anyone who has ever worked under pseudonyms, in windowless rooms, and now finds themselves wondering how to grow old with dangerous knowledge without being consumed by it. For that reader, this paper is a companion piece, a way of saying: you are not alone, and it is possible to walk away without entirely looking away.

2. Origins in Restraint: Berkeley, Asilomar, and Early Recombinant DNA

The story begins in an era when restraint was not merely an afterthought but a conscious, collective project. At Berkeley and other centers of molecular biology, cloning was new enough to feel miraculous and frightening at the same time. The tools were crude by today's standards, but the conceptual leap was immense: to cut and paste genetic material, to move bits of one organism into another, to take the code of life and begin editing it in systematic ways.

Asilomar did not appear out of nowhere. It grew from a climate in which leading scientists recognized that their technical capacities were outrunning their institutional and ethical frameworks. The conference represented a rare moment when those in the vanguard called a halt on themselves, not because governments forced them to, but because they could see, dimly but clearly enough, the potential for harm.

For a young scientist learning to clone in that environment, the message was clear: power comes with limits, and those limits must be taken seriously. The specific rule of “we cannot put *Drosophila* in *E. coli*” may sound almost quaint now, given what is routinely done in modern genetic engineering labs, but it symbolized something larger. It said that there are combinations of host and donor that cross an invisible line from exploration into unacceptable risk, and that line can and should be drawn by the practitioners themselves.

This section explores what it meant to come of age in a scientific culture shaped by that ethos, and how those early experiences set expectations that would later be challenged and, in some cases, betrayed.

2.1 Learning to clone in an era of self-imposed caution

To learn cloning in the early days was to inhabit a paradox. On one hand, there was an almost childlike fascination with the mechanics: restriction enzymes cutting DNA, ligases stitching fragments together, plasmids carrying new sequences into bacterial hosts. On the other, there was an ever-present sense that these manipulations could have consequences beyond the bench.

Training did not consist solely of pipettes and protocols. It also involved briefings on what was allowed, what was forbidden, and why. Graduate students and postdocs were told not only how to make constructs, but which constructs must not be made. Discussions about experimental design were interwoven with discussions about risk categories, containment levels, and the rationale for constraints that sometimes felt conservative or even arbitrary.

For many in that generation, these constraints were not resented but respected. They became part of the craft identity: to be a good molecular biologist was not only to be clever in design and careful at the bench, but also to be mindful of possible harm, even when the probability seemed low. This habit of mind would later become both a strength and a source of friction when the external pressures of national security, funding, and geopolitical competition began to pull in a different direction.

2.2 Asilomar and Jonathan King: the first serious conversation about limits

Asilomar was more than a meeting; it was a symbolic act. Scientists gathered, argued, and produced a set of guidelines that, for a time, structured the entire field. Central to that moment were voices warning that recombinant DNA research, if pursued without constraint, could create new pathogens, new routes for disease, and new tools for weaponization.

Figures like Jonathan King articulated concerns that went beyond immediate biosafety. They pointed to the way military and intelligence interests might co-opt biological research, the dangers of normalizing dual-use work, and the risk that once the taboo against certain kinds of experiments was broken, it would be difficult to rebuild. Their warnings were not universally accepted, but they were taken seriously enough to influence both policy and culture.

For someone who later entered anti-BW work, these early messages carried a particular irony. The very caution that Asilomar embodied would, in time, be selectively bypassed under the banner of defense, and some of the fears voiced in those early debates would be realized in ways that blurred the line between protection and provocation.

2.3 “We couldn’t put drosophila in *E. coli*”: what that boundary symbolized

The prohibition against moving certain kinds of eukaryotic DNA into bacterial hosts was not primarily about drosophila as such. It was about setting a visible, enforceable boundary that signaled seriousness. The example lodged in memory because it was concrete enough to guide decisions at the bench. If one wondered whether a proposed construct crossed the line, one could ask, half-jokingly but with real intent, whether it belonged in the category of “drosophila in *E. coli*.”

This kind of rule did several things at once. It simplified complex risk assessments into manageable heuristics. It reminded practitioners that they were operating under a collectively

negotiated regime of restraint. And it served as a constant reminder that the field existed under a sort of social license, one that could be revoked if accidents or abuses eroded trust.

Looking back from the vantage point of modern synthetic biology, that boundary may appear narrow or outdated. But as an ethical artifact, it remains instructive. It shows that scientists can, at least for a time, agree to limit their own power, and that such self-limitation can be both practically effective and culturally formative.

2.4 How those early norms shaped expectations about science and responsibility

The Asilomar-informed norms of early recombinant DNA work created certain expectations about how science and responsibility should relate. They suggested that:

- Scientists have both the capacity and the duty to forecast possible harms arising from their work.
- Collective deliberation and consensus-building are legitimate tools for setting limits, not obstacles to innovation.
- Self-imposed moratoria, even temporary ones, can be signs of maturity rather than weakness.
- The legitimacy of a field depends in part on its willingness to take risk seriously and to act visibly on that concern.

For those who absorbed these expectations at the beginning of their careers, later developments would be disorienting. When anti-BW work demanded secrecy rather than open deliberation, when dual-use considerations were handled behind closed doors rather than in public guidelines, and when restraint seemed to give way to a kind of technological bravado, it felt less like natural progress and more like a departure from an earlier, more careful path.

These early experiences thus serve as a baseline. The rest of the story can be read as a series of departures from, attempts to preserve, and eventual partial returns to the Asilomar ethos, across different domains and decades.

3. Into the Tunnel: Entering Anti-Biological Warfare Work

The move from open recombinant DNA work into anti-biological warfare was not a single decision, but a series of nudges and invitations. It often begins with a phone call, a hallway conversation, or a “special” seminar that is not widely advertised. The language is careful. The work is framed as defensive, as necessary, as a contribution to national and global security. For a scientist who has internalized the Asilomar ethic of responsibility, the logic can be compelling. If

dangerous capabilities exist, someone must understand them well enough to anticipate and counter them. If hostile actors are not bound by restraints, perhaps the only responsible act is to ensure that those who are bound by law and ethics are not helpless.

The entrance to the tunnel is well lit. There are clearances to be obtained, briefings to attend, and acronyms to learn. There is a sense of importance, of being trusted with matters that most citizens will never know about. There is, often, a genuine desire to prevent harm. What is less visible at the start is how long the tunnel is, how few exits there are, and how easily the line between understanding a threat and recreating it can blur.

For someone already accustomed to thinking about dual-use in biology, the anti-BW world presents itself as a natural extension of earlier worries. Instead of trying to avoid dangerous constructs, the task becomes to understand the threats others might pose, and to build surveillance, detection, and countermeasures. The problem is not that this framing is false; it is that it can slowly shift the center of gravity from prevention to capability, from “what must never be built” to “what might be built elsewhere, and therefore we must explore it too.”

3.1 The call to defense: why a biophysicist steps into BW-adjacent domains

Biophysicists are trained to think in terms of mechanisms, models, and quantitative predictions. Their skills translate naturally into questions like how an agent moves through a system, how it interacts with targets, and how physical constraints shape possible scenarios. When security agencies look for scientific partners, they often seek exactly this kind of mind.

For the scientist, the first appeal is moral. The work is described as defensive: modeling threats, understanding vulnerabilities, designing detection methods, and informing policy. The specter of biological attack, whether from state programs or non-state actors, is held up as a justification for research that would otherwise remain unthinkable. The implicit proposition is that refusing to engage is a form of irresponsibility, a withdrawal from the duty to protect.

There is also a second appeal, quieter but real: the desire to see whether one’s skills can operate at the highest stakes. Academic work may shape understanding; anti-BW work can, in principle, shape outcomes in conflicts and crises. For a scientist who has spent years arguing about the ethics of recombinant DNA, the chance to make a concrete difference in preventing catastrophic misuse can be hard to resist.

3.2 Living in compartments: clearances, need-to-know, and divided selves

Once inside, the logic of compartmentalization takes over. Information is divided into compartments and access is granted on a need-to-know basis. This structure is meant to protect sensitive material from unauthorized disclosure and to reduce the risk that any single individual could compromise entire programs.

For the person working within such a system, compartmentalization also divides the self. There is the part that lives in public, with a curriculum vitae, teaching history, and visible publications. There is the part that lives in secure spaces, working under different project names, reporting through separate chains of command, and often interacting with colleagues who are known only by first names or pseudonyms. There may be additional layers associated with specific clearances or collaborations.

Maintaining these divisions requires constant mental bookkeeping. One learns to answer questions selectively, to redirect conversations away from forbidden topics, and to carry multiple narratives of one's own life without allowing them to intersect. Over time, this becomes habitual. What begins as a necessary discipline for protecting information can harden into a general reluctance to speak freely about anything that feels important.

3.3 Pseudonyms upon pseudonyms: survival by fragmentation

In some settings, pseudonyms are formal tools: code names for projects, operations, or individuals. In others, they are informal masks used in publications, online forums, or side channels where sensitive topics are discussed without linking directly back to a real-world identity. For someone involved in anti-BW work, the use of multiple identities can become a default pattern.

Each pseudonym carries a fragment of the person's knowledge and experience. One might be associated with a particular research line, another with a set of internal reports, another with cautious interventions in public debates. Over time, the number of names grows. Some are used only once or twice, then abandoned. Others become long-running personas that accumulate their own histories and relationships.

Fragmentation of identity can be protective, but it also has costs. It makes it difficult to feel fully seen or understood by anyone outside the compartmented world. It complicates ordinary friendship and family life, as large parts of one's daily work cannot be shared. It can even erode memory, as the effort to keep track of what was said under which name leads to a kind of deliberate forgetting. In later years, looking back, one may find it difficult to reconstruct a coherent story of one's own life.

3.4 The slow burn of moral injury and operational fatigue

Working in anti-BW spaces involves repeated exposure to worst-case scenarios. Even when one's role is purely defensive, the mind must explore failure modes, imagine deliberate harm, and contemplate the vulnerability of civilian populations. Over time, this can produce a cumulative moral and psychological burden often referred to as moral injury.

Moral injury is not simply disagreement with policy or discomfort with particular decisions. It arises when one's core values are confronted by actions or possibilities that feel fundamentally wrong, and when one feels complicit, even indirectly, in a system that normalizes those possibilities. It can coexist with a continued belief that the work is necessary, which makes it harder to recognize and address.

Operational fatigue adds another layer. High-stakes work conducted under conditions of secrecy and constant alertness drains energy in ways that ordinary academic or industrial research rarely does. Sleep patterns are disrupted. Relationships strain under the weight of unspoken stress. Over years and decades, these effects accumulate. The point at which the burden becomes too great is rarely marked by a single dramatic event. It is reached instead through a slow, uneven realization that the costs now exceed the benefits, both to the individual and perhaps to the broader ethical landscape.

4. Life in the Bunker: Psychology, Culture, and Cost

To describe life in the bunker is to describe a particular combination of physical, social, and psychological conditions. Physically, the spaces are controlled: limited access, surveillance, enforced protocols for entry and exit. Socially, they bring together people from different institutions, disciplines, and agencies, each carrying their own cultures and agendas. Psychologically, they foster both intense camaraderie and deep isolation.

In such environments, the usual boundaries between scientific, military, and intelligence cultures blur. Conversations about experimental design may be interspersed with discussions about geopolitics, law, and covert operations. Ethical questions may be raised and then shelved for lack of time or because "policy has already decided." Humor tends to be dark, a release valve for tensions that can rarely be acknowledged plainly.

The cost of living in this mixture is not paid in a single currency. It comes in the form of chronic stress, ethical compromise, sleep disturbances, distrust of institutions, and, eventually, symptoms that can be recognized as post-traumatic stress. These costs are often minimized in the moment, both by individuals who want to believe they are coping and by organizations that need people to remain functional. They become fully visible only later, when one attempts to step back.

4.1 Working without windows: isolation, secrecy, and distorted normality

The absence of windows in a physical bunker is a simple security measure. It reduces lines of sight into the space and limits physical vulnerabilities. As a psychological fact, however, it also signals that what happens inside is separate from ordinary life. Daylight, weather, and the

rhythms of the outside world are replaced by fluorescent lights, regulated temperatures, and the artificial cycles of shift work and briefing schedules.

Secrecy reinforces this separation. People inside the bunker know things that others do not. They are not permitted to share those things with family or friends, and sometimes not even with colleagues in adjacent compartments. The result is a form of isolation that is not resolved by going home at night. A portion of one's mental life remains behind closed doors, inaccessible even in private reflection, because it has been trained to be inaccessible.

Over time, this environment can distort one's sense of normality. Threats that are statistically unlikely but operationally vivid may come to feel ever-present. Policies that would seem extreme or unacceptable in public debate can appear reasonable within the acontextual frame of briefings and simulations. Conversely, ordinary concerns about community life, local politics, or personal well-being can come to seem trivial by comparison. The world is divided into the urgent and the unimportant, and most of what ordinary people care about falls into the latter category.

4.2 Colleagues and archetypes: agents, analysts, digital skeptics, and “wet works”

The bunker brings together a cast of recurring archetypes. There are agents whose job is to operate in the field, analysts who digest data and produce assessments, scientists who provide technical expertise, and digital specialists who manage systems and signals. Each group has its own subculture, stereotypes about the others, and characteristic attitudes towards risk and ethics.

For example, some agents may present themselves as pragmatists, more concerned with outcomes than with abstract moral questions. Analysts may cultivate skepticism, always looking for alternative explanations and hidden assumptions. Digital specialists often view traditional agencies as clumsy and behind the times, dismissing their tradecraft as obsolete. Scientists may oscillate between fascination with technical problems and discomfort with the purposes to which their work may be put.

Within this mix, there are also attitudes towards “wet works,” a euphemism for direct physical operations that may involve lethal force. Many in the digital and analytical communities despise such activities, seeing them as crude, messy, and ethically indefensible. Others accept them as unfortunate necessities. Still others are attracted to their drama. These differences can lead to sharp internal divisions, even as the overall system continues to function.

For the scientist in this environment, it is easy to become a bridge figure, understood imperfectly by all sides, and therefore slightly alienated from each. One is expected to be technically precise, operationally useful, and ethically resilient. When those expectations conflict, there is rarely a clear path forward.

4.3 Carrying threat models home: when the perimeter follows you

Threat models do not stay in the bunker when a person goes home. The mind that has spent hours imagining how an adversary might misuse a biological tool, exploit a vulnerability, or infiltrate an institution does not easily turn that imagination off. Ordinary public spaces can start to look like potential targets. Everyday illnesses can be scrutinized for signs that they are not so ordinary. Travel routes, social gatherings, and even casual conversations may be scanned for patterns that resemble briefed scenarios.

This hypervigilance is adaptive in the short term for certain roles, but harmful in the long term for personal life. It can lead to withdrawal from community, increased suspicion of strangers, and difficulty trusting institutions that do not operate with similar threat awareness. It can also create friction in close relationships, as partners and family members may interpret the resulting caution as paranoia or emotional distance.

In some cases, the burden can be mitigated by peer support, counseling, or institutional programs that acknowledge the psychological toll of such work. In many others, however, it is simply carried in silence, adding to the layers of unsaid experience that define the lives of those who have spent extended time in the bunker.

4.4 PTSD as an occupational hazard: how it appears, how it is hidden

Post-traumatic stress disorder is often associated with acute events: explosions, attacks, or personally experienced violence. In anti-BW and related work, PTSD can arise more subtly. It may come from chronic exposure to catastrophic scenarios, repeated encounters with evidence of atrocities, or the cumulative effect of moral injury when one feels complicit in systems that run counter to deeply held values.

Symptoms can include intrusive thoughts, nightmares, emotional numbing, irritability, difficulty concentrating, and a persistent sense of threat. Because the triggers are often tied to classified material or to experiences that cannot be easily shared, those affected may find it hard to seek help. Even when they do, they may minimize the severity of their distress, comparing themselves unfavorably to colleagues who have experienced more visible forms of trauma.

Institutions sometimes recognize these issues and offer support. However, they also have an interest in maintaining operational capacity and may unconsciously valorize stoicism. Individuals may fear that acknowledging their struggles will jeopardize their clearance, their assignments, or their reputation. The result is an environment in which an occupational hazard is both widespread and systematically underreported.

5. The Breaking Point: Why I Walked Away at Seventy

No one announces in advance that a particular assignment, briefing, or consultation will be the one that tips the balance. The breaking point only becomes visible in retrospect, when one realizes that a line has been crossed internally. For someone who has spent decades in anti-BW work, the final straw may not be the most dramatic scenario, but the most repetitive one: yet another discussion of anthrax, yet another proposal to “proactively” explore a hypothetical threat, yet another instance of hearing enthusiastic plans from people who do not seem to carry the same weight of history.

In this case, the realization came in stages. First, being pulled back into anthrax-related conversations after having stepped away. Then, being presented with new ideas about climate-driven vectors and hemorrhagic viruses, framed as innovative, forward-looking, and necessary. The content was not shocking in itself. What made it unbearable was the sense that the field had not learned enough from earlier eras, that the language of defense was again being stretched to justify work that moved closer to offense in all but name.

At the same time, there was growing internal resistance to the role of perpetual expert. Colleagues expected continued participation. Some were angry at the idea of withdrawal, framing it as abandonment of the cause or dereliction of duty. Against that pressure, a different conviction emerged: that self-preservation and ethical clarity now required refusal.

5.1 The anthrax echo: being pulled back “for one more consult”

Anthrax is emblematic in anti-BW work. It stands for a certain class of threats: serious, well understood at a technical level, and repeatedly revisited in planning exercises. Even years after stepping away from operational roles, a person with relevant expertise may be asked for “just one more consultation,” “one more review,” or “one more sanity check.”

At first, such requests can feel modest and manageable. They tap into muscle memory and seem to ask little more than an informed opinion. Over time, however, each re-engagement pulls the mind back into familiar grooves: thinking at scale about infection, mortality, spread, and countermeasures. The emotional and cognitive patterns associated with these topics reassert themselves, often bringing with them memories and associations from earlier, more intense periods of work.

Recognizing the anthrax echo as a trigger is part of recognizing one’s own limits. It marks the realization that even minor engagements are not harmless; they reopen channels that had begun to close. The decision to decline further involvement is not about withholding useful knowledge out of spite or fatigue, but about acknowledging that indefinite reactivation is neither sustainable nor ethically necessary.

5.2 Climate ticks and hemorrhagic fevers: recognizing insanity in new clothing

The language of “proactive construction” of scenarios involving climate-shifted vectors and hemorrhagic pathogens can sound responsible on the surface. It suggests an attempt to stay ahead of emerging threats, to understand what might happen as environmental changes alter the distribution of diseases and their carriers. Taken in isolation, such studies may be justified and even valuable.

What raises alarm is the way this language can be used to normalize more aggressive forms of exploration. Proposals may move from modeling natural dynamics to suggesting experimental constructs that would never arise in nature, all under the banner of preparedness. The presence of high-containment laboratories and advanced genetic tools makes such work technically feasible, while the rhetoric of climate change and global health can provide a convenient cover of urgency.

For someone who remembers earlier warnings about dual-use, this pattern is all too familiar. It represents a shift from studying what might happen on its own to imagining and building what would otherwise never exist. The fact that such proposals are made confidently, sometimes by younger scientists who speak in upbeat tones about innovation, can make them feel all the more disturbing. The recognition of insanity here is not a personal judgment on particular individuals, but a diagnosis of a cultural drift towards treating dangerous constructs as acceptable research objects.

5.3 Colleagues’ anger and accusations of betrayal

Stepping away from anti-BW work does not happen in a vacuum. Colleagues who remain in the field may react with confusion, disappointment, or anger. They may feel that someone they relied on has abandoned them. They may interpret refusal as a moral judgment on their continued engagement. They may also fear that losing experienced voices will leave the field more vulnerable to errors or capture by less cautious actors.

These reactions can be painful, especially when they come from people who have shared long periods of intense work. They can trigger guilt and self-doubt. Is walking away selfish? Does one have an obligation to continue as long as one is capable? Are others now exposed to risks one could have helped mitigate?

The answer that eventually emerges is that responsibility does not require indefinite service. There is a difference between temporarily stepping back and making a reasoned decision that continuing to participate would do more harm than good, both to oneself and potentially to the integrity of the field. Accepting colleagues’ anger without internalizing their accusations is part of the price of claiming that boundary.

5.4 Choosing self-preservation: drawing a final ethical and psychological line

At seventy, the calculus changes. The remaining years are finite, and the cumulative effects of decades of high-stress work are impossible to ignore. Self-preservation is not simply about personal comfort; it is about maintaining enough health and clarity to contribute constructively in other ways, including in domains like AI where one's experience can help shape emerging guardrails.

Drawing a final line means stating plainly, to oneself and to others, that certain topics, roles, and environments are now off limits. It means declining invitations, no matter how "small" they seem. It means accepting that some problems will have to be addressed by others, and that one's own contribution will take different forms.

This choice is not an act of denial. One does not forget what one knows. The bunker remains in memory. The threats remain real. What changes is who bears the direct weight of that knowledge in operational settings. Walking away, in this sense, is not looking away. It is a decision to shift the mode of engagement from active participation to reflective testimony and to focus on preventing the repetition of past mistakes in new technological contexts.

6. Lessons from Anti-BW for the AI Century

The parallels between anti-biological warfare and artificial intelligence are not exact, but they are close enough to be instructive. Both domains involve powerful technologies whose full implications are not yet understood. Both sit at the intersection of civilian and military interests. Both attract significant funding from states and corporations. Both produce capabilities that can be used for great benefit or great harm.

From the vantage point of someone who has lived through the evolution of anti-BW work, certain lessons stand out as especially relevant to AI. They are not technical recipes, but patterns of institutional behavior and ethical drift that tend to recur when new capabilities emerge under conditions of geopolitical competition and incomplete governance.

These lessons can be grouped around several themes: the importance of recognizing dual-use realities early, the dangers of excessive secrecy, the value and limits of self-imposed guardrails, and the need to define a negative design space of things that must not be built, even if they are technically possible.

6.1 Dual-use déjà vu: why AI feels like recombinant DNA plus BW all over again

In the early days of recombinant DNA, the dual-use nature of the tools was obvious. They could be used to produce medicines, vaccines, and basic knowledge of life processes. They could also,

in principle, be used to create novel pathogens or enhance existing ones. Anti-BW work pushed this dual-use reality to the foreground, sometimes in ways that risked normalizing exploration of dangerous possibilities.

AI now sits in a similar position. Systems capable of modeling biological processes can accelerate drug discovery, vaccine design, and understanding of disease. The same systems can, at least in principle, lower the barrier to designing harmful agents or optimizing delivery mechanisms. More broadly, AI tools that enhance surveillance, decision-making, and automation can be used to strengthen public health and security, or to enable new forms of coercion and control.

Recognizing dual-use nature is not enough. The lesson from anti-BW is that once a technology is framed primarily in terms of defending against hypothetical adversaries, it can become easier to justify exploring dangerous options under the label of preparedness. AI discussions already show tendencies in this direction, with talk of “red-teaming” models against biological misuse and building internal capabilities “so that we know what others could do.” Without clear limits, such efforts can slide into the same pattern of expanding capability that characterized parts of anti-BW work.

6.2 Secrecy versus transparency: what must never be hidden from public debate

Security work requires some degree of secrecy. However, anti-BW experience shows that when too much of the debate over what should or should not be done is carried out behind closed doors, important ethical and political questions remain unexamined. Decisions about acceptable risk, lines not to be crossed, and trade-offs between capability and safety are not purely technical. They are societal choices that require democratic input.

In AI, there is already a push towards classified or proprietary programs, especially in national security and corporate contexts. Some of this is unavoidable. What must be resisted is the tendency to move entire categories of questions into hidden spaces. For example, whether it is acceptable to use AI to automate certain forms of targeting, surveillance, or manipulation should not be decided solely by defense agencies or product teams. Anti-BW teaches that leaving such decisions to small, self-selected groups can lead to outcomes that would not survive public scrutiny.

At the same time, transparency must be calibrated. Just as detailed publication of certain biological methods can be dangerous, full public release of certain AI models or techniques may pose real risks. The challenge is to find governance mechanisms that allow open discussion of principles and boundaries without enabling misuse. Anti-BW history, including both successes and failures in this regard, offers a reservoir of practical experience for AI governance designers.

6.3 Guardrails that work: Asilomar-style restraint versus performative ethics

Asilomar represented a form of self-imposed restraint that was substantive. Concrete limits were set, experiments were deferred, and containment levels were raised. Over time, however, some guardrail efforts in biology and anti-BW became more performative: statements of principle without corresponding changes in practice, review committees that rarely said no, and guidelines that provided cover rather than constraint.

AI is now entering its own guardrail phase, with companies and consortia issuing principles, safety frameworks, and alignment statements. The anti-BW lesson is that the value of such efforts depends on whether they lead to actual decisions not to build certain systems or deploy certain capabilities. If every proposal eventually finds a way through the process, the guardrail has failed.

Substantive restraint may involve moratoria on particular applications, refusal to work with certain clients, or limits on deployment scope and scale. It may also require building institutional cultures in which saying no is not career-ending. Anti-BW veterans can testify that without such cultures, formal mechanisms of review and oversight tend to become mere rituals.

6.4 “What we must never build”: the importance of negative design space in AI and biology

One of the most important concepts carried over from anti-BW work is the idea of a negative design space: a set of things that should remain unbuilt. In biology, this includes certain classes of constructs, experiments that bring together dangerous traits in novel ways, and operational programs that move from defense into offense. The content of this space is debated, but the category itself is crucial.

In AI, defining a negative design space is equally important. It might include systems designed explicitly for autonomous killing, large-scale behavioral manipulation without consent, or tools that make certain forms of biological or cyber attack trivially easy to execute. It might also encompass architectures that deliberately aim to evade human control or to maximize influence regardless of context.

The lesson from anti-BW is that if the only constraints on technology are reactive—responding after something has been built and perhaps misused—it will always be difficult to regain control. Proactive definition of what must not be pursued, backed by real enforcement and cultural buy-in, is one of the few durable guardrails available. Those who have lived through previous cycles of dual-use expansion are well positioned to argue for such negative spaces and to help articulate their boundaries.

7. Boundaries as Ethics: Saying No in a Culture of Yes

In high-risk scientific and security domains, the ability to say yes is routinely celebrated. Yes to new projects, new capabilities, new collaborations. Yes to calls from agencies and funders. Yes to questions framed as “if we do not, someone else will.” Against that cultural background, refusal can look like weakness, obstruction, or betrayal. Yet the experience of anti-biological warfare work suggests that boundaries are not simply personal preferences; they are ethical acts.

To say no in such a context is to assert that one’s responsibility includes limits on what one will help bring into being. It is to recognize that professional competence does not erase human finitude, and that a life spent constantly near the edge of catastrophe has its own costs and distortions. Boundaries are also a form of truth-telling. They make visible tensions that institutions would prefer to keep abstract, forcing the question of whether certain lines really should be crossed.

This section explores how the myth of the indispensable expert can trap individuals in roles that are no longer healthy or wise, how to decline requests without unnecessary drama, how to distinguish prudence from paranoia, and how age and vocation alter the meaning of continued engagement.

7.1 The myth of the indispensable expert: why no one is obligated forever

Institutions often act as though certain individuals are irreplaceable. They develop narratives in which a particular scientist, analyst, or operator is uniquely positioned to address specific threats. The individual may internalize this narrative, feeling that stepping away would place others at unacceptable risk. It is a powerful story, especially for those who entered the field out of a sense of duty.

The myth of indispensability is dangerous because it obscures institutional responsibility. No system should be built that depends on one person’s continued willingness to serve. When a program’s safety or effectiveness rests on the presence of a few experienced voices, the real problem is not the risk of their departure, but the failure to cultivate broader competence and robust oversight.

Recognizing this myth allows individuals to reassess their obligations. It becomes possible to see that others can and will step into roles, perhaps bringing fresh perspectives. It also highlights that the burden of correcting dangerous trends should not fall solely on those who have already carried it for decades. Stepping away is then reframed not as abandonment, but as a necessary transition that forces institutions to confront their overreliance on particular people.

7.2 How to decline: practical scripts for refusing dangerous work

Refusal does not always need grand gestures. Often, it begins with simple, clear statements. For example:

- “I am no longer available for work in that domain.”
- “I have decided not to participate in projects involving that class of agents or scenarios.”
- “I can offer general ethical perspectives, but I will not review or advise on operational details.”

These scripts avoid justifying the refusal in terms that invite argument. They state a boundary as a fact, not as a debate. If pressed, one can add that long experience has led to the conclusion that further involvement would not be wise. One can acknowledge that this may be disappointing while insisting that it must be respected.

In some cases, refusal may be coupled with constructive suggestions: recommending other experts who are better placed to engage, pointing to existing literature, or encouraging broader ethical review. The key is to resist being drawn into discussions that treat the boundary as negotiable. Once the decision is made, the purpose of conversation is to communicate it, not to reopen it.

7.3 Distinguishing prudence from paranoia in high-risk domains

In worlds where real threats exist and are actively studied, distinguishing prudent caution from paranoia is not straightforward. Both involve heightened awareness of risk and a willingness to act on that awareness. The difference lies partly in proportionality and partly in flexibility.

Prudence responds to specific patterns and evidence. It adjusts behavior in ways that are calibrated to actual exposure and context. It allows for revision when circumstances change or when new information emerges. Paranoia, by contrast, generalizes from limited experiences to broad, fixed beliefs, treats all ambiguous signals as confirmation, and resists correction.

For someone with a long history in anti-BW work, it can be easy to doubt one’s own judgment. Habits of scanning for hidden threats can feel excessive when measured against ordinary life. Yet some precautions are objectively reasonable: limiting disclosure of sensitive past work, avoiding certain locations or relationships, maintaining digital and physical hygiene at levels appropriate to one’s visibility and history.

The task, then, is not to abandon caution, but to subject it to periodic honest review. Conversations with trusted peers who understand both the real risks and the psychological toll can help. So can explicit recognition that personal safety and mental health are themselves legitimate priorities, not signs of weakness.

7.4 The role of age and vocation: what changes when you reach seventy

Age changes the terms of engagement. At seventy, time is no longer an abstract concept but a concrete constraint. The remaining years are limited, and the cumulative impact of past stressors is undeniable. Physical resilience may diminish; psychological burdens may weigh more heavily. Continuing to operate in high-intensity environments carries risks that were easier to absorb at forty.

Vocation also shifts. Early in a career, the primary focus may be on learning, contributing, and building. Later, it often turns toward mentoring, reflecting, and distilling. For someone who has spent decades in anti-BW and related work, this means that the most valuable contribution may no longer be detailed technical engagement, but the articulation of lessons learned, warnings, and guidance for those coming after.

Recognizing this shift allows a different kind of freedom. It becomes possible to step back from operational roles without feeling that one has ceased to serve. Instead, one can see the act of walking away itself as part of a vocation: modeling that it is permissible to leave, that lifelong service in dangerous domains is not a requirement for integrity, and that there is dignity in choosing how one's remaining energy will be spent.

8. From Bunkers to Black Boxes: Working with AI After Anti-BW

Transitioning from anti-BW work to AI is not as abrupt as it might seem. Both domains involve systems whose internal workings are not fully transparent to most people, whose outputs can have far-reaching consequences, and whose governance is contested. The bunker and the AI model are both black boxes of a kind: places where complex processes occur out of sight, and from which emerge decisions, predictions, or capabilities that others must live with.

For someone with a background in anti-BW, AI can feel, at least initially, like a safer space. The objects of study are models and algorithms rather than pathogens and human bodies. The risks, while serious, are often mediated through code and infrastructure rather than through direct biological harm. At the same time, the potential for misuse and unintended consequences is substantial, and the patterns of institutional behavior around AI echo earlier experiences.

This section explores why AI feels safer and how that perception can mislead, how AI might be used to reduce rather than amplify risk, how logistics and alignment frameworks can translate old safety instincts into new language, and what unique perspective an anti-BW veteran brings to AI ethics and governance.

8.1 Why AI feels safer and why that feeling can be misleading

Compared to handling biological agents or planning for catastrophic outbreaks, sitting at a keyboard working with models and datasets can feel benign. The immediate sensory environment is less charged: no containment suits, no specialized ventilation, no physical samples. The harms associated with AI often seem indirect, mediated through changes in institutions, markets, or information flows.

This relative calm can be soothing to someone coming from the intensity of anti-BW work. It may also lull one into underestimating the stakes. AI systems can influence public opinion, automate aspects of warfare, accelerate research in all domains, and entrench particular values and power structures at scale. The fact that the work is less visceral does not make its consequences less real.

The lesson from anti-BW is that distance from immediate physical hazard does not guarantee ethical safety. The key questions remain: what capabilities are being developed, in whose hands, for what purposes, and with what safeguards. The task is to carry the seriousness of earlier experiences into this new domain, without importing their secrecy or fatalism.

8.2 Using AI to reduce risk rather than amplify it

AI can be a risk multiplier or a risk reducer. In the context of biological threats, it might help identify vulnerabilities in health systems, optimize detection of outbreaks, and model the impact of interventions. It might also lower barriers to designing harmful agents or automating aspects of attack planning. The direction taken depends on choices made by researchers, funders, and policymakers.

An anti-BW veteran is well placed to advocate for risk-reducing applications. This can include focusing AI work on:

- Enhancing public health surveillance rather than optimizing offensive capabilities.
- Improving decision support for ethical review and oversight.
- Identifying patterns of dual-use concern before they reach implementation.
- Supporting transparency and accountability in high-risk research.

At the same time, one can argue against certain lines of AI research that explicitly target offensive optimization, even under defensive pretexts. The experience of watching “defensive” biological work drift toward capability expansion provides a concrete basis for skepticism about similar patterns in AI.

8.3 Logistics and alignment: rephrasing old safety instincts in new language

Conceptual frameworks like logistics and alignment provide an opportunity to integrate anti-BW lessons into AI discourse. Logistics, with its focus on trajectories of alignment between agents and living truth, invites questions about how AI systems are oriented: what they are optimized for, whose interests they serve, and whether they can be guided towards genuinely beneficial ends.

Alignment discussions in AI often revolve around the technical challenge of ensuring that models behave in accordance with specified goals. An anti-BW-informed perspective can broaden this to include upstream questions: are the goals themselves ethically sound, and who has the authority to define them. It can also emphasize the importance of negative space: not only ensuring systems do what they are supposed to do, but also ensuring they are not built to do things that should never be done.

Drawing from past experiences, one can point out that misalignment is not only a problem of rogue systems, but of institutional drift. Organizations can become misaligned with their own stated purposes under pressure from funding, competition, or ideology. AI alignment, in this broader sense, must therefore include governance of institutions as much as tuning of models.

8.4 What an anti-BW veteran uniquely brings to AI ethics and governance

The unique contribution of someone who has lived through anti-BW work to AI ethics and governance lies less in technical expertise and more in pattern recognition. They have seen how dual-use technologies evolve under pressure, how justifications for risky work are framed, how oversight mechanisms succeed or fail, and how human beings cope or break under sustained exposure to worst-case scenarios.

They can bring specific cautionary tales: about programs that began as defensive and drifted; about review bodies that rarely refused anything; about the personal costs of operating in permanent contingency mode. They can also bring a grounded sense that catastrophic misuse is not a theoretical possibility but a live potential that must be taken seriously without tipping into fatalism.

In practical terms, such a person can serve on ethics boards, advisory groups, and design teams, not as an ornament but as someone authorized by experience to ask uncomfortable questions. They can insist that discussions about risk include the long-term psychological and social impacts on those asked to steward dangerous systems. And they can model what it looks like to set and enforce personal boundaries even while remaining committed to broader safety goals.

9. Mapping the Ghosts: Integrating Lives Lived Under Pseudonyms

Living under pseudonyms for extended periods creates a trail of partial selves. Each name corresponds to a context, a set of relationships, and a slice of history. Some of these identities are tied to specific projects that are now defunct. Others to debates in which the pseudonymous voice was more candid than any official persona could be. Over time, the proliferation of such identities can make it difficult to feel like a single person with a coherent life story.

Mapping the ghosts is an attempt to bring these fragments into some kind of order. It does not require revealing true names or classified details. Instead, it involves acknowledging the existence of these past selves, locating them in time and function, and deciding which aspects of their experience to carry forward explicitly and which to leave in peace.

This section considers what pseudonymous lives leave behind, how continuity can be reclaimed without breaching necessary secrecy, how spiritual and psychological integration can proceed, and how to discern what should remain permanently veiled.

9.1 The trail of forgotten names: what compartmentalized lives leave behind

Looking back after decades, one may find that some pseudonyms are vividly remembered while others have faded. Certain code names are still associated with intense periods of work or crisis; others are vague labels attached to minor assignments. Online handles used in obscure forums or early digital communities may no longer be traceable. The net effect is a sense of having left behind a crowd of ghosts, each carrying a piece of one's life.

These compartmentalized lives leave behind more than memories. They may have generated reports, analyses, software, or conceptual frameworks that continue to circulate without obvious attribution. They may have influenced policies or decisions that have since taken on lives of their own. The person behind the pseudonyms lives with the knowledge that parts of their work are out in the world, detached from any public narrative of who they are.

There can be both pride and unease in this. Pride in having contributed to efforts that mattered. Unease in not being able to fully account for where certain ideas have gone, how they are being used, or how they might be misunderstood.

9.2 Reclaiming continuity: telling one coherent story without breaching secrets

Reclaiming continuity does not require disclosing sensitive details. It requires recognizing that all these fragments belonged to a single person, and that their experiences, even if partially hidden, shaped the present self. Telling one coherent story involves identifying major phases: early scientific training, entry into anti-BW work, key transitions, moments of disillusionment, shifts towards AI and other domains.

Within each phase, one can describe the kinds of roles played, the types of problems addressed, and the ethical tensions encountered, without specifying operational particulars. The aim is not to produce a full autobiography, but to allow the mind and, if appropriate, close others to see that there was an arc, not just a sequence of disconnected incidents.

This process can be private, undertaken in journals or conversations with trusted confidants. It can also be partly public, as in the present paper, where the goal is to extract lessons and patterns that may be useful to others. In both cases, the act of narrative integration can reduce the sense of fragmentation and help alleviate some of the alienation that comes from having lived too long in compartments.

9.3 Spiritual and psychological dimensions of decompartmentalization

Integration is not only an intellectual exercise. It has spiritual and psychological dimensions. Spiritually, it involves asking how all these experiences fit within a larger understanding of meaning, vocation, and accountability. Were they merely a series of accidents dictated by circumstance, or was there a thread of calling running through them? How does one understand, forgive, or redeem decisions that now seem questionable?

Psychologically, decompartmentalization involves allowing feelings that were suppressed for operational reasons to surface. Grief for harms witnessed or narrowly averted, anger at institutional failures, shame at perceived complicity, relief at having survived. These emotions do not vanish when ignored; they accumulate in the background. Bringing them into awareness, even partially, can be painful but ultimately healing.

Support from therapists, spiritual directors, or peers with similar backgrounds can be invaluable. The key is to recognize that the skills that served well in the bunker—suppression, compartmentalization, brisk movement past discomfort—are not the skills that serve best in later life when the priority shifts to wholeness.

9.4 What to keep classified forever, and what to offer as public testimony

Not everything should be brought into the light. Some details remain dangerous to disclose, either because they could enable harm or because they implicate people and institutions in ways that cannot be addressed constructively. Determining what should remain classified forever is part of the ethical work of integration.

At the same time, there are aspects of experience that deserve to be shared. Public testimony about patterns of risk, institutional drift, and personal cost can help others navigate similar terrains. It can also serve as a warning against repeating mistakes. The challenge is to find the line where testimony enlightens without endangering.

A guiding principle is to focus on structures rather than secrets. Describing how certain dynamics tend to emerge when science and security intersect can be done without revealing operational specifics. Naming the psychological and ethical pressures at work can help others recognize them, even if the concrete scenarios differ. In this way, the ghosts can be honored without being fully exhumed.

10. An Ethical Will to the Next Generation

An ethical will is not a legal document. It is a way of passing on convictions, cautions, and hopes. From the vantage point of someone who has worked in molecular biology, anti-BW, and AI, such a will has multiple intended recipients: young biologists, AI researchers, and security professionals. Each group faces choices that will shape not only their own lives, but the trajectories of powerful technologies.

This section offers distilled messages to these groups, grounded in lived experience. It does not seek to dictate specific policies but to highlight principles and attitudes that may help them navigate their own challenges. At its core is a plea for humility, courage, and the willingness to refuse certain paths even when they seem technically exciting or institutionally rewarded.

10.1 To young biologists: what Asilomar still has to teach you

To young biologists entering a world where synthetic biology, genome editing, and high-throughput analysis are standard tools, Asilomar may appear distant. Yet its central lesson remains relevant: scientists can and sometimes must pause their own progress to consider risks.

The message is not to fear new techniques, but to recognize that they exist in a social and political context. When an experiment or project carries plausible pathways to significant harm, the responsible response may be to slow down, seek broader input, or decline altogether. Do not let institutional or peer pressure convince you that raising such concerns is naïve or disloyal.

Remember that restraint exercised early can prevent worse constraints imposed later by scandal, accident, or exploitation. It is better to be part of a community that chooses limits than to find yourself working in a field whose boundaries are suddenly enforced from outside with little regard for nuance.

10.2 To AI researchers: how not to repeat the worst mistakes of anti-BW

To AI researchers, the parallel with anti-BW is an invitation and a warning. You are working with systems that can act as force multipliers in many domains, including those involving physical harm. The way you frame your work now will influence how governments, corporations, and adversaries see its potential.

Resist the temptation to justify risky lines of research by appealing solely to defensive needs. Once capabilities exist, they tend to be used, and the distinction between offense and defense becomes blurred. Advocate for clear, enforceable boundaries on applications that should not be pursued. Support structures that can say no, not just offer guidance on how to proceed.

Engage seriously with ethical and governance questions, not as bureaucratic hurdles but as integral parts of your work. Seek out voices from other fields, including those who have lived through earlier cycles of dual-use expansion. Let their caution inform your optimism, without extinguishing it.

10.3 To security professionals: why human limits matter as much as technical ones

To security professionals, the message is that human limits are not obstacles to be managed but realities to be respected. People working with dangerous knowledge and capabilities are not infinitely resilient. Secrecy, high stakes, and constant exposure to worst-case scenarios take a toll.

Design systems and policies that acknowledge this. Provide real support for those showing signs of moral injury or burnout. Do not penalize those who choose to step away after long service. Ensure that oversight and decision-making structures do not rely on a small group of veterans to provide ethical ballast indefinitely.

Recognize that the culture you cultivate around refusal, dissent, and caution will shape the long-term health of your organizations. Technical safeguards and procedures are necessary but insufficient. The character and well-being of the people you ask to steward dangerous tools are equally critical.

10.4 The case for refusal, humility, and quiet courage in an age of extreme capability

Across all these audiences, the core virtues are similar. Refusal: the willingness to say no when a line should not be crossed, even at personal cost. Humility: the recognition that neither individuals nor institutions can foresee all consequences, and that this uncertainty should push towards caution rather than bravado. Quiet courage: the stamina to hold these positions without fanfare, in the face of criticism or misunderstanding.

In an age where capabilities in biology and AI are expanding rapidly, these virtues may seem out of step. Much of the visible culture rewards boldness, disruption, and speed. Yet without refusal, humility, and quiet courage, societies drift into arrangements that only later reveal their full danger. The ethical will here is simple: do not underestimate the importance of these seemingly modest qualities. They are the thin line that separates wise stewardship from catastrophic hubris.

11. Conclusion: Walking Away, Without Looking Away

In the end, this paper is an attempt to reconcile two truths. The first is that there comes a time when one must walk away from certain kinds of work, for the sake of one's own sanity, health, and integrity. The second is that the knowledge and experience gained in that work continue to matter, perhaps more than ever, as new technologies emerge and old patterns repeat.

Walking away is not an abdication of concern. It is a change in posture. The person who leaves the bunker does not suddenly forget the risks that exist. They do, however, choose to engage those risks in different ways: through writing, mentoring, advising at a distance, and advocating for guardrails in adjacent fields like AI.

This conclusion revisits the themes of leaving without abandoning, the value of lived experience in shaping future guardrails, the metaphor of bunkers and windows, and the meaning of growing old with dangerous knowledge.

11.1 Why leaving the field is not abandoning the fight

Leaving anti-BW work at seventy does not mean ceasing to care about biological threats or national security. It means recognizing that direct operational engagement is no longer the best use of one's remaining energy. It also means trusting that others will step into those roles, bringing their own strengths and perspectives.

The fight, if it must be called that, is broader than any particular program or agency. It includes the struggle to keep powerful technologies aligned with humane values, to prevent their worst uses, and to protect those who work at their edges from being consumed by them. One can contribute to that broader effort from outside formal structures, sometimes more effectively than from within.

11.2 The enduring value of lived experience in shaping future guardrails

Lived experience is not data in the usual sense, but it is evidence. Someone who has seen how dual-use technologies evolve, how oversight fails, and how individuals cope under stress brings insights that cannot be fully captured in reports or models. As AI and modern biology advance, such experience can inform the design of guardrails that are realistic rather than naïve.

This requires that institutions make space for testimony that is sometimes uncomfortable. It also requires that those who hold such experience are willing to speak, within appropriate bounds, and to resist the pressure to minimize or sanitize what they have seen. The present paper is one small example of such testimony.

11.3 Bunkers, windows, and the choice to see clearly

The metaphor of the bunker with no windows captures both the necessity and the danger of certain forms of work. There are times when closing the windows is needed to protect people and information. There are also times when the lack of windows becomes a problem in itself, cutting off those inside from wider perspectives and accountability.

Growing older with dangerous knowledge involves a deliberate effort to open windows where possible: to let in light from other disciplines, from younger generations, from public debate, and from one's own neglected emotional responses. This does not mean flinging open the doors and abandoning all protection. It means choosing places and moments where visibility serves safety better than secrecy.

11.4 Final reflection: what it means to grow old with dangerous knowledge

To grow old with dangerous knowledge is to carry a weight that cannot be fully set down. There will always be memories of briefings, scenarios, and programs that others do not share. There will be moments of wondering whether different choices could have led to better outcomes, or fewer harms. There will also be gratitude for having survived and for the chance to contribute in new ways.

The hope embodied in this paper is that such lives can end not in embittered silence or in the reckless shedding of secrets, but in a kind of measured testimony. By speaking about patterns rather than particulars, about ethics rather than operations, those who have walked in the bunker can help illuminate paths for those now working with AI and other emerging technologies.

Walking away, without looking away, is a discipline. It honors both the need for boundaries and the continuing responsibility to bear witness. In that sense, it is a fitting final posture for those who have spent their lives at the intersection of science, security, and conscience.

12. References

Note: This list is illustrative rather than exhaustive. It gathers key works that inform the themes of recombinant DNA governance, dual-use biosecurity, moral injury, and AI safety and alignment.

12.1 Recombinant DNA, Asilomar, and Early Bioethics

- Berg P, Baltimore D, Brenner S, Roblin RO, Singer MF. 1975. Summary statement of the Asilomar Conference on recombinant DNA molecules. *Proceedings of the National Academy of Sciences of the USA* 72(6):1981–1984.
- National Research Council. 1975/2024. *Summary statement of the Asilomar Conference on Recombinant DNA Molecules 20 May 1975*. Washington, DC: National Academies Press.
- Berg P, Singer MF. 1995. The recombinant DNA controversy: Twenty years later. *Proceedings of the National Academy of Sciences of the USA* 92(20):9011–9013.
- Kevles DJ. 1992. Out of eugenics: The historical politics of the recombinant DNA controversy in the United States. In: Kevles DJ, Hood L, editors. *The Code of Codes: Scientific and Social Issues in the Human Genome Project*. Harvard University Press.
- Krimsky S. 1982. *Genetic Alchemy: The Social History of the Recombinant DNA Controversy*. MIT Press.

12.2 Dual-Use Biology, Biosecurity, and Governance

- National Research Council. 2004. *Biotechnology Research in an Age of Terrorism*. Washington, DC: National Academies Press (Fink Report).
- National Science Advisory Board for Biosecurity (NSABB). 2007. *Proposed Framework for the Oversight of Dual Use Life Sciences Research: Strategies for Minimizing the Potential Misuse of Research Information*. US Department of Health and Human Services.
- NSABB. 2023. *Proposed Biosecurity Oversight Framework for the Future of Science: A Draft for Public Comment*. (Updated recommendations on dual use and enhanced PPP oversight.)
- US NSABB. 2007. *Strategic Plan for Outreach and Education on Dual Use Research Issues*. (Education and awareness framework for DURC.)
- World Health Organization. 2004. *Laboratory Biosafety Manual*, 3rd ed. Geneva: WHO.

- World Health Organization. 2010. *Responsible Life Sciences Research for Global Health Security: A Guidance Document*. Geneva: WHO.
- Selgelid MJ. 2009. Governance of dual-use research: An ethical dilemma. *Bulletin of the World Health Organization* 87(9):720–723.
- Koblenz GD. 2010. From biodefence to biosecurity: The emergence of a new policy paradigm. *Biosecurity and Bioterrorism* 8(1):24–29.
- Gillum DR. 2025. Balancing innovation and safety in dual-use research of concern. *Applied Biosafety* 30(1):xx–xx.

12.3 Science, National Security, and Human Subjects

- Moreno JD. 1999. *Undue Risk: Secret State Experiments on Humans*. New York: W. H. Freeman / Routledge.
- Moreno JD. 2004. Bioethics and the national security state. *Journal of Law, Medicine & Ethics* 32(2):198–208.
- Annas GJ, Grodin MA, editors. 1992. *The Nazi Doctors and the Nuremberg Code: Human Rights in Human Experimentation*. Oxford University Press.
- National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research. 1978. *The Belmont Report: Ethical Principles and Guidelines for the Protection of Human Subjects of Research*.

12.4 Moral Injury, PTSD, and Occupational Stress

- Shay J. 1994. *Achilles in Vietnam: Combat Trauma and the Undoing of Character*. New York: Atheneum.
- Shay J. 2014. Moral injury. In: Currier JM, Drescher KD, Nieuwsma JA, editors. *Addressing Moral Injury*. (Reprinted essay discussing the evolution of the concept.)
- Litz BT, Stein N, Delaney E, Lebowitz L, Nash WP, Silva C, Maguen S. 2009. Moral injury and moral repair in war veterans: A preliminary model and intervention strategy. *Clinical Psychology Review* 29(8):695–706.
- Molendijk T, Kramer EH, Verweij DE. 2022. Contextual dimensions of moral injury: An interdisciplinary perspective. *European Journal of Psychotraumatology* 13(1):2054515.
- American Psychiatric Association. 2013. *Diagnostic and Statistical Manual of Mental Disorders*, 5th ed. Arlington, VA: APA.

- Nash WP, Litz BT. 2013. Moral injury: A mechanism for war-related psychological trauma in military family members. *Clinical Child and Family Psychology Review* 16(4):365–375.
- Hoge CW. 2010. Once a warrior, always a warrior: Understanding the military experience. *Daedalus* 139(1):17–26.

12.5 Laboratory Safety, High-Containment Work, and Professional Culture

- Harding A, Parker M. 2015. The ethics of high-containment biological laboratories. *Health Security* 13(1):21–28.
- Klotz LC, Sylvester EJ. 2012. The unacceptable risks of a man-made pandemic. *Bulletin of the Atomic Scientists* 68(3):8–21.
- Millett P, Snyder-Beattie A. 2017. Existential risk and cost-effective biosecurity. *Health Security* 15(4):373–383.
- National Academies of Sciences, Engineering, and Medicine. 2018. *Biodefense in the Age of Synthetic Biology*. Washington, DC: National Academies Press.

12.6 AI Safety, Alignment, and Dual-Use in the Digital Domain

- Amodei D, Olah C, Steinhardt J, Christiano P, Schulman J, Mané D. 2016. Concrete problems in AI safety. *arXiv preprint arXiv:1606.06565*.
- Bostrom N. 2014. *Superintelligence: Paths, Dangers, Strategies*. Oxford: Oxford University Press.
- Russell S. 2019. *Human Compatible: Artificial Intelligence and the Problem of Control*. New York: Viking.
- Christian B. 2020. *The Alignment Problem: Machine Learning and Human Values*. New York: W. W. Norton.
- Gabriel I. 2020. Artificial intelligence, values, and alignment. *Minds and Machines* 30(3):411–437.
- Brundage M, Avin S, Clark J et al. 2018. The malicious use of artificial intelligence: Forecasting, prevention, and mitigation. *arXiv preprint arXiv:1802.07228*.
- National Security Commission on Artificial Intelligence (NSCAI). 2021. *Final Report*. (Sections on dual-use AI, military applications, and governance.)

12.7 Bridging Science, Security, and Ethics

- Jasanoff S. 2005. *Designs on Nature: Science and Democracy in Europe and the United States*. Princeton University Press.
- Guston DH. 2000. *Between Politics and Science: Assuring the Integrity and Productivity of Research*. Cambridge University Press.
- Rappert B. 2007. *Experimental Secrets: International Security, Codes of Conduct, and the Regulation of Dual-Use Biotech*. University Press of America.
- Miller S, Selgelid MJ. 2007. Ethical and philosophical considerations of the dual-use dilemma in the biological sciences. *Science and Engineering Ethics* 13(4):523–580.
- Kassimeris G, Buckley J, editors. 2012. *The Ashgate Research Companion to Modern Warfare*. (Chapters on weapons of mass destruction and ethical challenges.)

These works collectively frame the historical, ethical, psychological, and technical background for the reflections offered in this paper: from Asilomar’s early self-governance, through dual-use biosecurity and moral injury, to contemporary debates on AI safety and alignment.

The author has published over 4,500 AI-assisted papers at:

<https://www.researchgate.net/profile/Douglas-Youvan/research> . Google Scholar has indexed these preprints, and if you want a particular reference, the AI mode of a Google search (Gemini) has efficiently catalogued these preprints.