


Phantom Exits: The Secondary Market Illusion and the Quest for Liquidity In Healthcare Angel Investing

OCT 24, 2025 • PAID



Share

DISCLAIMER: The views and opinions expressed in this essay are solely my own and reflect the positions, strategies, or opinions of my employer.

ABSTRACT

Angel investors in healthcare startups face a distinctive challenge that distinguishes them from investors in other venture categories: the extended timeline to liquidity, coupled with the binary nature of healthcare outcomes creates portfolios that are simultaneously illiquid and difficult to value. This essay examines the emerging secondary market infrastructure for private healthcare companies, the mechanisms which sophisticated investors are engineering liquidity in otherwise frozen positions, and the fundamental tension between price discovery and information asymmetry in healthcare venture investments. Through analysis of market structure, transaction mechanics, and the unique challenges of healthcare venture secondaries, this work argues that the promise of liquid secondary markets for healthcare angel positions is largely illusory, and that investors who orient their strategy around secondary liquidity are optimizing for the wrong variables. The central thesis holds that liquidity engineering in healthcare angel investing is less about creating functional markets and more about managing the psychological and portfolio construction challenges of decade-long hold periods in an asset class characterized by high mortality and long gestation.

TABLE OF CONTENTS

- Introduction: The Liquidity Fantasy
- The Time-Value Problem in Healthcare Angels
- Secondary Market Infrastructure and Its Limitations
- Information Asymmetry as Market Failure
- Price Discovery in Illiquid Markets
- Liquidity Engineering: Mechanisms and Realities
- The Portfolio Construction Challenge
- SPVs, Rolling Funds, and Structural Solutions
- Why Healthcare Is Different
- Conclusion: Building for Illiquidity

Introduction: The Liquidity Fantasy

There exists within the venture capital ecosystem a persistent fantasy that appeals to the part of our brain that wishes reality were more convenient than it actually is. It is the fantasy of the liquid secondary market for private company equity, where investors can trade positions in pre-exit startups with the same ease and price transparency that characterizes public equity markets. The fantasy is particularly seductive in healthcare angel investing, where the pathway from seed investment to liquidity event can stretch across ten to fifteen years, where binary regulatory outcomes can shift company valuations by orders of magnitude overnight, and where the opacity of clinical development makes it nearly impossible for outside investors to accurately assess company progress. The promise of secondary liquidity suggests that investors need not wait a decade for an exit, need not suffer through the anxiety of holding illiquid positions through multiple market cycles, and can actively manage

their portfolios by trimming winners, cutting losers, and rebalancing exposure as information emerges. This promise is largely false, or at minimum severely overpromised, and investors who construct their healthcare angel strategy around the assumption that secondary liquidity is readily available are setting themselves up for disappointment.

The reality of secondary markets for healthcare angel positions is that they are typically inefficient, characterized by extreme information asymmetry, and liquidity is often available when you do not want them and unavailable when you do. Liquidity appears when companies are hot and trading at inflated valuations where selling generates minimal returns or opportunity cost. Liquidity disappears when companies hit rough patches when clinical trials fail, when reimbursement proves elusive, or when the broader market sentiment shifts against digital health or biotech. The investors who successfully transact in secondary markets tend to be either insiders with superior information looking to reduce exposure before bad news becomes public, or outsiders who are systematically overpaying for positions in companies they do not fully understand. The bid-ask spreads are wide, the transaction costs are high, and the friction involved in finding counterparties, negotiating terms, and executing transactions often makes the juice not worth the squeeze for positions below several hundred thousand dollars in current value.

Yet despite these limitations, the secondary market for private healthcare company equity has grown substantially over the past decade, driven by several converging forces. The explosion of angel and micro-VC investing has created a large population of investors holding small positions in many companies with limited visibility into company progress and uncertain pathways to exit. The extension of time to IPO and the preference for companies to stay private longer has increased the duration of illiquidity that investors must endure. The rise of specialized secondary platform funds has created infrastructure where previously none existed. And the generational wealth transfer and increasing sophistication of high-net-worth individuals has created a class of investors who think about private company positions as assets to be actively managed rather than lottery tickets to be held until they pay off or expire worthless. Understanding how secondary markets actually function in healthcare venture, what mechanisms exist for engineering liquidity, and what the fundame

constraints are that prevent these markets from achieving anything approaching efficiency is essential for any serious healthcare angel investor.

The Time-Value Problem in Healthcare Angels

The core challenge that secondary markets purport to solve is the time-value problem inherent in healthcare angel investing. When an investor commits capital to a seed stage digital health or biotech company, they are typically looking at a minimum seven to ten years before a liquidity event, and in many cases significantly longer. A digital therapeutics company raising a seed round in twenty twenty-four might not achieve regulatory clearance until twenty twenty-six, might not secure reimbursement until twenty twenty-eight, might not demonstrate sufficient traction for a Series A until twenty thirty, might not grow into an attractive acquisition target until twenty thirty-two, and might not IPO or get acquired until twenty thirty-four or beyond. An investor who writes a fifty thousand dollar check in twenty twenty-four is making a commitment to have that capital locked up for a decade or more, during which time they will have no access to the appreciation in value, no ability to redeploy capital into better opportunities emerge, and significant uncertainty about whether the investment will ultimately succeed.

This extended duration creates several problems for portfolio construction and investor psychology. From a portfolio construction perspective, angel investors who are actively deploying capital into new opportunities every year will find that their portfolios become increasingly dominated by older vintage investments that have appreciated in paper value but cannot be monetized. An investor who deploys one hundred thousand dollars per year into angel investments will, after ten years, have a portfolio of one hundred investments assuming ten investments per year at ten thousand dollars each. If the portfolio performs well and the successful companies have marked up significantly in subsequent financing rounds, the investor might have several million dollars in paper value concentrated in a handful of positions. But paper value is essentially fictional until a liquidity event occurs. The investor cannot use it to make mortgage payments, cannot diversify into other asset classes, and

cannot even reliably borrow against it since private company equity is not accepted collateral by most lenders. The result is that successful angel investors often find themselves paper-rich and cash-poor, unable to realize the value they have created and increasingly constrained in their ability to make new investments.

From a psychological perspective, the extended illiquidity creates anxiety and regret that can lead to poor decision-making. When an investor watches a company they have invested in at a five million dollar valuation raise subsequent rounds at twenty million, fifty million, and one hundred million dollar valuations, they experience an emotional high of seeing their position appreciate. But this appreciation is theoretical until it is realized, and the path from one hundred million dollar private valuation to actual cash in hand is long and fraught with risk. Companies that reach one hundred million dollar valuations can still fail, can still be acquired for less than invested capital, can still languish in private markets for years unable to find an exit. The investor who is watching their position appreciate on paper is simultaneously aware that this paper value could evaporate, and this creates a powerful desire to somehow lock in gains, to crystallize at least some portion of the appreciation, to reduce exposure before something goes wrong. Secondary markets offer the promise of satisfying this desire, but the reality is more complicated.

The time-value problem is particularly acute in healthcare because the binary nature of many healthcare outcomes means that companies can go from seeming winners to total losses very quickly. A digital health company that has raised a Series B at an attractive valuation and appears to be on a strong trajectory can see its value crash if it fails to secure reimbursement from CMS or if a clinical trial comes back negative. A biotech company in Phase II trials can be worth hundreds of millions one day and worth close to zero the next if the trial fails to meet endpoints. This binary risk profile means that paper markups in healthcare can be especially illusory, and investors are overly focused on capturing secondary liquidity at marked-up valuations may be optimizing for exactly the wrong thing. The companies that are easy to sell in secondary markets are often the companies that have the most frothy valuations and the highest risk of disappointment. The companies that would actually be wise to avoid exposure to are often the ones where no secondary buyer can be found.

Secondary Market Infrastructure and Its Limitations

The infrastructure for trading private company equity has evolved significantly over the past fifteen years, moving from an entirely informal network of personal relationships and one-off transactions to a more structured ecosystem with dedicated platforms, specialized funds, and established practices. Understanding this infrastructure and its limitations is essential for healthcare angels considering secondary liquidity strategies.

The most visible components of the secondary market infrastructure are the platforms that attempt to create marketplaces for private company equity. Companies like Global, Hiive, and EquityZen have built technology platforms that allow shareholders to list positions for sale and potential buyers to browse available opportunities. These platforms provide some basic price discovery mechanisms, handle some of the logistics of share transfers, and attempt to create enough liquidity to make transactions feasible. For later-stage companies with significant employee equity and large valuations, these platforms can facilitate meaningful transaction volume. For early-stage healthcare companies, the platforms are far less useful. The companies listed tend to be technology unicorns or late-stage companies that employees want to diversify from, not seed-stage digital health or biotech companies that angels have invested in. The transaction sizes tend to be smaller than institutional buyers want and larger than most individual buyers can afford. And the information available about listed companies is typically minimal, making it difficult for potential buyers to conduct adequate diligence.

Beyond the platforms, there exists a growing number of specialized secondary funds that purchase positions in private companies from early investors, employees, or founders looking for liquidity. These funds range from large multi-billion dollar vehicles operated by firms like Lexington Partners and Coatue to smaller special funds focused on specific sectors or stages. For healthcare angels, the relevant secondary funds tend to be those that focus on venture-stage companies and have expertise in evaluating healthcare businesses. But even these specialized funds have

strong preferences about what they will buy. They want concentrated positions large enough to be meaningful, typically at least several hundred thousand dollars in current value. They want companies that have demonstrated some level of traction and de-risking, not seed-stage experiments. They want situations where they can access to meaningful information about company performance, ideally through conversations with management or through their own portfolio companies that interact with the target company. And they want to buy at a discount to the last financing round, often twenty to forty percent below the most recent preferred price.

These preferences create a structural mismatch between what healthcare angels want to sell and what secondary buyers want to buy. The typical angel investor has many small positions, most of which are in early-stage companies that have not yet meaningfully de-risked, and has limited information about company performance beyond what is shared in occasional investor updates. The secondary buyer wants large concentrated positions in later-stage companies where they can conduct thorough diligence and have conviction about future outcomes. The result is that most angel positions are simply not attractive to secondary buyers at any reasonable price. The positions that are attractive to secondary buyers are typically in companies that have become obvious winners and where selling means leaving significant value on the table. The positions that angels would most like to sell, companies that are struggling or where the path forward is uncertain, are exactly the positions that a rational secondary buyer will touch.

The mechanics of executing secondary transactions add further friction. Most private companies have right of first refusal provisions that allow the company and existing investors to block secondary sales or to purchase the shares themselves on the same terms as the proposed secondary transaction. Many companies actively discourage or prohibit secondary sales, particularly for early-stage investors, because they want to control their cap table and avoid having unknown entities as shareholders. The legal documentation for transferring shares can be complex and expensive, often requiring legal fees that are prohibitive for small transactions. And the tax implications of secondary sales can be unfavorable compared to holding until a qualified exit, particularly if the shares have appreciated significantly and the investor faces long-term capital gains tax.

term capital gains on the sale but would face potentially more favorable treatment in a merger or acquisition.

Information Asymmetry as Market Failure

The fundamental challenge that prevents secondary markets for private healthcare companies from achieving anything approaching efficiency is information asymmetry. Efficient markets require that buyers and sellers have access to similar information and can make informed decisions about valuation. Public equity markets approximate this through mandatory disclosure requirements, analyst coverage, and the continuous price discovery provided by high-frequency trading. Private company equity markets have none of these features. Information about company performance is closely held and shared primarily with existing investors and board members. There is no regulatory requirement for disclosure and no standardized reporting format. Many factors that drive value in healthcare companies, such as clinical trial results, reimbursement negotiations, and regulatory strategy, are often subject to confidentiality agreements and competitive considerations that prevent broad disclosure.

This information asymmetry creates a market for lemons problem in secondary transactions. Sellers typically know more about the company than buyers, and rational buyers must assume that sellers are more likely to want to exit positions when they have negative private information about company prospects. A founder selling shares in their own company before a liquidity event is a red flag. An early investor selling a position in a company that just raised a successful Series B at a marked-up valuation is suspicious. Why would they sell if they thought the company was going to continue performing well? The most innocent explanations are portfolio rebalancing, diversification, or personal liquidity needs, but buyers must always consider the possibility that the seller knows something they do not. This adverse selection problem depresses prices in secondary markets and can lead to market failure where buyers are unwilling to transact at prices that sellers find acceptable.

The information asymmetry is particularly acute in healthcare because so much value in healthcare companies is tied to binary events that are subject to confidential information. A digital health company negotiating a reimbursement contract with a major payer has private information about how those negotiations are progressing, whether the payer is likely to agree to favorable terms, and what the timeline looks like. An investor who sits on the board or maintains close contact with the CEO has much better information than an angel investor who receives quarterly updates. If that well-informed investor decides to sell their position in the secondary market, the buyer must ask why. If the investor decides not to sell and in fact wants to increase their position, the buyer has received valuable information even if the transaction does not occur.

Some secondary market participants attempt to overcome information asymmetry through diligence processes that involve speaking with company management, reviewing financial statements, and conducting expert interviews with industry participants who can provide context. But this level of diligence is expensive and time-consuming, and it is only feasible for large transactions where the potential profit justifies the cost. For small angel positions worth less than one hundred thousand dollars in current value, no rational secondary buyer will spend tens of thousands of dollars on diligence. The result is that small positions trade at steep discounts to reflect the information asymmetry and risk, or more commonly they do not trade at all because the bid-ask spread is too wide to allow a transaction.

The rise of data providers and private company intelligence services has made significant inroads into reducing information asymmetry. Companies like PitchBook, CB Insights, and Carta provide data on financing rounds, valuations, and investor composition that can help secondary buyers assess opportunities. Some specialized services track clinical trial progress, regulatory submissions, and reimbursement developments for healthcare companies. But the most valuable information, the information that actually drives returns in healthcare venture investing, is not available from data providers. It is private conversations with CEOs, board-level visibility into strategic decisions, and the pattern recognition that comes from sitting on multiple company boards and seeing how companies actually operate. This

information is available to professional venture capitalists who are deeply embedded in the ecosystem, and it is largely unavailable to secondary buyers trying to assess opportunities from the outside.

Price Discovery in Illiquid Markets

Even if information asymmetry could be overcome, the challenge of price discovery in illiquid markets would remain. In public equity markets, prices are discovered continuously through the interaction of thousands of buyers and sellers trading millions of shares. The current market price represents a rough consensus about the value of the company based on all publicly available information. In private company markets, prices are discovered episodically through financing rounds that might occur once every twelve to twenty-four months, and even these prices are complicated by the preference structures of venture capital investments.

When a healthcare company raises a Series A at a fifteen million dollar post-mortem valuation with a one-point-zero liquidation preference, what does this tell us about the value of common stock held by angels who invested in the seed round? The answer is not straightforward. The preferred stock purchased by the Series A investors has downside protection that the common stock does not. If the company is sold for less than fifteen million dollars, the Series A investors get paid back first and the common stockholders might receive nothing. This means that the common stock is worth less than the preferred stock, but how much less depends on the probability distribution of exit outcomes. If the company is very likely to exit for more than fifty million dollars, the difference between common and preferred is minimal. If there is a significant probability of a disappointing exit below the liquidation preference, the common stock could be worth fifty percent or more less than the preferred.

Secondary markets typically price common shares at a discount to the most recent preferred round, with discounts ranging from twenty to fifty percent depending on the stage of the company, the time since the last round, the performance relative to plan, and the general market sentiment. But these discounts are rules of thumb rather than results of rigorous valuation analysis, and they can be wildly off in either

direction. A company that raised a Series A at fifteen million dollars two years ago and has since tripled its revenue and is in active discussions for a Series B at sixty million dollars should have common stock worth more than the Series A preferred price less. But a company that raised at fifteen million dollars and has since failed to hit milestones and is struggling to raise its next round might have common stock worth twenty percent of the Series A price or less.

The challenge for secondary buyers is that they often do not have enough information to make these determinations accurately. They know the last financing round price and they might know some high-level metrics about company performance, but they typically do not have the detailed understanding of company trajectory, competitive dynamics, and exit prospects that would allow them to triangulate a fair value. The result is that secondary transactions often occur at prices that are either significantly too high or significantly too low relative to the intrinsic value of the shares, and the party with better information captures value from the party with worse information.

Some sophisticated secondary buyers attempt to solve the price discovery problem by constructing detailed financial models that project company performance, estimate exit valuations under different scenarios, probability-weight those scenarios, and calculate a present value for the equity. This is the theoretically correct approach to valuation, but it requires inputs that are often not available or highly uncertain. For example, what is the probability that a digital health company secures favorable CMS reimbursement? What is the likely exit multiple for a company in a particular category given current M&A market conditions? How should we discount cash flows for a company that might not be profitable for five or more years? Reasonable people can disagree on all of these inputs, and small changes in assumptions can lead to radically different valuations.

The practical reality is that price discovery in secondary markets for healthcare companies is driven more by supply and demand dynamics, market sentiment, and negotiating leverage than by rigorous fundamental analysis. A seller who desperately needs liquidity will accept a low price. A buyer who has conviction about a particular company and fears missing the opportunity will pay a premium. A platform that facilitates the transaction will push both parties toward the midpoint of the range.

earn their fee. The resulting transaction price tells us something about the balance of these forces at that particular moment in time, but it does not tell us much about the intrinsic value of the company.

Liquidity Engineering: Mechanisms and Realities

Given the limitations of traditional secondary markets, sophisticated investors have developed various mechanisms for engineering liquidity or liquidity-like outcomes for their healthcare angel portfolios. These mechanisms range from portfolio construction strategies that reduce the need for secondary liquidity to financial engineering that creates synthetic liquidity without requiring actual share sales.

One approach is to invest through structures that have built-in liquidity provisions. Some rolling funds and evergreen vehicles allow investors to redeem their capital on a periodic basis, subject to certain restrictions. The fund manager might facilitate liquidity by finding new investors to provide capital for redemptions, by managing the portfolio in a way that generates regular distributions from exits, or by maintaining a small cash buffer to handle redemptions. This approach can provide some liquidity to limited partners, but it does not solve the fundamental problem that the underlying portfolio companies are illiquid. If many investors want to redeem simultaneously, the fund manager may have no choice but to sell positions in the secondary market at unfavorable prices or to gate redemptions entirely.

Another approach is to construct portfolios with explicit vintage diversification, where an investor makes new investments every year and expects that older vintage investments will start producing liquidity events after seven to ten years. By the time an investor has been deploying capital for a decade, they should be receiving regular distributions from exits in their older vintage investments, and these distributions can be used to fund new investments, provide cash for living expenses, or simply accumulate as dry powder. This approach requires patience and discipline, and it requires that the investor has sufficient capital to continue making new investments even while their portfolio is illiquid. It also requires that the portfolio actually

produces exits at a reasonable cadence, which is not guaranteed in healthcare where many companies take longer than expected to reach liquidity and some never exit at all.

Some investors attempt to engineer liquidity through lending arrangements that allow them to borrow against their private company portfolios. The challenge is that most traditional lenders will not accept private company equity as collateral because it is difficult to value, difficult to liquidate, and subject to extreme volatility. Some specialized lenders will provide financing against private company portfolios, but typically only for large portfolios with meaningful concentrations in late-stage companies, and at loan-to-value ratios of twenty-five to forty percent with high interest rates. For most angel investors, borrowing against their portfolio is not a viable liquidity mechanism.

A more practical form of liquidity engineering for many healthcare angels is simply to sell partial positions when opportunities arise rather than selling entire positions. When a company has performed well and a secondary buyer approaches with an offer to purchase shares at an attractive price, an investor might sell twenty-five or fifty percent of their position to take some chips off the table while maintaining upside exposure. This approach provides some liquidity and reduces risk while preserving exit optionality. The challenge is that secondary buyers often want to purchase entire positions rather than partial stakes because the transaction costs are similar regardless of size and they want to have meaningful exposure if they are going to invest time in diligence.

Some angel investors have experimented with creating internal secondary markets within their own portfolios by syndicating portions of their positions to other accredited investors. The original investor acts as a quasi-fund manager, curating their portfolio and offering other investors the opportunity to purchase stakes in specific companies. This approach can provide some liquidity to the original investor while giving other investors access to curated deal flow. But it requires significant operational overhead to manage the cap table complexity, handle legal documents, and maintain relationships with co-investors. For most angels, the juice is not worth the squeeze.

The reality is that most liquidity engineering mechanisms in healthcare angel investing are marginal improvements rather than fundamental solutions. They can provide some incremental liquidity, some reduction in portfolio concentration, and the ability to rebalance and manage risk. But they do not transform the fundamental nature of angel investing, which is that it requires patient capital willing to be locked up for extended periods with significant risk of loss and uncertain timing of return. Investors who cannot accept this reality should not be angel investors in healthcare, regardless of what liquidity engineering mechanisms are available.

The Portfolio Construction Challenge

The illiquidity of healthcare angel positions creates distinctive portfolio construction challenges that differ from other asset classes and require careful thought about diversification, sizing, and pacing.

The traditional venture capital portfolio construction model assumes that most investments will fail, a small number will return the fund, and the goal is to capture exposure to the power law distribution of outcomes by making enough investments to ensure participation in the rare outlier successes. For institutional venture capital firms with large funds and professional staff, this might mean investing in forty to sixty companies per fund and maintaining active relationships with portfolio companies through board seats and regular interaction. For individual angel investors with limited capital and limited time, achieving adequate diversification is more challenging.

An angel investor who can afford to deploy fifty thousand dollars per year into healthcare investments might make five investments per year at ten thousand dollars each, or ten investments per year at five thousand dollars each, or some other combination. After five years of consistent investing, this investor would have a portfolio of twenty-five to fifty companies, which begins to approach reasonable diversification. But the challenge is that healthcare angel investing has lower hit rates than traditional software investing, particularly in categories like digital therapeutics or diagnostic companies where reimbursement is uncertain. An investor might e

that ten to twenty percent of their investments will be total losses within the first three years, another thirty to forty percent will struggle along without achieving meaningful traction, twenty to thirty percent will perform modestly and return one to three times capital, and ten to twenty percent will be genuine winners returning five times or more. With this kind of outcome distribution, an investor needs exposure to at least thirty to fifty companies to have high confidence of capturing some of the winner outcomes.

The pacing of investments creates further challenges. An investor who deploys capital too quickly will have a portfolio that is heavily weighted toward a single vintage, with all the attendant risks that come from investing during a particular market period. An investor who deploys too slowly will not achieve diversification quickly enough and will spend many years with a small portfolio where individual investment outcomes will have outsized impact on overall returns. The conventional wisdom is to invest consistently over time, maintaining discipline about deal flow and avoiding the temptation to hold capital back waiting for perfect opportunities.

But the illiquidity of healthcare angel positions means that portfolio construction is not a one-time decision but an ongoing challenge. As older investments appreciate in value and new investments are added, the portfolio becomes increasingly weighted toward a small number of winners. An investor who started with a well-diversified portfolio of fifty positions might find after five years that three of those positions represent seventy percent of their total portfolio value on paper. This concentration increases risk and creates the desire to trim winners and rebalance. But if those winners cannot be sold in the secondary market at attractive prices, the investor is stuck with a concentrated portfolio regardless of their preferences.

Some sophisticated angels address this through explicit portfolio construction rules. They might decide to never allow a single position to represent more than twenty percent of portfolio value, and they commit to selling portions of winning positions in the secondary market as opportunities arise to bring concentration back in line. They might decide to always maintain at least fifty percent of their portfolio in newer vintage investments to ensure exposure to emerging categories and avoid being too heavily weighted toward older companies. They might decide to allocate a certain percentage of the

angel portfolio to providing follow-on capital to their best-performing companies rather than only making new seed investments. All of these rules represent attempts to impose some structure and discipline on portfolio management in the face of illiquidity.

The challenge is that many of these rules become difficult to implement when secondary liquidity is not available or is only available at prices that make selling unattractive. An investor who wants to maintain diversification but cannot sell their concentrated positions is left with the uncomfortable choice of either accepting concentration than they would like or making fewer new investments to avoid increasing the overall portfolio size. An investor who wants to provide follow-on capital to winning companies but has limited dry powder because their capital is locked up in illiquid positions must either find new sources of capital or pass on follow-on opportunities they find attractive.

SPVs, Rolling Funds, and Structural Solutions

The structural limitations of individual angel investing have led to the development of pooled investment vehicles that attempt to solve some of the liquidity and portfolio construction challenges while introducing new complexities.

Special purpose vehicles have become a popular mechanism for angel investors to pool capital for specific investments. A lead investor who has access to an attractive opportunity might create an SPV and invite other investors to participate, allowing the group to write a larger check than any individual could write alone. SPVs provide several benefits: they allow smaller investors to access deal flow they would not otherwise see, they simplify the cap table for the company by consolidating many small investors into a single entity, and they allow the lead investor to build their reputation and demonstrate ability to bring capital to deals. For liquidity purposes, SPVs provide limited benefits. The underlying investments are just as illiquid as they would be if held directly, and the SPV structure actually adds some friction because

transfers of SPV interests typically require consent from the SPV manager and can trigger tax complications.

Rolling funds represent a more sophisticated structural solution. These funds allow investors to commit capital on a periodic basis, typically quarterly, with each commitment representing a separate legal fund with its own lifecycle. An investor might commit to invest fifty thousand dollars per quarter for four years, creating sixteen separate mini-funds that invest in deals as they arise. The rolling structure provides some natural diversification across vintages, and it creates the possibility for earlier funds to start generating distributions while later funds are still being deployed. Some rolling fund managers also build in modest redemption rights that allow investors to stop making new commitments or to receive pro-rata distributions from exits earlier than they would in a traditional fund structure. These provisions provide some liquidity-like benefits, but they do not fundamentally change the underlying illiquidity of the portfolio companies.

Evergreen fund structures take the liquidity engineering a step further by creating a single continuous fund that makes new investments, manages existing positions, and returns capital to investors through periodic distributions or redemptions. Unlike traditional venture funds that have a defined lifespan and return all capital through liquidation, evergreen funds continue indefinitely and attempt to provide ongoing liquidity through a combination of exit distributions and managed redemptions. A challenge with evergreen structures in venture is that they create potential mismatches between investor redemption requests and the actual liquidity of the underlying portfolio. If the fund has a strong year and many investors want to redeem, the fund manager may need to sell positions at inopportune times or gate redemptions, creating frustration and potential legal complications.

Some fund managers have experimented with hybrid structures that attempt to capture the benefits of multiple approaches. A fund might have a core evergreen strategy that holds long-term positions and provides periodic liquidity, combined with a secondary strategy that actively buys and sells positions to generate returns and facilitate redemptions. Or a fund might have a primary strategy focused on new investments and a separate continuation vehicle that holds mature positions pas

traditional fund term to avoid forced sales. These structures add complexity and require sophisticated fund administration, but they can provide better alignment with investor liquidity needs than traditional structures.

The reality is that no structure can fully solve the fundamental illiquidity of healthcare angel investing. Structures can smooth the rough edges, can provide some incremental liquidity, can reduce the psychological discomfort of decade-long hold periods. But at the end of the day, someone in the capital structure must be willing to accept illiquidity. Either the individual investors accept it directly by holding illiquid positions, or they transfer it to a fund manager who accepts it on their behalf, or they attempt to transfer it to secondary buyers who may or may not be willing to provide liquidity at acceptable prices. The illiquidity does not disappear; it just moves around the system.

Why Healthcare Is Different

Healthcare angel investing presents distinctive challenges for liquidity and secondary markets that distinguish it from angel investing in other sectors and warrant extra consideration.

The first difference is the extended timeline to meaningful outcomes. A software company raising a seed round might demonstrate product-market fit within two to three months, achieve meaningful revenue within twenty-four months, and reach an exit or late-stage financing within four to six years. The timeline is long but not extraordinarily so, and there are clear interim milestones that allow investors to track progress and update their views on company prospects. Healthcare companies, particularly those in digital health categories that require clinical evidence and reimbursement, operate on much longer timelines. A company might spend two years developing and refining their product, another two years running clinical studies, another year or two navigating regulatory clearance, and another two to three years securing reimbursement and achieving commercial traction. By the time the company reaches an inflection point where outcomes become clearer, seven to ten years may have passed. This extended timeline increases the carrying cost of illiquidity and

makes it more difficult for investors to maintain conviction through inevitable setbacks.

The second difference is the binary nature of many healthcare outcomes. A software company that fails to achieve product-market fit can often pivot to an adjacent opportunity, can find a smaller niche that still generates acceptable returns, or can be acquired for modest value that returns some capital to investors. Healthcare companies facing binary regulatory or reimbursement decisions have less room to maneuver. A digital therapeutic that fails to secure FDA clearance cannot simply pivot to a different indication without running new studies. A remote monitoring platform that fails to secure adequate reimbursement from CMS cannot easily pivot to a capitated pay model because patients will not pay out of pocket for something they believe should be covered by insurance. The result is that healthcare portfolios have more extreme outcome distributions with more total losses and potentially higher returns on winners but less middle ground.

The third difference is the challenge of external investors conducting adequate diligence. In software investing, a potential secondary buyer can often assess company quality by examining product demos, speaking with customers, reviewing public information about usage and traction, and triangulating competitive positioning through market research. In healthcare investing, the most important factors driving value are often opaque or confidential. Is the company's clinical evidence actually compelling, or are there methodological limitations that will become apparent under scrutiny? How are reimbursement negotiations with major payers actually progressing? What is the probability that a competitor receives regulatory approval first and captures the market? These questions are difficult to answer without insider access, and this information asymmetry makes secondary markets less efficient and less attractive to potential buyers.

The fourth difference is the concentration of eventual exits in strategic acquisitions rather than IPOs. Software companies have multiple exit pathways including IPOs, strategic acquisition by technology companies or other software firms, and in some cases private equity buyouts. Healthcare companies, particularly in digital health and diagnostics, exit primarily through strategic acquisitions by pharmaceutical

companies, medical device firms, health systems, or payers. This concentration in strategic exits means that exit timing and valuation are driven by the strategic priorities and M&A budgets of a relatively small number of acquirers, creating additional uncertainty and making it harder to predict when liquidity will materialize.

These differences mean that secondary markets in healthcare will always be more challenging than in other venture categories, and investors who expect software-like liquidity in healthcare will be persistently disappointed. The structural features of healthcare markets create irreducible illiquidity that no amount of financial engineering can fully overcome.

Conclusion: Building for Illiquidity

The question facing healthcare angel investors is not how to engineer liquidity in portfolios that are fundamentally illiquid, but rather how to construct investment strategies that accept and embrace illiquidity as a defining feature of the asset class. This requires a different mindset than prevails in much of modern finance, where liquidity is treated as an unalloyed good and illiquidity as a problem to be solved. Healthcare angel investing requires recognizing that illiquidity is the price of admission for exposure to a category of investments that offers distinctive returns and operates in markets with meaningful barriers to entry.

The first principle of building for illiquidity is to only invest capital that can genuinely be locked up for a decade or more without creating financial or psychological stress. This sounds obvious but is frequently violated in practice. Investors who are stretching to make angel investments, who are counting on being able to sell positions in a few years to fund other commitments, or who are investing money they might need for liquidity events in their own lives are setting themselves up for difficult decisions and suboptimal outcomes. The capital allocated to healthcare angel investing should be genuinely patient capital with no expectation of near-term access. If this constraint means making fewer or smaller investments, that is the correct choice.

The second principle is to construct portfolios with sufficient diversification to tolerate the high mortality rates and extreme outcome distributions characteristic of healthcare venture. An investor who makes three healthcare angel investments per year is not building a portfolio; they are buying lottery tickets. An investor who makes ten to fifteen investments per year over a period of five years, creating a portfolio of fifty to seventy-five positions, is beginning to have adequate diversification to capture power law outcomes. This level of diversification requires either substantial capital or small check sizes, and many aspiring healthcare angels lack one or both. For investors who cannot achieve adequate diversification through direct investments, participating in diversified funds or syndicates may be the more rational approach despite the added fees.

The third principle is to develop independent conviction about investment opportunities rather than relying on the existence of secondary markets as a safety net. The temptation when secondary markets exist is to lower one's diligence standards or conviction thresholds with the rationalization that if the investments do not work out, a secondary exit will be available. This is backwards reasoning that leads to investing in companies one does not truly believe in and then discovering no secondary buyer can be found at any reasonable price. Better to invest only in companies where one has genuine conviction about the long-term opportunity, the quality of the team, and the probability of eventual success, and to enter the investment with the mental model that the capital will be locked up until the company exits or fails.

The fourth principle is to maintain realistic expectations about the role of secondary markets in portfolio management. Secondary markets are useful at the margin for opportunistic transactions when attractive buyers emerge, for trimming excessive concentration in positions that have appreciated dramatically, and for providing modest liquidity in situations of genuine personal need. They are not useful as a portfolio management tool, as a mechanism for regularly rebalancing portfolios, or as a substitute for patient capital. Investors who orient their strategy around secondary market liquidity are optimizing for a capability that may not be available when needed and that may prove counterproductive if it leads to selling winning positions prematurely.

The fifth principle is to recognize that the discomfort of illiquidity is a feature rather than a bug of healthcare angel investing. The inability to mark positions to mark daily, the lack of constant price feedback, and the necessity of maintaining conviction through uncertainty are not unfortunate byproducts of market inefficiency but rather are fundamental characteristics that help generate returns. Public market investors must contend with constant noise, with daily price fluctuations that reflect sentiment more than fundamentals, and with the psychological challenge of watching positions move dramatically on no new information. Private market investors are insulated from much of this noise and can maintain long-term focus without the distraction of meaningless price movements. This insulation is valuable and should be recognized as such rather than treated as a constraint to be overcome.

The healthcare angel investors who are most successful over long time horizons are not those who have become experts at engineering liquidity or navigating secondary markets. They are those who have built portfolios and lives that can absorb a dose of illiquidity without stress, who have developed conviction and patience to allow investments to mature, who have achieved sufficient diversification to tolerate inevitable failures, and who have cultivated the relationships and judgment to identify opportunities that will generate exceptional returns when they eventually exit. These investors treat secondary markets as an occasionally useful tool rather than a central component of strategy, and they build portfolios designed to succeed on the assumption that no secondary liquidity will be available.

The promise of liquid secondary markets for healthcare angel positions is seductive precisely because it suggests that one can have exposure to an illiquid asset class without accepting the full cost of illiquidity. This promise is largely false. The secondary markets that exist are thin, inefficient, expensive, and available primarily in circumstances where using them is suboptimal. The mechanisms for engineering liquidity provide marginal benefits but do not fundamentally change the nature of the asset class. The structural features of healthcare markets create irreducible illiquidity that persists regardless of financial innovation.

For investors who can accept this reality and construct strategies accordingly, healthcare angel investing offers compelling opportunities to deploy patient capital.

into companies solving important problems in large markets with meaningful barriers to entry. The returns can be exceptional for investors who achieve diversification, maintain conviction, and accept decade-long hold periods. But these returns are compensation for bearing genuine illiquidity risk and for accepting that capital will be locked up for extended periods with uncertain timing of eventual liquidity events. Investors seeking liquid alternatives should look elsewhere. Those willing to embrace illiquidity as the price of admission to an attractive asset class will find that healthcare angel investing, approached with appropriate humility and patience, generate meaningful returns and provide exposure to the transformation of one of the largest and most important sectors of the economy.

The secondary market infrastructure that has emerged over the past decade provides some incremental benefits and some optionality for investors facing unusual circumstances. But it does not and cannot transform healthcare angel investing from an illiquid asset class requiring patient capital into a liquid asset class amenable to active trading and frequent rebalancing. Investors who design strategies around secondary liquidity are building on a foundation that may not be there when needed. Those who design strategies around fundamental illiquidity are building on bedrock. The difference in outcomes over multi-decade investment careers is substantial, and the choice of which foundation to build on may be the most important decision a healthcare angel investor makes.

If you are interested in joining my generalist healthcare angel syndicate, reach out to treyrwales@gmail.com or send me a DM. We don't take a carry and defer annual fees for six months so investors can decide if they see value before joining officially. Accredited investors only.

← Previous

Next

Discussion about this post

Comments Restacks



Write a comment...



B. Isaac Francisco  Oct 26

Great thoughts and perspective in this article.

Appears there is a logical contradiction in the following passage though: "The companies that are easy to sell in secondary markets are often the companies that have the frothy valuations and the highest risk of disappointment. The companies that would

actually be wise to trim exposure to are often the ones where no secondary buyer found."

Alternatively, it seems it "would actually be wise to trim exposure to" exactly "the companies that have the most frothy valuations and the highest risk of disappointment?"

Not that it would be ideal for buyers of secondary shares if they are sold high-price risk shares, though that would seem optimal for the seller, on whose perspective passage is focused.

Thoughts?

♡ LIKE 💬 REPLY

1 reply

1 more comment...