

The 91% Question

A stock just fell 91% in a single day. Every textbook indicator says it is now "oversold" — the classic signal to buy. Should you?

Getting that answer right, automatically, every night, turned out to be a problem about *refusing* rather than predicting. This is what I built, what the data actually said, and what I would not let it claim.

Svetoslav Rankov (sved) · sved.net · Engineering & automated systems

The trap that catches almost everyone

The honest answer to the 91% question is **almost always no** — and the reason is the single most useful thing I learned building this. In small-company biotech, a stock that has collapsed is usually not a bargain. It is a company whose drug just failed, or whose regulator just said no, or that is about to issue a flood of new shares at a discount simply to survive. The price fell because the news was bad. The market was right.

Yet the standard tools of technical analysis — the indicators built into every brokerage app — will light up green on exactly that stock. They measure only how far and how fast a price has moved. They have no idea *why*.

A screener tells you what is cheap. It cannot tell you what is cheap for a reason. That gap is where beginners lose money.

So the system I built is organised around a mechanism most trading tools do not have: a **veto layer**. Before anything can be called an opportunity, the software goes and reads the company's regulatory filings and asks whether there is a known reason the price is low. If there is, no signal is produced — regardless of how attractive the chart looks.

59

COMPANIES
MONITORED NIGHTLY

\$0

DATA COST — EVERY
SOURCE IS FREE

~4 min

TO READ THE WHOLE
LIST OVERNIGHT

87

AUTOMATED TESTS
PROTECTING IT

Three things it caught in its first week

None of these required prediction. Each one required reading a document that was publicly available and that nobody reads.

1. The falling knife

A company called Sionna fell **91% in one session**. Its indicators were textbook oversold. The system found the regulatory filing published that same day, tied it to the collapse, and blocked the signal. A naïve tool would have called it a buying opportunity.

2. The silent dilution

Outlook Therapeutics looked cheap — down 19% on the day and 70% from its high. The filings showed two share sales in 48 hours and enough cash to last roughly **seven weeks**. When a company is that short of money it must issue new shares, and it must do so at a discount. That is not a dip. The system scored it *avoid*.

3. The insiders who left first

At Capricor, executives sold \$3.1 million of stock at \$30.38 in late June. Six weeks later the shares traded at \$4.21. Meanwhile at Summit Therapeutics, both chief executives bought **\$103.6 million** of their own company's stock on the open market. Both facts sat in public filings the whole time.

THE DETAIL THAT MAKES INSIDER DATA USABLE

Most people who track "insider buying" get it wrong, because a filing shows several kinds of transaction and only one of them means anything. An executive receiving shares as a bonus, or exercising options, both appear as "acquired". Only an open-market *purchase* — spending their own money — carries a signal.

One real filing showed an executive exercising options at \$0.99 and selling at \$5.06 minutes later. Counted carelessly, that reads as insider buying. It is precisely the opposite.

What it actually does, every night

Reads

Pulls prices, regulatory filings, insider trades and trial records from official sources.

Checks

Applies fixed arithmetic — no guesswork — to find anything unusual, and runs the veto layer.

Investigates

An AI model researches what was flagged, reads the source documents, writes the analysis.

Reports

Emails a readable briefing and sends a phone alert — but only when something changed.

The important design decision is the split between steps two and three. An AI model is excellent at reading a legal document and judging what it means. It is unreliable at remembering a number. So it is never allowed to be the source of one: every figure traces back to an official filing, or the report says the figure is unknown.

Most nights it sends nothing at all. That is the feature, not a fault.

Three lessons you can use without any software

- 1

Ask why it is cheap before you ask how cheap

Before looking at a single chart, check whether the company recently sold new shares, received bad news from a regulator, changed its accountants, or is running out of cash. All of this is public. If any of it is true, the low price is an explanation, not an opportunity.
- 2

Insider buying is informative. Insider selling mostly is not

People sell shares for a dozen innocent reasons — tax, a house, a pre-scheduled plan. They buy for one. Weight the two very differently. When I measured this, selling was so common it flagged more than half the companies being followed, which makes it useless as a signal.
- 3

Compare against the alternative you actually have

A 9% gain sounds good. If the whole sector rose 12% over the same weeks, you lost — you would have done better buying a fund that tracks the sector and doing nothing. Measuring against zero flatters every strategy ever devised. This one lesson changed my conclusions more than anything else.

The finding I did not want

Having built it, I tested whether the signals were worth anything — scoring every alert against the sector fund anyone can buy instead. Three widely used rules turned out to be worthless or harmful:

COMMON RULE	AFTER 3 MONTHS	WHAT IT REALLY MEANS
"Buy stocks near their yearly low"	-2.68%	Worse than doing nothing. Cheapness alone is noise
"The more oversold, the better"	-9.59%	The most beaten-down names are beaten down for a reason
Oversold, <i>with</i> heavy selling volume	+1.72%	The only version that survived scrutiny

THE HEADLINE RESULT, STATED PLAINLY

Measured against simply owning a biotech index fund, **the core signal did not beat it** — unless confirmed by unusually heavy trading volume. So the system now demands that confirmation before calling anything actionable, and the reports say openly that the technical trigger is a timing tool, never a reason to buy on its own.

I could have left this out. It is the most valuable thing the project produced.

How it is built — and why that matters

A system that runs unattended every night, that nobody watches, has to be built differently from a script you run by hand. Four decisions shaped it:

It has no dependencies

The entire thing uses only what ships with Python. Nothing to install, nothing that can break because an unrelated package was upgraded six months from now. An automated test fails the build if anyone adds a dependency.

Silence had to be distinguishable from failure

Since the normal output is "nothing to report", a broken system and a quiet market look identical. A second, independent process therefore watches the first and raises an alarm if no report appears for two working days.

The financial data is full of traps, and each one gave a wrong answer first

WHAT WENT WRONG	THE CONSEQUENCE BEFORE IT WAS CAUGHT
Only counting cash, not investments	Understated one company's finances six-fold , nearly rejecting a company with three years of funding
A company quietly stopped using a data field	Returned a six-year-old figure of \$24.9M for a firm holding \$816M
A company's own filed figure was wrong	Produced a nonsensical result of "15,401%"
A one-off payment read as ongoing spending	Implied a funding crisis at a profitable company

All are now guarded, and each guard has a test. The general lesson applies far beyond finance: **implausible data is far more often a broken feed than a broken reality** — and a system worth trusting proves which before acting.

Changes are verified, not hoped for

Restructuring the most complex part meant capturing the exact output for all 59 companies beforehand, making the change, and confirming the new output was byte-for-byte identical. Not "it looks fine". Identical.

WHAT THIS IS NOT

It is not financial advice, it places no orders, and it is not a product. It is a research tool built for one person, published openly so the reasoning can be examined. Small-company biotechnology has near-binary outcomes and any individual position can go to zero.

What I do

This project is a fair sample of how I work: find the question that actually matters, build something that answers it honestly, measure whether it worked, and report the result even when the result is inconvenient.

The domain here happens to be financial data. The same approach applies to any problem where information arrives constantly, most of it is noise, and the expensive mistake is acting confidently on something nobody checked — monitoring and alerting, data pipelines, automated research and reporting, or replacing a manual review process that has quietly become too big to do by hand.

Have a problem shaped like this?

If you have a process that runs on someone reading the same sources every day, or a system whose output nobody quite trusts, that is the kind of thing I build.

sved.net

The complete source, including the unflattering results:

github.com/svedbg/pharma — MIT licensed

Written by Svetoslav Rankov (sved) · sved.net · The system described here is a personal research tool. Nothing in this document is financial advice or a recommendation to buy or sell any security. Figures are drawn from the project's own measured results and are historical; they carry known biases, including a rising market during the measurement period, and do not indicate future performance. Trading small-company biotechnology stocks carries a genuine risk of total loss.