

EPOCH

Structural Mispricing

EPOCH CAPITAL · RESEARCH

OCTOBER 2026

ADITYA NIDMARTI

EPOCH CAPITAL · RESEARCH · ESSAY III

Structural Mispricing

Venture investing is fundamentally a game of maximizing upside. The consequence of success must be large enough to matter to the portfolio.

ADITYA NIDMARTI

OCTOBER 2026 · 11 MIN READ

CONTENTS

Introduction

The Future Inside the Price

Two Ceilings

Connecting the Evidence

Establishing the Higher Ceiling

Investing Before the Outcome Is Known

Working Backward to Price

Paying More Than Competing Investors

A Discipline for Every Decision

Picking Your Spots

A company can succeed commercially without producing an exceptional investment outcome. The investor needs to understand how large and valuable the business could become, and whether the ownership acquired would make that success consequential. The search must therefore begin in areas where the ceilings are high.

The first installment described the [Productive Capacity Supercycle](#): the reallocation of capital toward the systems that expand and secure the economy's productive base. The second developed [Category Foresight](#): the discipline of identifying the markets forming within that transition and the companies capable of defining them. Together, they provide a way to look for unusually consequential outcomes.

Finding a high ceiling is the beginning of the investment task. The investor also needs a price that leaves a meaningful share of that upside. As the same future becomes widely understood and reflected in valuations, the opportunity to acquire it at a structurally undervalued price diminishes.

This final installment completes the three-part argument. Once Category Foresight establishes a credible thesis, the task is to identify structurally undervalued opportunities within the Productive Capacity Supercycle. Which high ceilings remain inadequately reflected in the price, and how does an investor recognize them?

Structural mispricing begins with a gap between the future business consensus expects and the future business the research supports. Consensus may recognize only growth within an existing market, while the evidence supports a larger category or a more valuable role for the company. The investment is structurally undervalued if that researched view supports an attractive prospective return at today's price, after accounting for the range of outcomes, their probabilities, time, and financing.

The investor's advantage comes from developing a view of the higher ceiling before consensus does. That requires identifying something specific that others are missing about the category's scale, the company's position, or the likelihood of reaching that future. A different opinion becomes useful when the research substantiates it and the available price leaves room to profit from it.

I call the work of developing that understanding and determining what to pay for it Structural Mispricing.

The Future Inside the Price

Every valuation contains assumptions about the business a company will become. Conventional analysis often begins with its existing market: revenue expands from the current customer base, margins approach those of established suppliers, and demand remains bounded by familiar budgets. That approach can encourage linear extrapolation from today's visible business.

Category formation can change those relationships. A technology may cross a threshold that makes previously impractical applications viable. Falling costs can bring new customers into the market. Greater production volume can reduce costs further, opening additional applications and accelerating adoption.


These changes can produce nonlinear growth, including periods of exponential adoption. The eventual market may be much larger than today's customer base or spending category suggests. Category Foresight investigates why that expansion could happen: which changes in cost, capability, or customer need could cause adoption to accelerate, and what would eventually constrain it.

The pricing disagreement is therefore about the business that could emerge: a company growing within the existing market, or one serving a new category with much greater scope.

Two Ceilings

The **consensus-implied ceiling** is the potential size and profitability of the business that the prevailing valuation appears to recognize.

The **category-underwritten ceiling** is the potential size and profitability supported by research into how the category could develop and the role the company could occupy within it.



The current price is only a clue to consensus. Different combinations of outcomes, probabilities, and timing can produce the same valuation. The investor must investigate which assumptions explain the price and where the research supports a different view.

Consider, in simplified terms, a company developing enabling technology for laser defense. It could be evaluated as a laser-source supplier, selling a critical component to established defense contractors. The analysis would estimate component sales into existing programs and apply the economics of comparable suppliers.

A different future becomes plausible if advances in the technology make a broader class of laser-defense systems viable. Demand could then extend beyond the applications and programs visible today. If the company can also build, deliver, and support complete systems, it could become a laser-defense prime: the contractor responsible for supplying the overall capability.

The higher ceiling depends on both changes. The category would become larger, and the company would capture the economics of supplying the complete system, with the greater responsibility and capital requirements that entails. Success as a prime could therefore have much greater economic consequences than success as a component supplier. Increasing the growth rate of component sales would not capture that difference.

The company might instead remain a valuable supplier while another contractor captures the prime relationship. The investor therefore needs evidence that the technology can be manufactured economically, that customers will adopt the complete system, and that the company can deliver and finance it. The investment edge would lie in substantiating that larger future while the prevailing valuation still treats the company principally as a supplier.

The current price is only a clue to consensus.

Connecting the Evidence

The insight that reveals a higher ceiling has to be developed. It often comes from connecting facts that are understood separately. A technical advance may appear in scientific literature, a changing operational requirement in procurement plans, and an improvement in production economics in manufacturing data. Their combined implication for a company can remain unclear even when the individual facts are widely known.

The investor must connect these disparate dots into a causal argument. The technical advance makes a capability possible; the customer's emerging requirement creates a reason to acquire it; manufacturing progress makes it economical to supply. The research then asks whether this company can bring those developments together and retain a durable share of the value created.

Specialization can leave these connections underexamined. Technical experts may focus on performance, customers on immediate operating needs, and investors on existing budgets and comparable businesses. The opportunity can remain open while evidence for a consequential future accumulates across these fields but has yet to be assembled into the market's valuation of the company.

The investor should be able to state the disagreement precisely: what does consensus assume, what connection does the research reveal, and how does that change the business the company could become? The next task is to test the conditions required for that outcome and the alternatives if some of those conditions fail.

Establishing the Higher Ceiling

Research must establish both why success could be exceptional and how much weight to assign that outcome. The higher ceiling must be developed independently of the price the investor hopes to justify. The analysis has to connect potential adoption to company revenue, revenue to durable cash generation, and cash generation to equity value.

If the thesis depends on exponential adoption, the investor must explain what causes it and test whether the business can support it. Which improvements change customer economics enough to prompt purchases? Can manufacturing and deployment keep pace? Where do procurement constraints, competition, or market saturation slow growth? Demand can expand rapidly while falling prices or rising delivery costs prevent the company from earning attractive returns.

Pressure testing means examining what happens when the thesis's assumptions fail or take longer to materialize. Slower cost reductions may delay adoption. Delayed adoption may require another financing. A competitor may reduce both market share and pricing power. Each change can affect the eventual value, the time needed to reach it, and the ownership an investor retains.

The analysis should produce a set of distinct, plausible outcomes, each with an explicit probability, capital requirement, and expected time to realization. One may involve category leadership. Another may produce a narrower supplier business. Others may end in weak economics or failure. The category-underwritten ceiling describes a plausible upside within this range. No particular future company outcome should be treated as certain.

Both the outcomes and the probabilities assigned to them require research and debate. Evidence of technical progress, customer adoption, and competitive strength should inform the estimates. Alternative interpretations should be challenged, and the valuation tested across reasonable probability estimates. This work gives the investor a basis for taking meaningful exposure to an exceptional outcome before it is established.

Only after those outcomes have been developed and challenged should the investor calculate an entry price. Raising the ceiling or the probability of success merely to make an asking valuation deliver the desired return reverses the discipline. The research establishes the outcomes. The return requirement determines what the investor can pay for exposure to them.

The research establishes the outcomes. The return requirement determines what the investor can pay for exposure to them.

Investing Before the Outcome Is Known

Even after the research, much of the value being underwritten remains to be created. Venture investing is akin to buying an out-of-the-money call option: paying today for the possibility of a substantial gain if the future unfolds favorably.

A lower valuation buys more ownership, but the company still has to create the value on which the investment depends. An arbitrary discount cannot make that future certain. The investor should embrace uncertainty with conviction grounded in the researched ceiling, the range of outcomes around it, and the price paid for exposure to them.

The investor must also be prepared for capital to remain tied up. Venture shares cannot be assumed to be immediately saleable at the price paid, and the company may need further financing before it produces a return. The underwriting must therefore consider both the time required and the capital needed along the way.

For a venture fund, this includes planning reserves: capital set aside for possible follow-on investments. Reserves give the fund capacity to maintain or increase ownership as evidence develops and the terms justify further investment. Without that capacity, an investor can identify the right company yet retain too little ownership for its success to matter. Reserve planning must account for the likely size and timing of future financings across the portfolio.

Working Backward to Price

The researched outcomes determine what the investor can afford to pay. The calculation starts with what each outcome could return to the investor, given the ownership retained and the capital required to get there.

For this framework, I use a 5× probability-weighted gross multiple as a pricing reference, before fund fees and carried interest. The successful outcome may be much larger. The reference tests the return across the full range of outcomes, including failure.

Suppose, purely for illustration, an investment has a 50% chance of returning 10× invested capital and a 50% chance of losing it all. The probability-weighted return is 5×. These multiples assume no follow-on investment and already account for dilution.

Pay twice as much for the same eventual ownership, and the successful outcome returns 5× instead of 10×. With the probabilities unchanged, the expected gross multiple falls to 2.5×. The higher price reduces the return available from the same business outcome.

If the fund invests in later rounds, calculate the overall return using its final ownership and the total amount it invested in the company. Money set aside for future rounds counts only if it is actually invested.

The holding period also affects what the investor should pay. A fivefold return over ten years produces a lower annual return than the same multiple over five.

The outcomes and probabilities must withstand the research and debate described above. If the asking price requires assumptions the evidence cannot support, the investor should seek different terms or pass.

Paying More Than Competing Investors

A higher category ceiling can justify paying more than competing investors only when the probability-weighted range of outcomes supports the higher price. In the laser-defense example, the investor must weigh becoming a defense prime, remaining a supplier, or failing, with financing needs and timing assessed in each case. That assessment may support a higher bid than an analysis confined to component sales within existing programs.

If the higher bid still offers the required return after accounting for time and dilution, the investor can offer more and win the investment. This is how an investor can occasionally “overpay” relative to competing bids while entering a structurally undervalued opportunity. The competitors may be extrapolating today’s component business, while the research supports nonlinear adoption and a credible chance of a more valuable role for the company.

The same research establishes where to stop. If bidding pushes the price beyond what the outcomes support, the appropriate decision is to decline. Winning a deal does not validate the thesis or justify changing its assumptions.

Winning a deal does not validate the thesis or justify changing its assumptions.

A Discipline for Every Decision

Structural Mispricing applies whenever the investor must decide whether to commit capital or continue owning an investment. That includes the first investment, a follow-on financing, a competitive bid, and a choice between holding and selling when liquidity becomes available.

For new capital, the question is what return the proposed entry price offers. That applies equally to reserves: each follow-on investment needs to earn its place against

the fund's other uses of capital. A strong original thesis does not justify investing at every subsequent valuation. Reserves preserve the ability to act; updated underwriting determines whether to use them.

For an existing holding, the question is whether continuing to own the shares offers an attractive prospective return compared with the amount that could be realized through an available sale. The original purchase price cannot resolve that decision.

Each assessment requires an updated view of the outcomes and their probabilities. Progress toward reliable manufacturing may reduce execution risk. Repeat customer adoption may support a larger market. Another financing may reduce the ownership retained. The thesis needs observable milestones that show which assumptions to revise and why.

Operating progress does not always improve the investment opportunity. A company can improve its product while becoming less attractive because its financing needs rise or its valuation advances faster than the evidence. As other investors recognize the same ceiling, the price may also rise enough to absorb the original mispricing. Conversely, a higher valuation can accompany a better opportunity when meaningful progress reduces uncertainty or establishes a larger credible outcome.

The same comparison of price and prospective return applies across investment stages. An early company can be inexpensive in absolute terms yet offer poor compensation for its risks. A later investment can cost much more and offer a stronger prospective return.

Picking Your Spots

The discipline brings together three judgments: whether the company can capture enough of the category's value to produce a consequential outcome, whether the research reveals something about that outcome that consensus has yet to price, and whether the available terms allow the investor to own enough of it on attractive economics. The ceiling establishes the magnitude; the researched outcomes and their probabilities establish what it is worth paying.

Those judgments can support committing capital, declining a financing, or choosing whether to hold or sell, according to the return available at the moment of decision. Cheap and expensive acquire meaning through that analysis.

The Productive Capacity Supercycle directs the search toward the systems where consequential companies could emerge. Category Foresight identifies the categories that could support unusually high ceilings and the companies capable of reaching them. Structural Mispricing asks whether the investor can acquire meaningful ownership before consensus fully reflects that future in the price.

The work begins with the ceiling. It ends with the price.

Also published on Aditya Nidmarti's website: <https://nidmarti.com/writing/structural-mispricing/>