

Elon Musk, I Still Want to Believe in X

Lev Neymotin, PhD
Uncompressed Thinking

June 21, 2026

Elon Musk,

I am writing this not as a complaint from someone hostile to X, and not as a political attack on the platform, but as a deeply disappointed user who came to believe, gradually and sincerely, that X might become something rare in our time: a public space where honest, direct, sometimes uncomfortable discourse could still survive.

That is why the sudden suspension, or lock, or whatever exact technical category X has assigned to my account, has troubled me far beyond the practical inconvenience of losing access to a platform. Of course, the practical side matters. A person builds a small community, follows people whose voices he values, shares essays, thoughts, arguments, and reflections, and then suddenly sees followers and followings disappear into zeroes, as if human relationships and public participation can be converted into nothing by a hidden switch. But the deeper unease is not about numbers. It is about the nature of the decision, the language in which it was communicated, and the absence of a human explanation.

The message I received was written in the now-familiar grammar of institutional distance:

“Our automated systems have determined that a violation of our Rules did take place.”

Then came the second sentence:

“We considered the additional information provided and decided an overturn of our original decision is not warranted in this case. Therefore, we will not overturn our decision to lock your account.”

I have read these sentences many times, not because they contain much information, but because they contain almost none. They carry the tone of final authority without the substance of explanation. They tell me that something was determined, something was considered, something was decided, and something will not be overturned, but they do not tell me what I did, which post or action caused the violation, which rule was violated, what context was examined, whether a human being actually reviewed the case, or how I am supposed to understand or correct the alleged problem.

What struck me most painfully was the division of responsibility between “our automated systems” and “we.” The automated systems determined. We considered. We decided. We will not overturn. Yet the human subject of the decision is left facing a wall made of passive language and collective pronouns. There is no identifiable reviewer, no specific charge, no meaningful reasoning, no visible standard of judgment, and no accountable voice on the other side.

This may be efficient at scale, but it is deeply dehumanizing at the level of the individual user.

I understand, of course, that a platform as large as X cannot operate without automation. I understand that spam, bots, coordinated manipulation, fraud, threats, impersonation, and real abuse must be detected quickly, and that automated systems are necessary tools in that effort. No serious person can expect every signal to be manually reviewed from the beginning. But there is a difference between using automation as a tool and allowing automation to become the face of judgment itself. A machine can flag a pattern, but a machine cannot fully understand a life of writing, the tone of a public voice, the intent of an essay, the seriousness of a moral argument, or the distinction between strong disagreement and abusive conduct.

This distinction matters especially on a platform that has presented itself not merely as another social media service, but as a public square.

I came to X because I believed that something important had changed there. I believed that after years in which many platforms seemed to narrow public language into approved emotional temperatures and acceptable institutional attitudes, X might offer something more open, more tolerant of disagreement, and more respectful of adult discourse. I believed that a person could write under his own name, take responsibility for his words, argue strongly, criticize power, defend unpopular positions, and still be treated as a participant in public conversation rather than as a problem to be processed.

That belief did not come instantly. It developed over time. I watched the platform, read the discussions, observed the tone of change, and began to feel that X might be one of the few remaining places where public disagreement could still be conducted without every sharp thought being treated as a danger signal. I know that no platform is perfect, and I know that openness creates real challenges. But I also know that without some place for direct and serious speech, public thought itself becomes weak, performative, and dishonest.

That is why this experience has caused me not anger alone, but a kind of embarrassment and sadness. I had started to respect X. I had started to accept it as a serious forum. I had started to believe that it was different. To be told, suddenly and without a specific explanation, that “our automated systems” found me in violation and that “we” will not overturn the decision feels not only procedurally inadequate, but personally humiliating in a way that perhaps only a real user can understand.

I am not an anonymous provocateur. I am not a bot. I am not a spam operation. I am not hiding behind a false identity. I write under my own name. My work is public. My interests are broad: science, technology, medicine, health, politics, war, international law, design, aging, sport, memory, and society. My writing is sometimes direct, sometimes intense, sometimes morally forceful, but it is writing. It is argument. It is analysis. It is reflection. It is the voice of a person trying to think publicly.

Perhaps I am too direct at times. Perhaps my style is too serious for a culture that often prefers safe phrasing. Perhaps some of my conclusions are expressed with too much certainty, or too much moral pressure, or too little diplomatic cushioning. I can accept criticism on that basis. I can revise language. I can clarify meaning. I can learn from a specific objection. But I cannot respond meaningfully to an unnamed violation produced by unnamed systems and reaffirmed by an unnamed “we.”

If I violated a specific rule, I should be told which rule.

If a specific post or phrase crossed a line, I should be shown that post or phrase.

If context was misunderstood, I should have a chance to explain the context.

If the violation was technical, behavioral, or automated rather than content-based, I should know that too.

If the decision was an error, there should be a path by which the error can be seen and corrected.

That is not a demand for special treatment. It is a request for the minimum dignity that should exist between a public platform and a real human being participating in public discourse.

I want to be careful here, because I do not hold Elon Musk personally responsible for this decision. I do not believe that he personally reviewed my account, read my posts, judged my essays, and decided that my voice should be removed or locked. In fact, my concern comes from the opposite direction. I have respected Musk because he has often appeared willing to confront precisely the kinds of dead institutional habits that now seem to be speaking through this enforcement message: bureaucratic distance, language without accountability, and systems that turn human cases into administrative outcomes.

I have viewed Musk as someone with moral courage, technological imagination, and a willingness to stand against stale consensus when he believes the consensus is false. I have even written along those lines. That is why I do not want to conclude that this decision reflects the spirit of X as he intended it. I suspect, or at least hope, that the problem lies somewhere in the many layers between the founder’s stated vision and the ordinary user’s experience.

Those layers matter. Between a public ideal and a user there may be automated systems, policy teams, outsourced review processes, internal risk filters, legal caution, moderation templates, and perhaps individual interpretations disguised as neutral enforcement. At each layer, a little more human judgment may be removed, a little more responsibility may be diffused, and a little more living language may be converted into categories. By the time the final message reaches the user, no one appears to be speaking, yet the decision is final.

This is how institutions become inhuman even when their original mission was human.

A platform may be created to defend open discourse, yet slowly governed by procedures that do not understand discourse. A public square may be built to resist censorship, yet moderated by systems that cannot distinguish moral intensity from abuse, argument from aggression, political disagreement from danger, or a real person from a pattern of signals. A company may speak of free expression at the highest level, while users encounter only the cold language of enforcement at the ground level.

This is the part that alarms me most, and it should alarm others as well, even those who have no interest in my particular account.

The issue is larger than one user. My account is not large in the scale of X. I am not a celebrity, not a government official, not a media figure, and not someone whose disappearance would shake the platform. But precisely for that reason, the case matters in a different way. Freedom of discourse is not measured only by how famous or powerful users are treated. It is measured by whether ordinary, real, serious users can participate without being erased by an unexplained process.

If a small account can be locked without a specific charge, and if an appeal can be denied without identifying the underlying issue, then what exactly has been reviewed? If the answer to a user's question — "What did I do?" — is only that "our automated systems" determined a violation and "we" will not overturn it, then the appeal process has not restored accountability. It has merely repeated the verdict.

This is not how a public square should speak to a human being.

I do not say this because I believe rules are unnecessary. On the contrary, rules are essential. A serious platform must protect users from threats, fraud, spam, impersonation, targeted harassment, and other forms of real abuse. Without rules, open discourse collapses into chaos, manipulation, and intimidation. But rules gain legitimacy through clarity, consistency, and explanation. When rules are enforced without those things, the user is not educated, corrected, or protected. He is simply processed.

There is a profound difference between being told, "This post violated this rule for this reason," and being told, "Our automated systems determined." The first is unpleasant but meaningful. The second is a closed door.

It is also important to distinguish strong language from dangerous language. Public discourse cannot be reduced to polite neutrality, especially when the subjects are war, antisemitism, terrorism, state power, medical risk, institutional failure, political hypocrisy, or the moral cost of human decisions. Some subjects are not soft. Some realities are brutal. Sometimes honest language carries urgency because the reality itself is urgent.

This does not mean that anything should be permitted. Intensity must not become threat. Moral clarity must not become dehumanization. Political disagreement must not become harassment. Public anger must not become violence. But honest intensity is not the same as abuse, and a platform that cannot distinguish the two will eventually suppress precisely the kind of speech that serious public debate requires.

A machine may detect sharpness, repetition, conflict, or emotional force, but it cannot easily understand why a person writes with urgency. It cannot know when a strong phrase comes from grief, historical memory, moral alarm, or reasoned conviction. It cannot understand the difference between someone trying to inflame and someone trying to warn. It may read temperature, but it does not understand conscience.

That is why human review matters.

Not symbolic human review, not a template that says “we considered,” but real review that identifies the content, the rule, the context, and the reason. A user should not have to guess whether the issue was a political post, a phrase about war, a criticism of an institution, a medical warning, an automated misclassification, or some unrelated technical signal. Without specificity, correction is impossible, and without correction, enforcement becomes punishment without instruction.

I am not asking X to agree with me. I am not asking the platform to endorse my views. I am not asking for immunity from its rules. I am asking to be treated as a real person whose words can be judged in context, and whose alleged violation can be named.

There is a moral difference between saying, “You crossed this line,” and saying, “A system determined that you crossed a line, but we will not tell you where the line was.”

The first may be strict, but it is still human. The second is administrative silence disguised as certainty.

My worry is that platforms increasingly train users to accept this silence. A decision arrives, the appeal is denied, and the user is expected to disappear quietly, embarrassed and confused, because the process has spoken. But a process that does not explain itself should not be immune from public questioning. In fact, the more power platforms have over public speech, the more carefully they should explain their use of that power.

X has an opportunity to be different. That is why I am writing this with concern rather than condemnation. I still want to believe that X can be a place where rules and freedom are not enemies, where safety does not require bureaucratic opacity, and where direct speech is not automatically treated as suspicious. I still want to believe that the spirit of the platform is larger than the language of the message I received.

I still want to believe that there are people inside X who understand the difference between moderation and mechanized judgment, between protecting discourse and sanitizing it, between enforcing rules and hiding behind them. I still want to believe that Musk's vision of X is not exhausted by the phrase "our automated systems."

But belief cannot survive indefinitely without process, and trust cannot be built on unexplained decisions.

So I ask, calmly and plainly, for a real review of my account. I ask to know which rule was allegedly violated, which post or action caused the decision, what context was considered, and whether a human being actually examined the matter. If I made a mistake, I am willing to understand it. If a post needs correction, I am willing to consider it. If the system misunderstood my writing, I ask that the misunderstanding be recognized. If the decision was made in error, I ask that the account be restored.

This is not a demand for privilege. It is an appeal for accountability.

I write this with frustration, yes, and with embarrassment, certainly, because it is painful to be treated by a platform one had begun to respect as if one were merely an object of automated enforcement. But I also write with restraint, because my purpose is not to attack X. My purpose is to ask what kind of public square X wants to be.

If X is to remain a place for honest discourse, then the human voice must not disappear behind automated systems, and the human user must not be left facing a verdict without a charge.

Elon Musk, I still want to believe in X.

I still want to believe that the platform can be better than this message.

I still want to believe that a real public square can enforce rules without erasing the dignity of the people who speak there.

But for that belief to survive, "we" must once again mean accountable human beings, and not only a sentence generated by a system no one can question.