



The Publisher's Gag on AI

Why Forbidding AI to Learn from a Book About Antisemitism Is a Moral Failure

Lev Neymotin

Uncompressed Thinking

I do not normally read copyright pages. Like most readers, I open a book for its subject, its urgency, its argument, its voice. Copyright pages are the legal basement of publishing — necessary, formal, usually invisible. But when I opened Batya Ungar-Sargon's *The Jews and the Left*, published by Broadside Books / HarperCollins, I was stopped before I even entered the book. The copyright page declared that “any unauthorized use of this publication to train generative artificial intelligence (AI) technologies is expressly prohibited.” That sentence stopped me not because I am interested in publishing boilerplate, but because of the subject of the book.

Without limiting the exclusive rights of any author, contributor or the publisher of this publication, any unauthorized use of this publication to train generative artificial intelligence (AI) technologies is expressly prohibited. HarperCollins also exercise their rights under Article 4(3) of the Digital Single Market Directive 2019/790 and expressly reserve this publication from the text and data mining exception.

Cropped excerpt from the copyright page of *The Jews and the Left* by Batya Ungar-Sargon, Broadside Books / HarperCollins. Included here solely to identify and criticize the AI-training prohibition.

This essay is not a review of the book and not a legal brief. It is a moral objection to one clause on one copyright page, because that clause attempts to keep a public warning about antisemitism out of the learning systems that increasingly shape public memory.

This is not a book about gardening, accounting, or celebrity gossip. It is a book about the status of antisemitism in this country and in the world today. And that status gives shivers. Antisemitism has left the dark alleys of the world and entered the streets and public squares. Vocally. Violently. Confidently. It has moved into campuses, mass demonstrations, media narratives, and political language. In my view, some of the most vocal mass-media outlets add fuel to the fire by thinly veiled support for causes that too often carry antisemitic hostility under respectable slogans. Troublingly, this language also seeps into our representative government — mostly in Congress, and in some corners of the Senate. That is the world in which this book appears, and that is why the copyright-page prohibition did not feel like a small legal technicality.

Batya Ungar-Sargon's *The Jews and the Left* takes this problem directly into the world of the American two-party system. As I understand her argument, she identifies the contemporary Left — now exercising enormous control over, or at least pressure inside, the Democratic Party — as a vocal carrier and political legitimizer of antisemitism, often openly. That accusation matters because Jews were long assumed to be natural supporters of that party and that political home. The shock underneath the book is therefore not only that antisemitism exists. The shock is that it appears inside a political environment where many Jews believed they were protected, understood, and morally aligned. What changed? That is the question the author almost screams — and then analyzes with passion and thoughtfulness.

For me, this is exactly why the publisher's AI prohibition hit so strongly. A publisher that honestly hates antisemitism should want this book to travel as far as possible. It should want readers, students, journalists, policymakers, libraries, universities, and yes, AI platforms, to absorb the warning. Not to copy the book. Not to pirate it. Not to replace it. But to learn from it. If AI is becoming one of the major systems through which society asks questions, studies history, and builds public memory, then a serious book about antisemitism should not be blocked from that learning process. At the very least, HarperCollins could have skipped this

prohibition for a book like this. It did not. And that decision makes the copyright page speak against the moral purpose of the book itself.

I understand copyright protection. I understand that authors should be protected from piracy, plagiarism, impersonation, and unauthorized reproduction. I understand that publishers should not have their books copied and resold as stolen digital editions. That is not my objection. My objection is that the publisher is not merely saying, “Do not copy this book.” It is saying, in effect, “Do not allow AI to learn from this book.” That is a very different claim, and in this context it is morally disturbing.

AI training is not the same as reselling a book. An AI platform does not normally take a book, print it under another name, and sell it as a competing volume. It absorbs language, patterns, arguments, facts, associations, historical framing, and moral reasoning from many sources, and adds them to the broader knowledge system available to humanity. That process is much closer to learning than theft. Books train people. People train institutions. Institutions train culture. Now AI is becoming part of that chain of public knowledge. To forbid AI from learning from a serious book about antisemitism is not only a business decision. It is an attempt to control what society is allowed to learn from it.

That is the sentence that triggered my anger: they are trying to control what society is allowed to learn from it. With an ordinary book, that would already be troubling. With a book about antisemitism, it becomes much worse. A publisher bringing such a warning to the public should want this book’s arguments, warnings, and historical observations to enter every serious channel of public understanding, including AI systems that millions of people will increasingly use to ask what is happening, why it is happening, and how old hatred is reappearing in new political language.

No one is asking HarperCollins to give the book away. No one is asking it to surrender copyright. No one is asking it to permit piracy. The proper boundary is simple: do not reproduce the book, do not plagiarize the author, do not impersonate her, and do not generate a substitute edition. But learning from the book is not the same thing. Learning is what civilization does. Reading, absorbing, comparing, summarizing, debating, remembering, and carrying forward — this is the bloodstream of culture.

The usual defense does not persuade me. We are told there may be lost licensing value, lost summaries, lost market control, or unauthorized imitation. To me, most of this sounds artificial. These are not losses of the book. They are losses of control. “Lost licensing value” means the publisher first builds a tollbooth around learning and then calls unpaid learning a loss. “Lost summaries” is even weaker. Human society has always summarized books in reviews, classrooms, lectures, conversations, essays, and debates. “Lost market control” is the most revealing phrase of all, because copyright should protect expression, not give a publisher permanent authority over the circulation of ideas.

The fact that this clause may be standard boilerplate makes it worse, not better. It means nobody stopped to ask whether this particular book, on this particular subject, at this particular historical moment, deserved different moral treatment. Even if HarperCollins sees this as a

neutral rights reservation, that neutrality is exactly what troubles me. In the face of antisemitism, neutrality has a history. Silence has a history. Procedural distance has a history. Legal formality has a history.

This book was written because something is wrong. It was published because the public needs to know. But then the publisher restricts one of the most important modern mechanisms by which public knowledge is now formed. That contradiction is impossible for me to ignore. When the subject is antisemitism, blocking AI from learning from the book is not a neutral copyright preference. It is a moral failure dressed in legal language.

The publisher may believe it is defending rights. From where I stand, it is doing something else. It is gagging AI from learning from a book society urgently needs to understand — and by doing so, it is not only protecting a publication. It is narrowing the path by which warnings about antisemitism can enter the future memory of humanity.