

Do Hard-to-Trade Stocks Actually Pay Off?

By Yonatan Brunshtein · September 17, 2026

The short version

I spent several weeks testing a simple question on the TSX Venture Exchange — Canada's market for small, early-stage, mostly resource and mining companies: does a stock being hard to trade tell you anything about what it's going to do next? I tested five different ways of measuring “hard to trade,” on 1,127 companies, over nine years, and split my data so I couldn't fool myself. Four of the five measures said the same thing: the hardest-to-trade stocks did meaningfully better afterward than the easiest-to-trade ones. One measure said the opposite. I think I know why, and the reason turned out to be the more interesting finding.

Terms you'll need

I use a handful of finance and statistics terms repeatedly below. I'm defining all of them here, in one place, so the rest of this piece can just use the words without re-explaining them every time. Skip this section if you already know the jargon; come back to it if something later doesn't make sense.

Liquid / illiquid — A stock is “liquid” if you can buy or sell a meaningful amount of it quickly, without moving the price much. It's “illiquid” if trading is thin — maybe it goes days without a trade, or a modest order swings the price 10%. Illiquid stocks are usually small, unloved, or both.

Excess return — How much a stock made or lost, minus how much the overall market did over the same period. If a stock rose 20% while the market rose 5%, its excess return is +15 percentage points. This is the number I'm trying to predict throughout this piece — not “did the stock go up,” but “did it beat the market.”

The benchmark — The “market” I measure everything against is the S&P/TSX Venture Composite Index — essentially the average of every stock on this exchange. It's the yardstick.

Quintile — Split a group into five equal-sized buckets, ranked from lowest to highest on some measure. I constantly split my 1,127 companies into five illiquidity quintiles: quintile 1 is the most liquid fifth, quintile 5 is the most illiquid fifth.

“Q5 minus Q1” / the spread — The single number I care about most: take the average (or median) return of the most-illiquid group (quintile 5) and subtract the return of the most-liquid group (quintile 1). A positive number means illiquid stocks did better; negative means they did worse.

Median vs. average (mean) — The average adds everything up and divides by the count — one huge winner can drag it way up. The median is the middle value if you lined every result up in order — one huge winner barely moves it. When results are lopsided (a few massive winners, many small losers), the median tells you what happened to a “typical” stock; the average tells you what happened to the total pile of money. Both matter, but they can tell very different stories, and mine do — see the section on this below.

Statistical significance / t-statistic — A way of asking “could this result just be random noise?” The bigger the t-statistic (in either direction, positive or negative), the less likely the pattern is a fluke. As a rough rule of thumb, a t-statistic bigger than about 2 (positive or negative) is generally considered “real.” Most of my key results have t-statistics between 5 and 9, which is a strong signal by normal standards.

Out-of-sample test / holdout — The single most important safeguard in this whole project. Before I looked at a single result, I randomly set aside 30% of my 1,127 companies (380 of them) and locked that list away. I only used the other 70% while building and checking my approach. At the very end, I ran the finished method on the locked-away 30% exactly once. If a pattern only existed because I'd subconsciously tuned it to fit — a very easy trap when you're poking at the same data over and over — it should fall apart on the untouched group. It didn't. That's the difference between a real finding and a lucky-looking coincidence.

Winsorizing — Capping extreme outliers at a fixed percentile before analyzing, so one 3,000%-return stock doesn't single-handedly decide the answer. I capped everything at the 1st-to-99th-percentile range, a decision I made before running any numbers.

Survivorship bias — A distortion that creeps in when your data only includes things that are still around today. My list of companies is every stock currently listed on this exchange — which means every company that went bankrupt, got delisted, or quietly disappeared between 2016 and 2024 is missing. Since failed companies are usually the most illiquid ones right before they die, this tends to make illiquid stocks look better in my data than they really were. I come back to this a lot — it's the biggest weakness in the whole study, and I don't try to hide it.

Three specific ways of measuring “how illiquid” — these are the actual yardsticks I built the whole study around:

Measure 1 — the Amihud ratio — For each stock, I look at how much its price moved on an average day relative to how many dollars of it actually traded that day. A stock that swings 5% on just \$2,000 of trading is far more illiquid than one that swings 5% on \$2 million. This is the most widely used illiquidity measure in academic finance, first proposed by Yakov Amihud in 2002.

Measure 2 — zero-return days — The percentage of trading days where the closing price didn't move at all. A stock that closes unchanged 60% of the time simply isn't trading — nobody's showing up to buy or sell it.

Measure 3 — the Corwin-Schultz spread — An estimate of the bid-ask spread (the gap between what buyers offer and what sellers ask) built entirely out of each day's high and low price. I don't have access to real historical bid-ask data for these stocks, but I do have daily highs and lows, and a clever 2012 method by Corwin and Schultz can back out a spread estimate from just that.

I also used two supporting measures — how much of a stock's total shares trade in a day (“turnover”), and plain dollar trading volume — as a further check on measures 1 and 2.

The question I started with

In finance there are two completely opposite, both completely reasonable, stories about what it means when a stock is hard to trade.

Story one: if you own something that's hard to sell, you deserve to be paid extra for that inconvenience. So illiquid stocks should, over time, drift upward relative to liquid ones — a kind of compensation for the hassle and risk of being stuck holding them. Economists call this the illiquidity premium.

Story two: a stock that's hard to trade is usually hard to trade for a reason — nobody wants it. Thin trading is often the last symptom before a company runs out of money, gets diluted into oblivion by a desperate financing, or quietly delists. In this story, illiquid stocks should do worse, not better — the market is slowly recognizing they're in trouble.

Both of these are argued seriously in the academic literature, and which one wins isn't obvious — it depends on the market. I decided to actually test it, rather than assume an answer, on the single market where I thought the question would be sharpest: the TSX Venture Exchange.

Why the TSX Venture Exchange

The TSX Venture Exchange (TSXV) is Canada's exchange for small, early-stage companies — mostly mining and resource exploration, plus smaller pockets of oil and gas, technology, biotech, clean energy, and everything else. The typical company in my final dataset has a market value around C\$18 million and trades around 27 cents a share. If an illiquidity effect exists anywhere, it should show up loudly here — this is about as far down the size-and-liquidity spectrum as a regulated stock exchange goes. It's also a market that, as far as I could find, nobody has studied this exact way before.

The rule I made myself follow before looking at a single result

Here's the trap with a question like this: if you have thousands of stocks and a lot of flexibility in how you slice the data, you can always find some way of measuring “illiquid” that produces an exciting-looking pattern. That's not a discovery — it's just noise you kept digging until you found. It happens by accident more often than people admit.

So before I ran a single number, I wrote down every decision in advance: exactly how I'd measure illiquidity (three separate ways, defined above, formulas fixed), exactly which time windows I'd test (6, 12, and 24 months forward), exactly how I'd handle extreme outliers (capped at the 1st and 99th percentile), and — most importantly — I randomly locked away 30% of my companies (380 of the 1,267) before touching anything. I only ever looked at results on the other 70% while I was building and checking my method. I opened the locked-away 30% exactly once, at the very end, to see if what I'd found held up. Every headline number in this piece is from that locked-away group.

Where the data came from

This part is less exciting but matters for trusting the result. Reliable historical price data for tiny Canadian stocks is hard to get — the usual sources either don't cover this exchange, don't go back far enough, or don't handle stock consolidations properly (Venture-exchange companies routinely do 5-for-1 or 10-for-1 reverse splits, and if your price data doesn't adjust for that, it looks like the stock randomly jumped 500% overnight, which is nonsense).

I ended up building my company list from the exchange's own public directory (1,267 actual operating companies, after stripping out shell companies, warrants, and other non-stock instruments), and pulling daily prices, sector classification, and each company's share count from TMX's own market-data service — the same organization that runs the exchange. I cross-checked the price data for accuracy and ran automated checks that caught and fixed a handful of corrupted price records before any analysis began.

One data point I could not get historically: how many shares each company had outstanding at each point back in 2016–2024 (needed to size-adjust the comparison, since bigger companies behave differently than tiny ones). I used each company's current share count as an approximation, which I flag honestly, and I double-checked the finding using stock price alone as an alternative size measure — same answer either way.

What I found, part one: illiquid stocks did better

I sorted all 1,127 stocks with usable data into five buckets (quintiles) by how illiquid they were, at each of 108 monthly snapshots between 2016 and 2024, using the Amihud measure. Then I looked at what happened over the following 6, 12, and 24 months. On the locked-away 30% of companies I never touched while building this:

Table 1. The most-illiquid fifth of companies, minus the most-liquid fifth, in average return over the next N months. Positive = illiquid stocks won.

Measure	Months ahead	Illiquid group's edge	Signal strength (t-stat)
Amihud measure	6	+0.32	+9.19
Amihud measure	12	+0.65	+5.94
Amihud measure	24	+1.06	+5.83
Zero-return-days measure	6	+0.26	+6.25
Zero-return-days measure	12	+0.53	+4.40
Zero-return-days measure	24	+0.76	+4.76
Corwin–Schultz spread	6	-0.03	-0.58
Corwin–Schultz spread	12	+0.01	+0.07
Corwin–Schultz spread	24	+0.05	+0.34

By the Amihud measure, the most illiquid fifth of stocks beat the most liquid fifth by 32 percentage points over six months, 65 points over a year, and 106 points over two years. The signal strength on all of those is well above the “this is probably real” threshold of 2. The zero-return-days measure — a completely different way of defining “illiquid” — pointed the same direction, with similar strength. And this wasn't a fluke of the 70% of companies I built the method on: it showed up, just as strongly, on the 30% I'd locked away and never looked at until the end.

Figure 1. Forward excess return by illiquidity quintile (holdout sample)




Figure 1. Typical (median) 12-month return, relative to the market, for each of the five illiquidity buckets — bucket 1 is the most liquid, bucket 5 the most illiquid — on the locked-away 30% of companies. For the Amihud and zero-return measures the bars climb steadily left to right. For Corwin–Schultz they don't.

What I find most convincing here isn't just that bucket 5 beat bucket 1 — it's that the climb from bucket 1 to bucket 5 is smooth and step-wise (Figure 1). It's not that one weird group of stocks dragged the whole thing; each step up the illiquidity ladder adds a bit more return. That's what a real, gradual relationship looks like, as opposed to a coincidence driven by a handful of stocks.

Why am I suddenly talking about the median instead of the average?

Here's something I have to be upfront about, because it changes how big the “real” effect looks. Table 1 above uses averages. Averages on this kind of data are dangerously easy to misread.

The typical stock in my sample, over any 12-month stretch, actually did worse than the market — the median 12-month result was -13 percentage points versus the market. Meanwhile the average 12-month result was positive, at +10 points. Both of those are true at the same time, because a small number of stocks went up 300%, 500%, even 900%, and those extreme winners pull the average way up while barely nudging the median. Figure 2 shows this directly — most of the distribution sits to the left of zero (below-market), but there's a long thin tail stretching far to the right.

Figure 3. Forward excess returns are heavily right-skewed — mean $\neq$ median




Figure 2. How 6-, 12-, and 24-month returns (relative to the market) are actually distributed. The dashed line is the median (typical outcome); the solid line is the average. They point in opposite directions because of the long right-hand tail of big winners.

And here's the catch that matters for this whole study: the most illiquid stocks are also, almost by definition, the cheapest and most speculative ones — which means they're exactly where those rare 500%+ winners tend to live. So an average-based comparison between “illiquid” and “liquid” stocks is partly just measuring which bucket has the fatter jackpot tail, not what happened to a typical stock in each bucket.

So, throughout the rest of this piece, I lean on the median comparison — “what happened to the typical illiquid stock vs. the typical liquid stock” — as my main number, and I show the average alongside it for context. On the median basis, the Amihud gap is smaller but still large and still statistically solid: about 24 points over six months, 43 over a year, 67 over two years. Smaller than the headline average numbers, but still a real, large, statistically strong gap.

What I found, part two: one measure disagreed, and that's the interesting part

If I'd stopped after the Amihud measure, I'd have a tidy, one-sided story: illiquid stocks win. But I'd pre-committed to testing three different definitions of illiquidity, specifically because I didn't want to trust just one, and the third one — the Corwin–Schultz spread estimate — told a different story.

On the same locked-away 30% of companies, using the same “typical stock” (median) comparison, the widest-spread stocks actually did worse than the narrowest-spread ones — by about 10 points over six months and 12 points over a year. Smaller in size than the Amihud effect, and less consistently statistically significant, but leaning the other way.

So which is it? I went looking for why two supposedly-related “illiquidity” measures would point in opposite directions, and found something worth knowing on its own: in this specific market, a stock's Corwin–Schultz spread estimate is actually negatively related to how illiquid it is by the other two measures. That sounds backwards, but here's the mechanism. The Corwin–Schultz method infers the spread from how wide a stock's daily high-low range is. A stock that barely trades often prints almost the same price all day — a narrow range, which the formula reads as a tight, cheap-to-trade spread. A stock that trades actively but swings wildly within the day — which, on this exchange, usually means a hot, speculative, momentum-driven name — prints a wide range, which the formula reads as an expensive, wide spread. In other words: on the TSX Venture Exchange, this particular measure ends up tracking day-to-day price swings (volatility) more than it tracks true trading cost. And it's true, separately, that hyped-up volatile stocks tend to cool off afterward — that's a well-documented pattern in its own right. I think that's largely what the Corwin–Schultz result is picking up.




Figure 3. All five ways I measured illiquidity, and what each one said about the 12-month return gap between the most- and least-illiquid stocks (locked-away 30% only). Four measures agree; one doesn't.

Making sure this is real and not a trick of the data

A result this size on penny stocks deserves real skepticism, so I spent a good chunk of this project trying to break my own finding. Here's the full stress test — every alternative version of the 12-month Amihud result I tried, each one designed to rule out one specific way the result could be fake. The pattern is the same at 6 and 24 months; I'm showing 12 months because it's the middle case.

Table 2. Every version of the 12-month Amihud test I ran, on the locked-away 30%.

Version of the test	Illiquid – liquid gap (12mo)	Signal strength
Main version (plain average)	+0.65	+5.94
Median within bucket	+0.43	+6.56
Weighted by company size	+0.54	+4.52
Using 3-day average prices	+0.63	+6.13
Only stocks priced $\geq$ C\$0.10	+0.32	+7.54
Tighter outlier caps	+0.45	+9.56
Median + priced $\geq$ C\$0.10 (strictest)	+0.26	+6.40

- **Median instead of average:** +0.43. The skew-resistant number I lean on throughout this piece.
- **Weighted by company size instead of treated equally:** +0.54. The effect isn't confined to the very smallest nano-caps inside the illiquid bucket.
- **Using 3-day average prices at both ends of the window:** +0.63. Rules out the result being a fluke of one lucky or unlucky closing price.
- **Dropping every stock priced under 10 cents:** +0.32. Roughly halves the gap but leaves it large and clearly significant — it's not purely a sub-dime phenomenon.
- **Capping outliers much more tightly:** +0.45. Clipping the extreme winners hard doesn't make the effect disappear.
- **The two strictest checks together (median, only stocks over a dime):** +0.26. Smaller, but still standing.

Is it just a “cheap stocks bounce back” effect in disguise?

The Amihud measure divides by trading dollars, and illiquid stocks tend to be cheap stocks, so there's a real risk that what looks like an “illiquidity effect” is actually just “beaten-down penny stocks tend to bounce” — a completely different, well-known pattern that has nothing to do with liquidity. I tested this directly: I split companies into three price groups (cheap, medium, expensive) and, separately, into the five illiquidity buckets, and looked at the illiquidity gap within each price group on its own.




Figure 4. The illiquid-minus-liquid return gap, calculated separately within each price group. If this were just a “cheap stock” effect, the gap should vanish among the pricier stocks. It doesn't.

Among the cheapest third of stocks, the illiquid ones still beat the liquid ones by about 36 points. But among the priciest third — stocks that are not penny stocks by any definition — illiquid ones beat liquid ones by an even larger 62 points. That rules out “it's just cheap stocks” as the explanation — if anything, the effect is stronger among the more expensive names. (The middle price group came out close to a wash, which I don't have a tidy explanation for and am flagging rather than glossing over.)

Is it just a bet on the mining sector?

Two-thirds of this exchange is mining and resource-exploration companies, so a natural worry is that I've just discovered “illiquid mining stocks did well,” dressed up as a general liquidity finding. I re-ran the whole sort comparing companies only against others in their own sector — illiquid mining stocks vs. liquid mining stocks, illiquid tech stocks vs. liquid tech stocks, and so on. The result barely moved (41 points at 12 months, versus roughly the same number in the sector-blind version). It's not a sector bet.

Does controlling for company size change the answer?

I ran a regression — a statistical technique that isolates the effect of one thing while holding others constant — controlling for both company size and sector at the same time, separately at each of my 108 monthly snapshots, then averaged. The illiquidity effect held up as a strong, independent factor even after accounting for size and sector; smaller companies also did better on their own (a well-known pattern in finance in its own right), but illiquidity wasn't simply standing in for “small.”

Table 3. What's left of the illiquidity effect, and the size effect, once I hold sector and company size constant (locked-away 30%). “Illiquidity effect” is what moving a company from the most-liquid to the most-illiquid ranking adds to its return; “size effect” is the same for market value.

Measure	Months ahead	Illiquidity effect	Signal	Size effect	Signal
Amihud measure	6	+0.19	+5.76	-0.04	-3.78
Amihud measure	12	+0.39	+7.67	-0.09	-2.70
Amihud measure	24	+0.64	+6.95	-0.17	-2.49
Zero-return-days measure	6	+0.15	+2.94	-0.06	-6.18
Zero-return-days measure	12	+0.32	+2.63	-0.11	-5.17
Zero-return-days measure	24	+0.20	+1.72	-0.25	-4.44
Corwin–Schultz spread	6	-0.01	-0.13	-0.07	-6.72
Corwin–Schultz spread	12	+0.10	+0.80	-0.14	-4.45
Corwin–Schultz spread	24	+0.21	+1.39	-0.27	-4.01

Is illiquidity even a stable trait, or does it change randomly month to month?

This mattered to me because if a stock randomly jumped between “liquid” and “illiquid” from one month to the next, a quintile sort wouldn’t mean much. It doesn’t jump around: a stock in the most-liquid group one month has about a 96% chance of still being there the next month, and a stock in the most-illiquid group has about a 94% chance of staying put. Illiquidity is a slow-moving, sticky characteristic of a company — closer to its size or its sector than to something that changes on a whim — which is exactly what makes it reasonable to sort on in the first place.




Figure 5. If a stock is in a given illiquidity group this month, here’s the chance it’s in each group next month. The heavy diagonal shows illiquidity barely moves month to month.

Does it only work during the crazy speculative years?

TSX Venture had two genuinely wild speculative runs in my sample window — a rally starting in 2016 and another during the 2020–2021 pandemic-era retail-trading boom. I worried the whole result might just be “junk stocks rip during manias” dressed up as a liquidity story. Splitting the data into three multi-year chunks, the effect is indeed largest during those two periods, but it doesn't disappear in the calmer 2022–2024 stretch — it shrinks by roughly half but stays clearly positive and statistically solid (21 points). So there's likely a real, smaller, year-in-year-out effect, with a much bigger boost layered on top of it during speculative frenzies.

Table 4. The 12-month illiquid-minus-liquid gap, split into three-year chunks (locked-away 30%). The Amihud gap shrinks but survives the calmer years; the Corwin–Schultz gap vanishes exactly during the wildest speculative period.

Measure	2016–2018	2019–2021	2022–2024
Amihud measure	+0.50	+0.52	+0.21
Zero-return-days measure	+0.34	+0.46	+0.16
Corwin–Schultz spread	-0.24	+0.03	-0.16




Figure 6. The 12-month illiquid-minus-liquid gap, broken into three time chunks. Biggest during the speculative booms, smaller but still positive afterward.




Figure 7. The same gap, month by month, for all three measures. The two big peaks line up exactly with the 2016 and 2020–2021 speculative rallies.

The weak spot: companies that didn't survive to be counted

This is the part of the project I'm least comfortable glossing over, so I'm giving it its own section rather than a footnote.

My list of 1,267 companies is every stock currently trading on the TSX Venture Exchange, as of when I pulled the list. That sounds thorough, and in one sense it is — it's the entire exchange, not some hand-picked subset. But it has a specific, unavoidable gap: any company that went bankrupt, got delisted, or otherwise disappeared at any point between 2016 and 2024 is simply not in my list, because it doesn't exist today. And on this particular exchange, where a meaningful share of listed companies are speculative exploration ventures that never find anything, that missing group is genuinely large — plausibly comparable in size to the group of survivors I actually studied.

I looked for a way to plug this gap — a historical snapshot of exactly which companies were listed at each point in time, or a comprehensive list of everything that delisted and when. I couldn't find one that's freely available. I found one partial third-party list, but it only covered 2022 onward and was heavily skewed toward companies that got bought out (a good outcome, not a bad one), so using it would have introduced a different, arguably worse bias. I chose not to use it rather than paper over one problem with another.

Here's why this matters for interpreting my results, and the direction is not ambiguous: companies that fail are disproportionately illiquid right before they die — that's often the last visible symptom. By construction, my “most illiquid” bucket is missing its worst-performing members: the ones that went to zero and vanished from the exchange entirely. That means my illiquid-stocks-do-better finding is very likely an overstatement of the true effect. I ran the numbers on a rough worst-case: if roughly a fifth of the most-illiquid group were secretly complete wipeouts that I couldn't see, it would erase about half of the 12-month gap I measured. A wipeout rate around 30% would erase most of it. A genuine effect likely survives realistic assumptions, but I can't tell you the exact true size, and neither can anyone else without better data.

One thing that partly reassures me: I checked separately whether companies that stopped trading partway through my own sample (as opposed to ones missing entirely from the start) were driving the result, and they weren't — re-scoring those handful of companies under a worst-case “total loss” assumption barely changed anything. So the part of survivorship bias I could measure turned out to be small; it's the part I couldn't measure — companies missing from the list altogether — that I have to flag and leave unresolved.

There's one silver lining buried in this problem. This missing-company bias should push every one of my illiquidity measures toward making illiquid stocks look better than they really are — including the Corwin–Schultz spread measure. And yet the Corwin–Schultz measure still came out negative, against that headwind. That actually makes me trust the Corwin–Schultz result more, not less: it's the one signal that shows up despite a bias working against it.

So what does this actually mean?

Two different kinds of “hard to trade”

Stepping back, I think the cleanest way to make sense of everything above is that my five measures split into two families that are measuring genuinely different things. The Amihud measure, the zero-return-days measure, turnover, and dollar volume are all asking some version of “how much does this stock actually trade?” The Corwin–Schultz measure is asking “how wide is the price gap between buying and selling it?” On most markets those two questions have similar answers. On the TSX Venture Exchange, in this period, they didn't — a stock could trade rarely but calmly (which reads as illiquid on the first four measures and as tight-spread on the fifth), or trade often but wildly (the reverse). A study that only used one measure and called it a day would have reported a confident, one-directional answer that happened to be a coin-flip away from the opposite conclusion.

Is this a real risk premium, or something else?

A genuine “you get paid for holding something illiquid” risk premium should show up fairly steadily over time. What I actually see (Figure 7) is a modest, fairly consistent gap most months, with two enormous spikes exactly during the 2016 and 2020–2021 speculative manias — periods when the cheapest, most obscure, hardest-to-trade names on the whole exchange were the ones getting bid up hardest by momentum and retail speculation. My honest read is that there's probably a real, modest illiquidity effect underneath all of this, but a much larger share of the headline number is something closer to “illiquid stocks have outsized exposure to speculative market moods,” which is a different — and less flattering — story than a steady risk premium.

This is not a trading strategy

I want to be direct about this because it's the most likely way this piece could be misread. The “most illiquid fifth” of my sample is, almost by definition, the group of stocks where a typical day's trading is a few thousand dollars. You cannot actually buy a meaningful position in that basket without your own buying pushing the price up and erasing the opportunity — the size of the effect and the size of the money you could put behind it are working directly against each other. This is a description of how returns are spread across a liquidity spectrum on one specific exchange, not a strategy anyone can execute.

What I'd want to test next

The natural next question, and the one I think is actually more important than anything above: do the most illiquid companies pay for it later, specifically when they need to raise money? Small resource companies constantly issue new shares to fund exploration, and if the most illiquid, least-trusted names have to raise that money on much worse terms — deeper discounts, more free warrants thrown in, more dilution per dollar raised — then the “value trap” story and the “illiquid stocks did better” story aren't actually in conflict. They'd be two stages of the same mechanism: illiquid companies get punished on financing terms, existing shareholders get diluted hard, and the stock's subsequent price move reflects a mix of that dilution and whatever speculative attention the company manages to attract. That's a

smaller, more hand-collected study — maybe 30-40 companies from the extreme ends of my illiquidity ranking — but it's the piece that would tell me whether I've found a genuine market quirk or just re-discovered how badly junior mining financing treats its most desperate participants.

The honest limitations list

- **Survivorship bias, unresolved:** explained above at length — the biggest open issue, and the reason I treat my “illiquid stocks win” numbers as an upper bound, not a precise estimate.
- **Company size is approximated:** I used each company's current share count applied to all historical dates, because I couldn't get historical share counts for free. I checked the result against stock price alone as an alternative size measure and got the same answer, which helps, but it's still an approximation worth knowing about.
- **This is one exchange, one time period:** 2016–2024, TSX Venture only. I have no idea whether this generalizes to other junior exchanges, other countries, or other decades, and I'm not claiming it does.
- **Correlation, not causation:** I've shown that illiquidity and future returns are related in a way that survives a lot of stress-testing. I have not identified the exact mechanism causing it, beyond the speculative-mood theory above, which is a reasonable read of the evidence, not a proven cause.
- **The middle-price-tercile result was a wash:** flagged honestly in the double-sort section rather than quietly dropped.

How this compares to what's already known, and one claim I'm walking back

The direction of my main result — illiquid stocks earning more — lines up with what the broader academic literature has found on larger, more liquid markets since Amihud and Mendelson's original 1986 paper, and with Amihud's own follow-up work extending it internationally. What's different here isn't the sign, it's the size and the circumstances: on established markets this shows up as a modest, steady premium; on TSXV it shows up as something much larger, much lumpier, and heavily concentrated in speculative booms — which is why I read it as mostly sentiment with a smaller genuine premium underneath, rather than a clean version of the same textbook effect.

I also want to walk back one phrase I've used casually elsewhere in describing this project: calling it “the cleanest test” of this relationship. It's a clean test in the sense that mattered most to me — the direction wasn't assumed, the holdout was real, and the specification was locked before I saw a result. It is not the cleanest test in any absolute sense: the survivorship gap is real and unresolved, the size control is an approximation, and nobody besides me has yet reproduced these numbers from scratch. “Disciplined” is the word I'd defend. “Cleanest” oversells it.

Where this leaves me

Testing this properly — with the direction left genuinely open, a chunk of the data locked away until the end, and every decision made before I could see how it would turn out — the honest answer is more interesting than either clean story I started with. On the TSX Venture Exchange, between 2016 and 2024, stocks that were hard to trade in the sense of “barely anyone shows up to trade them” went on to meaningfully outperform stocks that traded easily — a pattern that held up on data I never touched while building the method, survived every stress test I could throw at it, and shrank but didn't disappear outside the speculative booms. Stocks that were “hard to trade” in the sense of “wide gap between buy

and sell prices” went the other way, weakly. The size of the first effect is very likely inflated by the companies that failed and vanished before I could count them. And none of it is a strategy you could actually run, because the stocks in question are too thin to buy at any scale. What I'm left with is a clean, honestly-caveated picture of how illiquidity and returns actually relate at the very bottom of the market — plus a much better next question than the one I started with.

What this builds on

This work draws directly on a handful of foundational papers in market microstructure — the field that studies how trading actually happens, not just prices in the abstract:

- **Amihud & Mendelson (1986)** — established the core logic that illiquidity should be compensated with higher returns.
- **Amihud (2002)** — built the specific price-impact ratio I use as my primary illiquidity measure.
- **Amihud, Hameed, Kang & Zhang (2015)** — extended that illiquidity-premium evidence internationally.
- **Corwin & Schultz (2012)** — developed the high-low spread estimator I use as my fifth, dissenting measure.
- **Datar, Naik & Radcliffe (1998)** — established share turnover as a liquidity proxy.
- **Lesmond, Ogden & Trzcinka (1999)** — developed the zero-return-days measure.
- **Fama & MacBeth (1973)** — developed the regression method I use to control for size and sector.
- **Newey & West (1987)** — developed the statistical adjustment I use to handle overlapping return windows.
- **Pástor & Stambaugh (2003) and Petersen (2009)** — inform how liquidity risk and panel-data statistics are generally handled in this literature.

How this can be checked and reproduced

I built this as a deterministic pipeline: every parameter lives in one configuration file, and nothing downstream is hard-coded. The steps run in order — build the universe, pull daily prices and the benchmark, calculate the illiquidity measures, freeze the 70/30 holdout split (written once; the pipeline is built to refuse to overwrite it), calculate forward returns, run the portfolio sorts and regressions, run the full stress-test battery, and generate every figure and table in this piece. Re-running it from the same cached inputs reproduces every number in this document exactly — which is also how I caught and fixed a couple of small hand-transcription slips while assembling this plain-language edition.

- Monthly snapshots from January 2016 to December 2024 (108 dates).
- A 252-trading-day trailing window for every illiquidity measure, with a 126-day minimum to be scored.
- 6/12/24-month forward horizons.
- Outlier capping at the 1st/99th percentile within each snapshot-and-horizon group.
- The 30% holdout — 887 training, 380 holdout — from a fixed random seed, written once and never re-rolled.
- Quintile cutoffs computed from the training group only, at every single snapshot.

Data-quality screens

Two automated screens run before any metric is calculated. First, any company whose typical closing price exceeds C\$50 is dropped as a likely units artifact or a stack of unadjusted consolidations (three specific cases in my data: one stock quoted near C\$1,365, one near C\$135, one near C\$94). Second, within each remaining price series, an isolated one-day move greater than roughly 5x in either direction that reverses by more than 60% the next day is treated as a bad data print and smoothed out; genuine multi-day moves and real consolidations, which don't reverse, are left untouched. Of the 1,267 cleaned companies, 1,167 returned usable price history, and 1,127 had enough of it to be scored on at least one snapshot. My main results are unchanged whether or not these two screens are applied.

Where I deviated from my original plan, and why

Four things changed from my original methodology document, each forced by data access, and each logged here rather than quietly changed:

- **Price data source:** I originally planned to use Alpha Vantage; its adjusted-price data turned out to be premium-only, and my fallback, Yahoo Finance, rate-limited my access. I switched to TMX Money's own data service.
- **Shares outstanding:** I'd originally planned to collect this by hand, ticker by ticker. TMX Money's service returns current shares outstanding automatically for essentially every active company, so I automated it instead, with the limitation discussed above.
- **Sample size:** my original plan targeted a hand-picked sample of 150–300 companies — a cap that existed only to make the manual shares-outstanding work manageable. Since that step became automatic, I used the entire cleaned universe of roughly 1,130 companies instead, which removes sample-selection judgment calls entirely.

- **Sector classification:** comes from TMX Money's own industry mapping, matched to the exchange's six sector buckets.

Everything else — the research question itself, leaving the direction of the effect genuinely open, the three pre-registered illiquidity formulas, the portfolio-sort and regression methods, the 70/30 holdout discipline, the outlier-capping rule, the look-ahead protection, and treating survivorship as a disclosed limitation rather than something to paper over — stayed exactly as originally planned.

A note on how I built this

Claude (Anthropic) assisted with the data collection and analysis code, statistical structure, and figure generation, and cleaned up language and grammar in this paper edition, which I, the author, wrote in full. All research decisions — what to measure, how to define it, the interpretation of results, and the framing of the contribution were my discretion; I am responsible for the content in full.