

THE DON'T TRUST THIS BOOK SERIES

DON'T TRUST THIS BOOK



How to think with AI

STUDENTS & PARENTS EDITION

Carl Rosenbacke & Victor Rosenbacke
with Rikard Rosenbacke

do-not-trust.com

The “Don’t Trust This Book” series

DON’T TRUST THIS BOOK

How to think with AI

STUDENTS & PARENTS EDITION

Carl Rosenbacke & Victor Rosenbacke

with Rikard Rosenbacke

© 2026 Carl Rosenbacke, Victor Rosenbacke, and Rikard Rosenbacke.

First edition. All rights reserved.

ISBN: 978-91-992078-0-3

do-not-trust.com

This book is the education edition in the “Don’t Trust This Book” series. Its companions are in preparation: the generic version, *Don’t Trust This Book: How to Think with AI*, and editions for medicine and business. They appear at do-not-trust.com as they are released. The deeper book beneath all of them is *Balance at Speed: Being Human When Technology Accelerates*. Under everything sits a research series, free to read at arxiv.org (four papers published, a fifth on its way); search “Rosenbacke”.

The ideas here are helpful for learning, teaching, and parenting.

All books and tools in the series live at do-not-trust.com.

About the Authors

Carl Rosenbacke. As a teenager in elementary school, Carl created video lessons on quantum physics that were followed by university students. He performed strongly throughout school and is now completing a degree in medicine. Having been a student during the transition explored in this book, he experienced many of its changes firsthand.

Victor Rosenbacke. Victor has tutored students since he was fifteen. He still does, now at university level. He finished a degree in economics at 21 while simultaneously studying medicine. He has sat at more kitchen tables teaching and after-school sessions than most teachers meet in a decade. He knows exactly what it looks like when someone has learned something, and when they only look like they have.

Rikard Rosenbacke, their father, has spent thirty years around high-stakes decisions: in finance, in research, in medicine. He has a PhD in how people think and decide, and how machines fool us. He wrote *Balance at Speed* as a self-reflecting book, written from his own perspective. Here, he mostly stays out of the way except for a few short notes, where the science or the scars are his.

A Note Before We Start

Don't trust this book. Test it. That is not a joke; it is the whole method, and it starts now. Everything in these pages can be tried this week, at a kitchen table or in a classroom. Try it. Keep what works. Drop what doesn't.

Who is this book for? Three people, and we mean all three. If you are a parent of a kid between 10 and 18: this is for you, and you need no technical knowledge at all. If you are a teacher: everything here works in a classroom as well as tutoring in a kitchen, and the boxes tell you exactly how. And if you are a student, reading this yourself: welcome. You are not the problem in this book. You are the hero of it; the whole thing exists so that the AI makes you smarter instead of just faster. It was written so you can read it. Every word.

We will say “the student” a lot. The student is your kid, or your class, or you. The method is the same for all three.

One promise: What this book teaches is training wheels. Then one day the questions in this book have become habits, and the wheels come off by themselves. That is the plan.

How to Read This Book

The book has three parts and a conclusion. Part One is the problem: how a student can get great marks and learn nothing, and why nobody notices (not the parent,

not the teacher, not even the student). Part Two is the method: four steps, called BOLD, that anyone can run on any schoolwork. Part Three is what the adults do: how to coach without becoming the AI police, which never works, at home and in the classroom.

Every chapter ends with a box called TRY THIS. It gives one small, concrete thing to do: one line for parents; one for teachers; one for students. If you only read the boxes, the book still works.

The Appendix at the back is the working kit: a tool the student can paste into any AI, a one-page learning check, subject guides, and a checklist for the adults. Nothing in this book needs any of it to make sense, but once the ideas land, the kit is how they become habits.

Contents

Part One · The Failure	1
1 · The Homework That Looked Right	2
2 · It Is Not an Answer Machine	10
3 · How Learning Drifts from Reality	15
Part Two · The Method	23
4 · The Three Traps	24
5 · The Loop: Be BOLD	30
6 · The Essay Problem	41
Part Three · The Discipline Around You	47
7 · What Cybersecurity Already Knows	48
8 · Building the Brakes	52
9 · Show Your Working	57
10 · The Edge	60
Conclusion · The Whole Game	65
Appendix	69
Appendix A · The Student Tool	70
Appendix B · The Learning Check	73
Appendix C · Subject-by-Subject Guides	75
Appendix D · The Adults' Checklist	77
Appendix E · Where the Depth Lives	79
Appendix F · Plain-Word Glossary	80
Appendix G · References & Further Reading	82

Part One • The Failure

How a student can get great marks and learn nothing, and why nobody notices, not even the student.

1 • The Homework That Looked Right

Not long ago, an essay came home that looked exactly right. The paragraphs were tidy. The argument flowed. There was a clever point in the third paragraph that surprised everyone. The teacher gave it a high mark. The parents were pleased. The student was pleased.

And somewhere, quietly, a small alarm should have been going off. Because there was no way to tell from the essay alone whether anyone had learned anything at all.

For as long as school has existed, a good piece of work proved something real. To write a clear essay, a student had to wrestle the ideas into order. The wrestling was the learning. A clean, strong page told you the struggle had happened. You could read the learning right off the page.

That link has disappeared. A student can now hand in a clean, strong essay without doing any wrestling, because an AI will do it for them, in seconds, for free, at any hour. The page looks exactly the same as before, or even

better. But behind it there may be nothing. A well-made surface with no learning underneath.

Here is the plainest way to see what has changed.

Imagine parents who do their kid's homework, who write the essays, solve the problems, and let the kid hand it all in. The teacher is pleased. The grades are fine. And what has the student learned? Nothing.

Everyone can see the problem when it is the parents holding the pen. AI used the wrong way is exactly that: the same hollow grades, just without the parents' tired evenings.

Here is the part that matters most, and it is the reason this book exists. The student often doesn't know the difference either. They read the smooth explanation the AI wrote. It made sense. So they honestly feel they understand it. The feeling of understanding arrived. The understanding did not. Nobody is lying to anybody. That is exactly what makes it dangerous.

The good page used to prove the learning happened. That link has quietly snapped.

This is not a book about cheating, and it is not a book against AI. Most students who lean on the AI are not necessarily trying to cheat. They are doing what every human does when handed a tool that makes a hard job easy: they use it. The problem is not character. The problem is that, in learning, the hard part is not in the result. The hard part is the wrestling with the task. Take it away, and the marks stay high while the understanding quietly empties out. Nobody notices until much later, when the marks finally catch up with the truth.

THE NUMBERS

58% of 9–12-year-olds already use AI chatbots, below the official age limit of 13. (Internet Matters, UK, 2025)¹

64% of teenagers use AI chatbots; about three in ten use them every day. (Pew Research Center, USA, 2025)²

62% of parents worry about the accuracy of what AI tells their children, but only 34% have talked with their child about how to judge it. (Internet Matters, 2025)¹

Read those three numbers together and you have the whole situation. The students are already there (most of them before any adult officially allowed it). The adults are worried. And the conversation that would actually

help has mostly not happened. This book is that conversation.

CARL · THE QUANTUM TUTORIALS

As a teenager, I made video lessons on quantum physics. I remember sitting in front of the camera, trying to explain an idea I thought I understood, only to realise that I could not yet make it clear. I had to go back, work through it again, and rebuild the explanation step by step until someone else could follow it. Later, I discovered that some of the people watching were master's and PhD students. That made the standard very real. Smooth language was not enough. If I did not understand the physics, they would notice almost immediately. Teaching forced me to turn information into something I could actually defend, explain, and reconstruct for another person.

We open with that story for a reason. The deepest learning either of us ever did came in the moments we could not fake, where we had to rebuild an idea, out loud, for someone who would notice if we were bluffing. That is exactly the thing AI offers to remove. And it is exactly the thing this book protects. Hold on to one line: the test of real learning is whether the student can rebuild the idea themselves, for someone who is

watching. Everything else in these pages is a way of getting there.

The four steps, in one breath

The way there has a name, and you will meet it properly in Chapter 5. We give it to you now, because the chapters in between will make more sense if you can already see its shape. Four steps in a loop, and their first letters spell BOLD:

Build. Before the AI says a word, put down what you already think: your own idea, your rough guess, your first attempt. It can be thin. It has to be yours.

Oppose. Then turn the AI against it. Do not ask “is this right?” Order it to attack: find the weakest point, the missing case, the best counter-argument.

Look. Step outside the chat and check against something real: the textbook, a second source, a fresh problem solved with no machine open.

Decide. Choose what goes into your answer, and why. The decision is yours, and you can say why every piece earned its place.

That is the whole loop: your thinking first; the AI against it; the world outside it; the decision yours. When Build, Oppose, Look or Decide turn up in the next chapters, this is what they mean. Chapter 5 shows how to run the loop in one real evening.

A NOTE FROM RIKARD

I have spent thirty years watching capable adults approve things that looked exactly right and turned out to be hollow. Confident reports. Polished numbers nobody had checked. Our minds carry an old weakness: if a thing is well said, we assume it is well thought. For most of history that shortcut was safe, because sounding good took real work. It is not safe anymore, at any age.

What my sons describe in schoolwork is the same failure I have watched wreck grown-up decisions, just earlier in life, and far cheaper to fix since the scale was not as massive. A fourteen-year-old who learns to ask “is this actually right, and could I do it myself?” is building the most valuable habit I know. Most adults in serious jobs still don’t have it.

Rikard

One more thing, because it decides who needs this book most. This has nothing to do with good students and

bad students; those labels are not real, and this book will never use them. This trap has nothing to do with being lazy. It runs on pressure: the more a grade matters to you, the more tempting the shortcut that keeps it high. So, the students who care the most are exactly the ones it catches most easily, quietly, while everyone stays pleased. If you care a lot about doing well, this book was written for you. That is not a warning about who you are. It is a fact about how the trap works for all of us.

TRY THIS

If you are a parent: Ask about one recent piece of work: “Talk me through how you got to that point; I’m curious how you worked it out.” No screen; no page. If they can rebuild it in their own words, the learning is there. If they can only point at the page, now you both know, without a single accusation.

If you are a teacher: Pick one strong written piece and ask the student to explain their favourite point from it, out loud, one minute, no notes. Not as a check, as interest. What they can rebuild is theirs. What they cannot is where teaching starts.

If you are a student: Pick something you handed in recently and try to explain the main idea out loud, to nobody, without looking. If you can, it’s yours, machine

or no machine. If you can't, you just found out something the mark would never have told you.