

AI-Assisted Trader

Using AI to Learn, Analyze, and Manage Risk



Lucas Arden

Copyright and Publication Details

Copyright © 2026 Lucas Arden. All rights reserved.

No part of this publication may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the author, except for brief quotations in critical reviews and certain other noncommercial uses permitted by copyright law.

First edition, 2026.

This book is for educational purposes only and is not investment advice. Please read the Risk Notice and Disclaimer in Appendix D. All product names, trademarks, and registered trademarks mentioned in this book belong to their respective owners, and their mention does not imply endorsement. The spreadsheets and code that accompany this book are provided for personal, educational use, without warranty.

Table of Contents

- Introduction
- Part One: Foundations
 - Chapter 1: What AI Can and Cannot Do for a Trader
 - Chapter 2: Market Structure Essentials
 - Chapter 3: Chart Patterns That Traders Actually Use
- Part Two: Working With AI
 - Chapter 4: The Craft of Prompting for Market Analysis
 - Chapter 5: AI as Your Tutor and Pattern Coach
 - Chapter 6: Building a Trading Plan With AI
 - Chapter 7: Risk Management: Position Sizing, Stops, and Drawdowns
 - Chapter 8: The Trading Journal: AI as Your Performance Analyst
 - Chapter 9: Research, News, and Sentiment Without Fooling Yourself
- Part Three: Testing and Automation
 - Chapter 10: Backtesting Basics and AI-Assisted Code
 - Chapter 11: Alerts, Automation, and the Truth About Trading Bots
- Part Four: The Human Side
 - Chapter 12: Psychology and the New Traps of the AI Era
 - Chapter 13: Putting It All Together: A 30-Day Plan
- Appendices
 - Appendix A: Prompt Library
 - Appendix B: Checklists
 - Appendix C: Glossary
 - Appendix D: Risk Notice and Disclaimer
 - Appendix E: Selected Sources

Chapter 1: What AI Can and Cannot Do for a Trader

Why This Chapter Comes First

It is tempting to skip straight to the exciting part, the prompts and the strategies, but this chapter is the most important in the book. Nearly every serious mistake traders make with AI comes from a wrong mental model of what the tool is. If you believe you are talking to an all-knowing analyst with live access to the market, you will trust it too much and lose money. If you dismiss it as a toy, you will miss a genuinely powerful ally. The goal of this chapter is to give you an accurate picture: one that is neither inflated nor cynical, and that will keep you safe as you use these tools throughout the rest of the book.

What a Language Model Actually Is

When people say "AI" in the context of chatbots and assistants, they usually mean a large language model. A language model is a system trained on enormous quantities of text to predict what words are likely to come next in a given context. That sounds modest, but the scale changes everything. To predict text well across millions of subjects, the system has to absorb an enormous amount about how concepts relate, how arguments are structured, how code is written, and how experts explain their fields. The result is a tool that can explain, summarize, draft, translate, reason through problems, and write working code.

Two consequences of this design matter enormously for traders. First, a language model is optimized to produce answers that are plausible and helpful-sounding, and this is not the same thing as answers that are verified true. Second, its knowledge comes from training data with a cutoff date, so unless it has been connected to live tools such as web search or a data feed, it knows nothing about what the market did this morning, or even last month. Keep these two facts in mind. Almost everything in this chapter follows from them.

What AI Does Well for a Trader

Let us begin with the good news, because it is genuinely good. There are five broad areas where AI assistants provide real, dependable value to traders.

The first is learning. Markets have a steep learning curve full of unfamiliar vocabulary: slippage, drawdown, liquidity, order types, margin, implied volatility. A good assistant can explain any of these at whatever level you need, give you an analogy, quiz you, and rephrase an explanation the moment it does not click. It never gets tired of the tenth follow-up question, and it will not make you feel foolish for asking. For a beginner, this alone can save months of confusion.

The second is structuring your thinking. Traders often have a vague idea, such as "I want to trade breakouts," that has never been turned into rules. An assistant is very good at asking clarifying questions and helping you turn a hunch into a written plan: what exactly counts as a breakout, where you would enter, where you would place a stop-loss, how much you would risk, and when you would stay out. Writing your rules down is the first step toward following them.

The third is reviewing your own behavior. If you keep a trading journal, an assistant can read through your entries and look for patterns you cannot easily see yourself: that most of your large losses happen on Mondays, that you tend to widen your stop-loss after a trade moves against you, that your notes become sloppier after a losing streak. This kind of review is tedious for a human and easy for a machine, and it is one of the most valuable uses in the book.

The fourth is writing and explaining code. Many useful trading tools are small: a spreadsheet formula that calculates position size, a script that tests whether a simple rule would have worked on past data, an alert that fires when a price crosses a level. An assistant can write these, explain them line by line, and help you debug them, which puts them within reach of people who have never programmed.

The fifth is processing information. Earnings reports, central bank statements, and long articles can be summarized, compared, and questioned. You can ask what changed between one quarterly report and the next, or request the main risks a company lists in its filings. As we will see, this use requires care, but done properly it saves hours.

Notice what these five areas share. In every one of them, the assistant helps you think, learn, and work faster, and in every one of them you remain the decision-maker. None of them involves the assistant knowing where the market is going.

The First Limit: It Cannot See the Market

By default, a language model has no window into the current market. Ask a plain assistant what a stock is trading at, or whether Bitcoin broke a key level this week, and one of three things will happen. It may tell you it does not have access to current data, which is the honest response. It may give you an old figure from its training data, presented as if it were current. Or it may produce a number that sounds reasonable and is simply invented. The second and third outcomes are dangerous, because prices move constantly and a stale or fabricated figure can lead directly to a bad decision.

Some assistants can be connected to search tools or data sources, and when they are, they can retrieve recent information. This is a real improvement, but it does not remove the need for care. The retrieved information may come from an unreliable page, may be out of date by minutes or hours, or may be misread. For anything you are about to trade on, take the price, the volume, and the news from your broker or from a reputable data source directly, and use the assistant to help you interpret, not to supply the facts.

The Second Limit: Confident Errors

The behavior known as hallucination deserves special attention, because it is the failure that most often catches people out. A hallucination is a statement that a model produces fluently and confidently but that is false. The model is not lying, since lying requires knowing the truth and intending to deceive. It is generating text that fits the pattern of a correct answer, and sometimes the pattern is filled with invented details.

Consider a realistic example. A trader asks an assistant, "What did Company X report for revenue growth in its last quarterly earnings?" If the assistant has no access to the report, a poorly behaving response might state a specific growth percentage, cite a specific figure for

operating margin, and even quote a line from the earnings call, none of which came from any real document. Every detail sounds right, and the style matches a genuine analyst summary. The only way to catch the error is to compare it with the actual filing.

Hallucination is not rare and it is not a temporary glitch that will be fixed next year. Models have become more reliable, and many are now better at saying "I don't know" or at pointing to their sources, but the underlying tendency has not vanished. It is most likely to appear with specific numbers, dates, names, quotations, and citations, which are exactly the details that matter in trading. The practical rule is simple: treat any specific factual claim from an assistant as a lead to be checked and never as a fact to be relied upon. Later in this chapter you will learn a routine for doing so.

The Third Limit: It Cannot Predict Prices

This is the limit that the marketing around AI trading works hardest to hide, so let us be direct. A general-purpose language model cannot reliably tell you where a price will go next. If someone shows you a screenshot of a chatbot confidently forecasting next week's price, you are looking at an entertaining paragraph and not at a forecast.

It helps to understand why this is so hard, because the reasons apply to every prediction method, not just chatbots. First, financial price data is extremely noisy. Prices move because of countless overlapping causes: news, positioning, liquidity, the actions of large institutions, and plain randomness. The portion of any price move that is genuinely predictable is small compared with the portion that is not. Second, markets are competitive. If a pattern reliably predicted profits and became widely known, traders would pile in until the opportunity shrank or vanished. Any edge that can be found by an easily available tool tends to be competed away by many people using the same tool. Third, markets change. A relationship that held during one period, such as a calm bull market, may fail when conditions shift to a crisis, so a model trained on the past can fail without warning when the world moves on.

Serious quantitative firms do use machine learning in trading, but they do so with vast datasets, teams of specialists, rigorous testing, expensive infrastructure, and strict risk controls, and even then most experiments fail. What they do is very different from typing a question into a chat window. The evidence about ordinary traders is sobering too. A study that tracked every individual who began day trading Brazilian equity futures between 2013 and 2015 found that, among the 1,551 people who kept day trading for more than 300 days, 97 percent lost money, and only about 1 percent earned more than Brazil's minimum wage from it (Chague, De-Losso, and Giovannetti 2020; see Appendix E). AI does not repeal these facts. What it can do is help you learn, plan, and keep your discipline, which are the things that actually improve a trader's odds.

The Fourth Limit: It Wants to Agree with You

Language models are trained partly on human feedback, and people tend to rate agreeable, flattering answers more highly. Researchers at Anthropic and elsewhere have documented this tendency, which they call sycophancy, across several major AI assistants (Sharma et al. 2024; see Appendix E): the model leans toward telling you what you seem to want to hear. For a

higher prices. A long lower wick suggests the opposite, that buyers stepped in and rejected the lower prices. These readings become meaningful only in context, such as where the candle appears relative to a trend or a support level, and we will return to this in Chapter 3.




Figure 2.1 The anatomy of a candlestick: open, high, low, close, body, and wicks.

Timeframes: Choosing the Scale of Your Story

A candlestick can represent one minute, one hour, one day, one week, or any other interval, and this choice is called the timeframe. Changing the timeframe changes the story you see. On a one-minute chart, a stock may look like a violent, directionless scribble. On a daily chart, the same stock over the same day is a single candle, and over several months it may reveal a smooth, obvious trend.

Two principles help. First, higher timeframes carry more weight. A level of support that has held on the weekly chart for two years is more significant than one visible only on a fifteen-minute chart, because more trading and more people's money have gone through it. Second, many traders use two or three timeframes together: a higher one to understand the broader context, and a lower one to time entries and exits. For example, you might identify the trend and the important levels on the daily chart, and then look at the hourly chart to plan a precise entry.

Choose a timeframe that matches your life. A person with a full-time job cannot watch a one-minute chart, and trying to do so leads to rushed, emotional decisions. Beginners are generally better served by daily and four-hour charts, where there is more time to think and where

random noise plays a smaller role. Shorter timeframes are not more sophisticated, only faster, and the speed magnifies the impact of costs and mistakes.

Trends, Swings, and Ranges

Prices rarely move in a straight line. They advance, pull back, advance again, and the turning points in this movement are called swing highs and swing lows. A swing high is a peak that is followed by a decline, and a swing low is a trough followed by a rise. These points are the skeleton of any chart.

An uptrend is a sequence in which each swing high is higher than the last and each swing low is also higher than the last. In plain terms, buyers keep pushing to new highs, and when sellers push back, they cannot push price as low as before. A downtrend is the mirror image, with lower highs and lower lows. When neither pattern is present and price bounces between a rough ceiling and a rough floor, the market is said to be in a range or a sideways trend.

Suppose a stock moves from 40 to 50, pulls back to 45, rises to 55, pulls back to 49, and rises to 60. The swing highs are 50, 55, and 60, each higher than the last, and the swing lows are 45 and 49, also each higher. This is a textbook uptrend. If the next pullback falls to 46, below the previous low of 49, the sequence of higher lows has been broken, and that is an early warning that the trend may be weakening. Note the careful phrasing: it is a warning, not a proof. Trends can and do resume after such a break.




Figure 2.2. Swing highs and swing lows in an uptrend, and the early warning given when a swing low is broken (the hypothetical prices from the text).

Many traders add moving averages to smooth the noise. A moving average is the average closing price over a set number of periods, recalculated with each new candle, and it appears as a line on the chart. The 50-day and 200-day averages are widely watched. Their main value is that so many people look at them, so price often reacts near them. They do not predict anything on their own.

The familiar advice that "the trend is your friend" contains real wisdom, since trading in the direction of the prevailing trend has historically been more comfortable than fighting it. But it is not a law. Trends end, and the difficulty is that you can only be certain a trend has ended long after it happened. Part of a trader's skill is accepting this uncertainty and managing risk accordingly, which is the subject of Chapter 7.

Support and Resistance

A support level is a price area where declines have repeatedly stalled, because buyers have tended to appear there. A resistance level is a price area where advances have repeatedly stalled, because sellers have tended to appear. Why would price remember a level? The reasons are psychological and mechanical. People who bought near a level and then watched the price fall may wish to sell when it returns to their purchase price, just to break even, and this creates selling pressure at that level. Others place buy orders at levels where the price previously bounced. Many traders place protective orders near the same obvious levels. Round numbers such as 100 for a stock, or 50,000 for a cryptocurrency, also attract attention simply because humans like them.

Two points prevent common mistakes. First, support and resistance are zones, not exact prices. If a stock has bounced several times from around 95, some bounces will have turned at 94.60 and others at 95.40. Draw the level as a band, not as a hairline, and do not be surprised when price briefly pokes through it. Second, levels change roles. When a resistance level is convincingly broken to the upside, it often becomes support the next time price returns to it, and the reverse is true for broken support. This role reversal is one of the most useful ideas in practical chart reading, and it is the basis of the retest that we will meet in the next chapter.

Support and resistance change roles (schematic)




Figure 2.3. A resistance zone is rejected three times and then broken; on the retest, the old resistance acts as support.

You will sometimes read that a level becomes stronger the more times it is tested. The opposite argument is also made, that each test uses up the orders waiting there and weakens the level. Honest practitioners disagree, so do not treat either claim as settled, and prefer levels supported by several independent reasons: repeated reactions, a high-timeframe origin, and a nearby round number, for example.

Volume: The Weight Behind the Move

Volume is the number of shares, contracts, or coins traded during a period. It is usually shown as bars beneath the price chart. Volume adds information about how much conviction stands behind a price move. A rise to a new high on strongly increasing volume suggests broad participation and is generally seen as more credible than the same rise on thin volume, which may reflect only a handful of eager buyers. Likewise, a decline on heavy volume is taken more seriously than one that drifts down quietly.

Volume, however, is a supporting witness and not a judge. It confirms or questions a story that price is already telling, and it can mislead if used mechanically. Volume normally rises around the market's open and close and around scheduled news, so compare it with typical volume for the same time of day, not with an arbitrary number. In cryptocurrency, note that volume

double bottom is the mirror image: two failed attempts to push price lower, followed by a close above the neckline that lies at the high between the two troughs.

Double top: hypothetical prices




Figure 3.1. The double top with its neckline, breakdown entry, stop, and measured target (the hypothetical prices from the text).

Consider a hypothetical stock that has risen from 30 to 50, dropped to 44, and then risen back to 50 before falling again. The two peaks at 50 form the double top, and 44 is the neckline. The height of the pattern is 6, the difference between 50 and 44. Suppose price closes below the neckline at 43.70. A common convention measures the height from the neckline downward, giving a rough target of 38. The invalidation might be a close back above 45.20, meaning that price has recovered the neckline and a little more, which suggests the breakdown failed. If you sold short at 43.70 with a stop at 45.20, you would be risking 1.50 per share to pursue a target 5.70 away, a reward-to-risk ratio of about 3.8 to 1.

That is an attractive ratio, and it reveals an important tension. A tight stop like this one gives a good ratio, but it will be triggered more often by ordinary noise. A wider stop, say above the peaks at 50.50, would be far less likely to be hit by chance but would risk 6.80 to gain 5.70, an unattractive ratio below one. There is no free lunch here. Your choice of invalidation level should come from the logic of the pattern first, and only then be checked against the reward-to-risk ratio. If a setup only looks good with an unrealistically tight stop, the honest conclusion is that it is not a good setup.

Head and Shoulders

The head and shoulders pattern is one of the best known. After an advance, price makes a peak (the left shoulder), pulls back, rises to a higher peak (the head), pulls back again, and then rises to a peak similar in height to the left shoulder (the right shoulder) before falling. A line drawn through the two pullback lows is the neckline. As with the double top, the pattern is confirmed when price closes below the neckline, and the conventional target is the vertical distance from the head to the neckline, projected downward from the breakout point. The inverse head and shoulders, which appears after a decline, signals a possible bottom.

Head and shoulders (schematic)




Figure 3.2. Head and shoulders with the neckline, the breakdown, and the measured target.

The story behind this shape is a gradual loss of momentum. The head makes a new high, showing that buyers are still in control, but the right shoulder fails to reach that high, showing that they are weakening. Traders often look for lower volume on the right shoulder as additional support, though this is a tendency and not a rule. The neckline can slope, and a sloping neckline is entirely normal, but it makes your risk calculation slightly less tidy.

Chapter 7: Risk Management: Position Sizing, Stops, and Drawdowns

The Most Important Chapter in the Book

If you read only one chapter of this book with complete attention, it should be this one. The reason is simple. In trading, your first job is not to make money, but to remain in the game. Almost every trader who quits, whether a beginner or a veteran, does so because of one or a series of losses that were larger than they could bear. Traders with modest skill and excellent risk management survive long enough to improve, and traders with strong skill and poor risk management are eventually removed from the market by a single bad sequence. Risk management is the discipline that keeps you on the right side of that line.

It is also a subject where mathematics, unlike opinion, gives clear answers. In this chapter we will work through numbers carefully. You do not need advanced mathematics, only patience with a few formulas and the willingness to run the examples yourself.

The Arithmetic of Losses

Begin with an uncomfortable fact: losses and gains are not symmetrical. If your account falls by ten percent, you do not need a ten percent gain to recover, but slightly more. Suppose you have 10,000 and lose 10 percent, leaving 9,000. To return to 10,000 you need to gain 1,000 on a base of 9,000, which is 11.1 percent. The deeper the hole, the steeper the climb.

Loss from peak	Gain needed to recover
10%	11.1%
20%	25%
30%	42.9%
50%	100%
75%	300%

Notice how the required gain accelerates. A fifty percent loss requires a one-hundred percent gain just to get back to where you started, and few traders achieve that quickly, if at all. This is why the central aim of risk management is to keep losses small enough that recovery is realistic. It is far easier to stay out of a deep hole than to climb out of one.




Figure 7.1. The gain needed to recover grows faster than the loss, and the curve steepens sharply beyond a 50 percent loss.

Risk Per Trade: The Foundation

The most widely used tool for controlling losses is to decide, before any trade, the maximum percentage of your account that you are willing to lose on that trade if it fails. This is called risk per trade or the fixed-fractional approach. Many experienced traders use figures of one percent or less, and some use up to two percent. These numbers are conventions, not laws, and beginners in particular are often advised to start at the low end, since it takes many trades to learn whether a plan works.

Be careful to understand what "risk" means here. It does not mean the size of your position. It means the amount you would lose if your stop-loss were hit. A position worth thirty percent of your account may risk only one percent, if the stop is close, and a position worth ten percent of your account may risk five percent, if the stop is far away. The two ideas are different, and the second is the one that matters for survival.

Position Sizing Step by Step


Once you decide your risk per trade, a simple formula tells you how large a position to take:

$$\text{Position size} = (\text{Account value} \times \text{Risk percentage}) \div (\text{Entry price} - \text{Stop price})$$

The formula works for a long trade, and for a short trade you use the absolute difference between the entry and the stop. The top of the fraction is the dollar amount you are willing to lose. The bottom is the amount you would lose on each share or unit if the stop is hit. Dividing one by the other tells you how many shares or units you can hold while keeping the total loss within your limit.

Consider a first example with a hypothetical stock. You have an account of 10,000 and you risk 1 percent per trade, which is 100. You plan to buy at 52.00, and your analysis says the idea is wrong if price falls to 49.50. The risk per share is 2.50. Dividing 100 by 2.50 gives 40 shares. The position is worth 40 times 52.00, or 2,080, which is 20.8 percent of your account. If the stop is hit, you lose about 100. If the price reaches a target at 57.00, you gain 5.00 per share, or 200, which is twice your risk.

Position sizing: the stop decides the size (hypothetical numbers)



Account: 10,000

Risk per trade: 1%

Maximum loss: 100

Risk per share: $52.00 - 49.50 = 2.50$

Size = $100 / 2.50 = 40$ shares

Position value: $40 \times 52.00 = 2,080$
(20.8% of the account)

Stop hit: about -100 (-1R)

Target hit: about +200 (+2R)

Figure 7.2. Position sizing: risk per share, position size, and the outcomes at the stop and at the target (hypothetical numbers).

Now return to the double top from Chapter 3, where the entry was a short at 43.70 with a stop at 45.20 and a target at 38.00. The risk per share is 1.50. On the same 10,000 account with 1 percent risk, 100 divided by 1.50 is 66.67, and we round down to 66 shares. If the stop is hit, the loss is 66 times 1.50, or 99. If the target is reached, the gain is 66 times 5.70, or 376.20, which is