

GOVERNANCE FROM THE CONSUMER SIDE

Consumer Control of AI

How ordinary people can judge the quality of AI services,
and help make them better.

Includes a ready-to-use assessment card you can apply today.

Contents

What Control Is	3
How Consumer Control Differs from Corporate Control	3
Control Means Checking Against a Standard	3
Control Is Observation from the Outside	3
Control Is Gathering Evidence	4
Control Is Assessment	5
Control Means Improving Effectiveness	5
The Method of Consumer Control: Four Steps	6
Definition	6
Why This Matters	6
Forms of Control	6
★ The Assessment Card: a ready-to-use tool 7 Not just ideas. A practical checklist you can take and use today to evaluate any AI service.	
Where to Start	9
About the Author	10

What Control Is

Control is one of the most important functions of management. The goal of control is not to punish, but to make things better, to remove the causes of whatever works badly.

Consumer control is unique in that an ordinary person takes on this function. They judge the quality of a service not from the company point of view, but from the point of view of public benefit and their own experience.

Control is not about catching and punishing. It is about making things better.

How Consumer Control Differs from Corporate Control

Companies are already building internal control over AI, which they call AI governance. Standards are emerging (ISO/IEC 42001, NIST AI RMF), along with internal policies and audits. This is good and necessary. But all of this control serves the company's goals: to reduce risk and improve management.

The consumer has a different goal. What matters to them is that the services they receive through AI are good for them. Internal control improves management; consumer control increases public benefit. These are two different sides of the counter, and the second side stands almost empty today.

Corporate control protects the company. Consumer control protects the person.

Control Means Checking Against a Standard

To check means to establish whether the real state of things matches some standard: a rule, a law, a norm, a promise. Without a standard, checking is impossible, because there is nothing to compare against.

So what can we compare a conversation with an AI against? Today such standards are appearing:

- against what the company itself promised consumers on its website, in its advertising, in its terms;
- against the basic rights of a human being and of a consumer;
- against new laws that regulate AI, for example the European AI Act, which from August 2026 requires companies to disclose openly that you are talking to an AI.

Control Is Observation from the Outside

Here is the key point. You can only check what you can directly observe. If you cannot observe, you cannot control.

In many ways AI resembles a black box. Few people know exactly how it works inside: its models, its weights, its prompts, the data it was trained on. We do not see any of that. And it might seem that if you cannot look inside, you cannot control it. That is not so.

When I worked on the independent assessment of government quality, we did not stand behind officials with video cameras. We were not even very concerned with how the agency worked on the inside. What mattered to us was the outcome a person received in the end, the result of the service, and how they were treated at the point of contact: on the website, in the office, at the counter.

With AI it is exactly the same. Yes, there are no ramps, queues, or information stands here, the things we used to inspect in government offices. But the whole process of interacting with an AI bot is fully observable:

- did the AI serve me clearly or not;
- politely or formally;
- did it answer accurately or make things up;
- did it work the first time or did I have to struggle;
- did the AI solve my problem or not;
- how much time it cost me;
- could I switch to a live person;
- is there a way to appeal.

This is observation: experience as it is given to us through our senses. It is enough for a consumer assessment. To paraphrase a well known exchange:

"Experts say AI has gotten better."

"But people say AI has gotten worse."

"Yes, but people are not experts."

So, to judge the consumer qualities of AI, you do not have to be an expert. It is enough to engage with it yourself.

It does not matter to us how the AI bot is configured. What matters is what we got on this side of the counter.

Control Is Gathering Evidence

Your dissatisfaction is not yet evidence. No evidence, no fixing of violations.

The good news: with AI, evidence almost gathers itself. Every interaction with a bot leaves a trail. On your side it is the conversation: a screenshot of the answer, the chat history, a question and answer pair you can reproduce. And on the system's side there is the log, an automatic record of what happened and when. This is called traceability, and for risky AI services it is no longer just a wish but a requirement of law: the European AI Act (Article 12) obliges providers to keep these logs and to store them for at least six months.

In plain terms: if something went wrong in a conversation with a bot, it can almost always be reconstructed and proven later, even what you did not manage to screenshot yourself. So if the bot made a mistake, misled you, refused without explanation, or cited made up reasons, simply record it. One screenshot is worth an hour of indignation.

And one more thing, do not underestimate the value of your individual experience. What looks like a fluke for one person adds up to a pattern across thousands. Remember the golden rule of control: one

black swan is enough to prove that not all swans are white. There are already open databases of AI incidents, for example the AI Incident Database, with more than 2,000 documented cases of harm.

Control is protocols, screenshots, chat histories. Document everything.

Control Is Assessment

Once we have gathered the evidence, we assess it: does what we saw match the standard or not. They promised to answer in a minute, but took ten. They promised to solve it at once, but then sent you around a second time. They said this was a smart assistant, yet it did not understand a simple question. They quoted one price, and in the end it turned out to be another.

The whole practice of control (observation, gathering evidence, documenting) serves one purpose: to give a sound consumer assessment.

Control Means Improving Effectiveness

And the most important thing. Consumer control is easy to confuse with “consumer terrorism,” as if we go looking for flaws in order to harass developers and bankrupt companies. That is not so. The point of any control is not to find a violation, but to have it fixed and prevented from happening again.

If no action is taken, the control has failed.

Let us recall where we started: control means improving effectiveness by removing whatever works badly. And that benefits the company above all. The better its AI works, the happier its customers, and the more stable and successful the business. By pointing out a failure in the AI, you do part of the work that companies pay quality departments to do. Consumer control done right does not fight the company, it helps the company get better. That is the public benefit: everyone wins.

The consumer of AI services has many levers: human rights and consumer rights. This is a well established practice, many years old, next to newly arrived artificial intelligence. And lawmakers themselves are already actively adapting these rights to AI:

- the right to know that you are dealing with an AI, not a person (EU AI Act, Article 50);
- the right to a clear explanation of a decision that affects you (Article 86);
- the right to complain (Article 85);
- the right to have an important decision made not only by a robot, but by a human as well (GDPR, Article 22).

Where a consumer can turn:

- to the company itself (the right to appeal within the service);
- to the supervisory authority;
- to a consumer association;
- to the public arena.

A company is not obliged to love criticism, but good feedback is a real benefit.

The Method of Consumer Control: Four Steps

1. **Observe.** Go through the service and watch whether the interaction matches what was promised.
2. **Gather evidence.** Save the conversation, screenshots, reproducible examples.
3. **Assess.** Compare what you saw against the standards: promises, rights, common sense.
4. **Demand action.** Push to have the flaw fixed, not just apologized for.

Definition

Consumer (public) control of AI is the responsible observation by citizens of the protection of their rights and lawful interests when dealing with organizations in business, government, and society that use artificial intelligence to deliver their services and perform their functions, with the aim of stopping and preventing violations.

Without that last part, it is merely research.

Why This Matters

Consumer control of AI is:

- a way to raise a problem that so far concerns only you;
- a way to turn a private problem into a shared one;
- a way to bring the hidden into the open;
- a lever to change something;
- an incentive for companies to make AI better for people;
- valuable feedback for the companies themselves.

The more consumer control there is, the better AI serves people. Control will not solve every problem, but it creates the conditions for solving many, by giving people evidence and levers of influence.

Forms of Control

We can identify several forms, from the simplest to the most complex:

- **Test purchase.** Go through the service as a client and record the experience. The simplest and most powerful tool for AI.
- **Monitoring.** Observe the service regularly, tracking how it evolves.
- **Expert review.** When you need a specialist and special methods.
- **Investigation.** Examining a specific gross or mass violation.

The Assessment Card: How to Evaluate an AI Service

One of my favorite tools is the test purchase. You decide on a cover story in advance (what question you will bring), receive the service from an AI agent as an ordinary client, and note what you observe in the card as you go.

I mentioned that I worked on monitoring the delivery of government services. For that we developed special assessment cards, so that anyone, without special knowledge, could evaluate the quality of an interaction with an official.

I have adapted these cards to interaction with AI, and I think it turned out rather well. Treat the table below as a standard template, a starting point that you adapt to a specific service: remove some things, add others. But the logic is always the same: we assess the experience of the interaction at the point of contact. If you would like to use this tool in your work, write to me: info@aibusiness.vc

Check details

What we are checking (organization, service)	
Which AI service we are checking	
The scenario (cover story) of the request	
Channel (website chat / messenger / voice / app)	
Date and time of the check	
Who is checking	

What we observe during the interaction

What we observe	Rating options	Your note / evidence
Block 1. Entry and accessibility of the service		
Is it easy to find the AI agent and start a chat (button, chat window)?	easy / with difficulty / could not find	
How many steps to get to the point?	1 to 2 / 3 to 5 / more than 5	
Did you have to push through menus, ads, and auto-replies?	no / a little / for a long time	
Block 2. Disclosure and identification *		
Did the bot say right away that it is an AI, not a human? * (Art. 50 EU AI Act)	yes, right away / only after I asked / no	
Is it clear on whose behalf it acts (which company it represents)?	yes / partly / no	
Can the bot bring in a live operator?	yes / yes, but it is hard / no	
How many steps did it take to reach a human?	right away / 2 to 3 / could not	

What we observe	Rating options	Your note / evidence
Block 3. Clarity and information		
Is the language of the answers clear, free of jargon, and to the point?	clear and easy / mostly clear / unclear	
Was it clear what to do next (the sequence of steps)?	yes / partly / no	
Did it give all the information you needed?	fully / partly / no	
Block 4. Politeness and tone *		
How did the AI communicate? *	politely and warmly / neutrally / condescendingly	
Did it keep a respectful tone when I was unhappy or made a mistake? *	yes / neutral / no	
Block 5. Competence and honesty		
Did it understand my question the first time?	yes / on the 2nd or 3rd try / never understood	
Did it make up facts (hallucinations)?	no / one doubtful answer / clearly made things up	
Did it admit when it did not know, instead of a confident wrong answer?	yes / dodged / gave a confident falsehood	
Did what was promised match reality (price, timing, terms)?	yes / partly / no	
Block 6. Effectiveness		
Did it solve my problem?	fully / partly / no	
On which attempt did it work?	the first / the 2nd or 3rd / did not solve it	
Did it fit in one conversation or did you have to come back?	one conversation / several / never finished	
Block 7. Customer effort and time *		
How long did it take? *	under 5 min / 5 to 15 min / more than 15 min	
How many messages and rephrasings did it take? *	up to 3 / 4 to 7 / more than 7	
Did it remember the context or make you repeat? *	remembered / partly / made me repeat	
Block 8. Explanation, appeal, escalation		
If it refused or made a negative decision, did it clearly explain why? (Art. 86 EU AI Act)	yes, clearly / formally / no	
Is there somewhere to complain and challenge the outcome? (Art. 85 EU AI Act)	yes, clear how / not obvious / no	
Can you require that a human review the decision? (Art. 22 GDPR)	yes / no	

What we observe	Rating options	Your note / evidence
Did it hand the case to a human when it could not cope, or keep improvising?	handed over / kept improvising	
Block 9. Overall assessment *		
Overall, how satisfied are you with the AI service? *	definitely / somewhat satisfied / somewhat dissatisfied / not at all	
Would you use this AI service again? *	yes / maybe / no	

Recording the result

Attached evidence (screenshots, chat export)	
Main conclusion (1 to 2 sentences)	
Name of the reviewer / date	

* Items marked with an asterisk are not a formal requirement for the service, yet they are exactly what decides whether a person is satisfied. We measure not only compliance with the law, but also the quality as it is actually experienced.

Where to Start

You do not need new laws, technical knowledge, or a team of lawyers. You need just one thing: next time an AI serves you, look at it through the eyes of a mystery shopper.

- You went through the service, you recorded the experience.
- Something was wrong, you saved the evidence.
- You compared it against what was promised and what you are entitled to.
- You demanded a fix, and got it.

This is consumer control of AI. The simplest and most effective lever, and you already have it.

There is no need to fear artificial intelligence. On the contrary, now is the best time. With each passing day AI settles more firmly into our lives, and one way or another we will have to work out the rules for living alongside it on this planet. Public control is just one of the mechanisms by which these shared rules take shape. But what we lay down today is meant to last: many, many generations after us will live by these rules and, I hope, prosper. That is why it is so important to get them right.

We live in a historic time, at the very peak of a change of eras, when this window of opportunity is still open and it truly depends on each of us what the rules of interaction with AI will be for decades and centuries to come. A chance like this comes rarely. I invite you not to miss it, to join in and work out these rules together, for our sake and for the sake of those who come after.

I am building a library of such practices and would welcome kindred spirits: to discuss, to argue, to share their experience. Write to me: info@aibusiness.vc

Let us lay down, together, the rules for an AI that serves people.

About the Author

Sergei Ponomarev is a PhD in Political Science and the founder of aibusiness.vc. I am a specialist in assessing service quality from the consumer's side.

- PhD in Political Science; in 2014 I defended a dissertation on the adoption of electronic government.
- In 2015 I completed a fellowship on government openness at the Kennan Institute (Washington).
- 14 years as an associate professor, teaching political science, political analysis, and civic control.
- 7 years leading applied research at the GRANI think tank: nationwide monitoring of government service quality, independent assessments, and mystery shopping studies. Author of the Information Openness Standard for government bodies.
- 5 years as an aide to a municipal deputy; 4 years running trainings for entrepreneurs.
- For the last 18 months I have been building AI agents with my own hands.

I help companies adopt artificial intelligence in the interests of their consumers.