

July 20, 2026 Meeting Summary

July 20, 2026 | Chicago | By TIG

We held the latest meeting of The Investing Group on Monday, July 20, 2026, at the Union League Club of Chicago in Room 816.

Introduction

We held the latest meeting of The Investing Group on Monday, July 20, 2026, at the Union League Club of Chicago in Room 816.

Conversation moved across small cap market structure, geopolitics and energy, AI infrastructure, power generation, frontier model competition, China, education, robotics, market structure, and a lightning round covering individual ideas.

Several themes tied the evening together. The first was that power and physical infrastructure continue to define the pace of the AI buildout. The second was that execution, incentives, and management quality matter just as much as the underlying theme itself. The third was that competition at the frontier continues accelerating while companies with established distribution and cash-generating businesses may be best positioned to translate AI into long-term value. Finally, despite rapid technological progress, human judgment remains a durable competitive advantage in both investing and business.

1. Chinese Micro Cap Mania, Halt Mechanics, and Trading the Chaos

The discussion began with one of the day's most remarkable examples of low-float speculation, a reminder of how quickly speculative momentum can overwhelm fundamentals in today's market.

Discussion centered on ZYBT, a recently listed Chinese small cap and veterinary vaccine producer that surged from just above \$1 at the open to close at \$8.01, finishing at its high of the day after gaining more than 1,000% from the prior session's close.

Much of the conversation examined the market mechanics that drove the move, including float dynamics, volatility halts, liquidity, and order execution. With no identifiable company-specific catalyst released that day, the rally appeared to be driven almost entirely by momentum and market structure. ZYBT had roughly 47 million shares outstanding but a public float of only about 4.4 million shares, meaning less than 10% of the company's shares were freely tradable. More than 90 million shares changed hands during the session, turning over the entire public float roughly twenty times and illustrating how extraordinary price movements can develop when heavy trading volume collides with a limited supply of tradable shares.

The stock repeatedly triggered Nasdaq Limit Up Limit Down (LULD) volatility pauses throughout the day, often trading for only seconds before entering another five to ten minute halt. As each reopening established a new reference price, the volatility bands effectively stepped higher

throughout the session, creating a staircase effect that allowed the stock to continue advancing while repeatedly pausing trading. The halt sequence carried into the final minutes of regular trading, leaving the stock halted at its \$8.01 closing price. Because securities already halted during the last ten minutes of the session generally do not

reopen before the closing auction, anyone short the stock had little opportunity to manage the position before the market closed. When trading resumed after hours, where LULD halts no longer apply and liquidity becomes substantially thinner, the stock briefly traded into the \$11 range before reversing sharply and eventually fading into the \$3 range. The sequence illustrated how repeated volatility halts can amplify pressure on short sellers by limiting opportunities to cover during regular trading hours before forcing them into a far less liquid after-hours market. The discussion also covered order execution, noting that once trading moved into the after-hours session, stop orders were generally no longer available, leaving limit orders as the primary tool for managing risk in a much thinner and more volatile market.

Discussion also turned to JLHL's trading activity on July 9. The recently listed Chinese small cap rallied from just above \$3 into the low teens during regular trading before extending further after hours. Although the percentage gain did not match ZYBT's, the move demonstrated many of the same structural characteristics, including a very small public float, exceptionally high turnover relative to shares available for trading, and rapid price expansion that far exceeded what company fundamentals alone could reasonably explain. Unlike ZYBT's halt-driven advance, JLHL climbed steadily throughout most of the trading session, illustrating how stocks with similar structural characteristics can produce very different trading patterns.

The conversation also noted that these episodes continue appearing within a fairly recognizable segment of the market. Many involve recently listed Chinese Nasdaq companies with very small public floats and significant insider ownership, creating conditions where relatively modest amounts of buying or short covering can produce outsized price movements. In these situations, sentiment, float dynamics, short positioning, liquidity, and halt mechanics often become the primary drivers of price action while underlying business fundamentals play only a limited role over the course of a single trading session. Capturing substantial gains requires not only identifying the opportunity but also successfully exiting before liquidity deteriorates or trading becomes restricted, illustrating both the exceptional upside and the considerable risks associated with this corner of the market.

2. Iran, Oil Markets, and the Iraq Pipeline

The group revisited the conflict involving the United States, Israel, and Iran following the resumption of hostilities after the collapse of the ceasefire earlier in July.

As of the meeting, reporting indicated the United States had completed roughly nine consecutive days of strikes against targets in Iran, with CENTCOM stating the operations were intended to impose continued costs on Iranian forces while degrading their ability to threaten commercial shipping and civilian vessels operating near the Strait of Hormuz. The conflict also continued expanding across the region. Two American service members were killed in action during the July 17 attack on Muwaffaq Salti Air Base in Jordan, while another remained missing as recovery efforts continued following the strike. A separate service member was killed at Erbil Air Base in northern Iraq during the controlled detonation of unexploded munitions recovered from a downed Iranian drone. Nearly one hundred U.S. service members had sustained some level of injury, mostly minor, since the ceasefire collapsed, although military officials indicated additional injuries could still be identified as reporting continued. Missile interceptions and air defense activity also extended across Jordan, Bahrain, Kuwait, Qatar, and Iraq.

Energy markets remained a central theme throughout the discussion. Brent crude had fallen to roughly \$70 per barrel in early July, its lowest level since the day before the conflict began, before rallying back into the high \$80s following renewed fighting and Iran's efforts to disrupt shipping through the Strait of Hormuz. The rebound of more than 20 percent in less than three weeks illustrated that energy markets continued responding quickly to changes in geopolitical conditions. Major indexes remained relatively resilient over the period despite weakening market breadth, while energy shares generally outperformed during periods of heightened tension.

Discussion then turned toward the theory that Iran may simply be attempting to extend the conflict. One view

circulating publicly holds that Iran is preserving military resources while attempting to raise economic and political pressure over time instead of seeking a decisive military outcome. Reserve calculations were also reviewed during the meeting, including a theoretical scenario of roughly sixteen days of U.S. consumption using the Strategic Petroleum Reserve alone and broader estimates suggesting global inventories could cover approximately ten to twenty weeks depending on assumptions. Those figures were not viewed as precise forecasts, but they highlighted both the importance of existing inventories and the limits of relying on strategic reserves if supply disruptions become prolonged. The Strategic Petroleum Reserve had declined from roughly 414 million barrels at the beginning of the conflict to approximately 316 million barrels by mid-July following coordinated strategic reserve releases by the United States and other countries. The conversation also acknowledged that U.S.

intelligence continued assessing Iran as retaining meaningful missile-launch capability despite significant damage to portions of its defense industrial base.

The discussion also prompted a review of several longer-term developments that help provide context for the evolving energy landscape. On July 17, Iraq and Syria signed an agreement during a U.S.-Iraq business summit in Washington to begin reconstructing the long-idled Kirkuk- Baniyas pipeline, with a U.S.-led consortium that includes Chevron (CVX) participating through non-binding agreements while technical studies and commercial terms continue to be developed. The project has been described as targeting as much as 2 million barrels per day of capacity once fully developed, compared with the pipeline's historical design capacity of roughly 700,000 barrels per day. The system has been out of service since 2003, making this a multi-year infrastructure project rather than a near-term supply solution. The agreement followed a White House meeting earlier that week between President Trump and Iraqi Prime Minister Ali al-Zaidi and formed part of a broader package of 48 agreements, memoranda of understanding, and partnership declarations intended to expand economic cooperation between the United States and Iraq.

These developments also provide useful context for current energy markets. Iraq's oil production remains roughly half of its pre-conflict level, while vessel traffic through the Strait of Hormuz has fallen dramatically from historical norms during periods of heightened military activity. Together, those developments reinforce the importance of alternative export routes while helping explain renewed interest in projects such as the Kirkuk-Baniyas pipeline. They also illustrate why energy markets remain highly sensitive to any meaningful deterioration in global supply flows as strategic reserve buffers continue to decline.

3. Bloom Energy, Scandium, and the Anatomy of a Short Report

Bloom Energy (BE) generated one of the longest discussions of the evening following the publication of a high-profile short report earlier in the month.

Hunterbrook Media published its report on July 8 while disclosing that it held a short position. Shares fell as much as 12% intraday before closing down more modestly. Although the stock traded below \$200 by the time of the meeting, much of that decline occurred after the initial short report as valuation multiples across portions of the AI infrastructure and power sector compressed. Before the selloff, Bloom had climbed from roughly \$90 at the beginning of the year to just over \$350 in late June, reflecting the type of valuation typically associated with near-perfect execution.

The core of the short thesis centered on scandium, a rare earth element used in Bloom's solid oxide fuel cells. Hunterbrook constructed a supply and demand model estimating that

increasing annual deployments from roughly 1 gigawatt to a long-term 5 gigawatt production scenario would require approximately 220 metric tons of fuel cell grade scandium oxide, compared with projected global supply of roughly 240 metric tons. Under that framework, Bloom alone would require more than 90% of projected global supply, leaving little margin for competing demand or supply disruptions. The report also projected total market demand

approaching 310 metric tons annually, suggesting the market could face a supply deficit even before Bloom reached that level of production. Given that China controls the overwhelming majority of global scandium production and processing, the report argued that the key question was not whether Bloom's supply chain had Chinese exposure, but to what extent. Hunterbrook further alleged that trade data, corporate filings, and other supply chain research pointed to China-linked sourcing, including claims involving Hunan Oriental Scandium, a major Chinese producer of fuel cell grade scandium oxide that the report identified as a key supplier.

The report extended beyond scandium. Additional concerns focused on the concentration of revenue generated through joint ventures in which Bloom holds ownership interests, the relationship between the company's reported backlog and its remaining performance obligations, and permitting questions surrounding Oracle (ORCL)'s Project Jupiter data center in New Mexico. Together, the report argued that Bloom's valuation depended on assumptions that warranted greater scrutiny.

On July 9, Bloom responded by calling the allegations false and misleading. The company reaffirmed its audited financial statements, stated that it holds sufficient scandium oxide to meet current demand and its existing backlog, denied any dependence on China, and said its supply chain provides visibility to support production of up to 25 gigawatts annually. Management also referenced access to several hundred metric tons of scandium oxide while explaining that it does not publicly disclose detailed sourcing information because of supply chain considerations. Shares recovered intraday following the response before closing modestly higher. Within twenty-four hours, Crossroads Capital released its own short thesis, describing the overlap as coincidental while arguing that it reinforced its conclusions. Crossroads also highlighted the challenge of sourcing sufficient non-Chinese scandium should Bloom eventually scale toward a 5 gigawatt production level.

Much of the discussion focused on the wording of Bloom's rebuttal, as situations like this are often decided by the quality of a company's response. While management addressed the allegations in qualitative terms and referenced access to significant scandium supplies, it did not publish detailed quantitative data directly addressing Hunterbrook's 220 versus 240 metric ton framework or the underlying supply assumptions. When a company facing this type of criticism responds to a question centered on tonnage without providing corresponding tonnage data,

that omission can become part of the story. It was also noted that management's denials relied primarily on categorical statements rather than a detailed numerical rebuttal.

Research surrounding the short thesis also identified yttrium oxide as the primary alternative electrolyte. The tradeoff is significant. Yttrium-based cells generally operate closer to 800 to 1,000 degrees Celsius compared with roughly 700 to 800 degrees for scandium-based designs, increasing thermal stress on surrounding components, shortening expected lifespan, requiring more expensive alloys and thermal shielding, and reducing power density. While yttrium is substantially more abundant than scandium, it does not fully eliminate the geopolitical supply question because it is also subject to China's April 2025 rare earth export licensing regime. Bloom's use of scandium reflects the performance advantages associated with lower operating temperatures rather than the absence of alternative materials.

One valuation framework discussed during the meeting suggested fair value was roughly 60% below Bloom's June peak, implying a value near \$140 per share. Based on the information available through the meeting, Bloom's operating results continued supporting the underlying business. Earlier in the year, first-quarter revenue more than doubled year over year, the company returned to profitability, earnings substantially exceeded expectations, full-year guidance was raised, and Bloom continued reporting a rapidly expanding backlog tied to AI infrastructure and power demand.

The discussion ultimately settled on the view that the underlying business remained strong while the valuation had become far more difficult to justify after the extraordinary rally. Data centers continue requiring reliable power, Bloom's fuel cells address a genuine infrastructure bottleneck, and demand remains robust. Even so, strong businesses do not always represent attractive investments at every price, particularly when expectations already assume years of exceptional execution. With second-quarter earnings scheduled for July 28, management's

discussion of scandium supply, backlog growth, and long-term production plans was viewed as one of the most important near-term developments to monitor.

4. Oracle's Project Jupiter, Power Constraints, and the Nuclear Question

The Bloom discussion led directly into a recent SemiAnalysis report examining Oracle's Stargate related data center project in New Mexico, known as Project Jupiter. The report suggested the project faces a potential one to two year delay due to permitting challenges and natural gas pipeline constraints. Shortly before the meeting, New Mexico's Land Commissioner denied the proposed pipeline rights-of-way for a second time, reinforcing concerns that the timeline could continue slipping. Oracle plans to deploy up to 2.45 gigawatts of Bloom Energy fuel cell capacity at the campus to support OpenAI infrastructure, replacing a previously planned natural gas

generation facility. While the transition to fuel cells addressed many of the environmental concerns surrounding the original proposal, it did not eliminate the project's dependence on pipeline natural gas, leaving fuel delivery as one of the primary bottlenecks. Oracle shares had continued declining as investors weighed the company's leverage profile and substantial AI infrastructure spending. The conversation also explored whether much of that negative sentiment may have already become reflected in the share price, with the possibility that sentiment could begin stabilizing over the next two to three months if long term demand for AI compute continues developing as expected.

The conversation then expanded into power infrastructure. Even if a group of investors identified an attractive data center site and had sufficient capital available, the primary constraint would still be electricity. In many regions, utilities simply cannot deliver power at the scale these facilities require, forcing developers to secure on site generation. Coal is no longer viewed as a practical option, while conventional nuclear projects remain constrained by lengthy permitting and construction timelines. As a result, much of the industry has shifted toward fuel cells, natural gas turbines, solar, and other on site generation approaches.

Nuclear policy also became part of the discussion. Four executive orders signed in May 2025 identified nuclear power as a central component of the current administration's energy strategy, calling for accelerated reactor deployment, reforms to the Nuclear Regulatory Commission's permitting process, expanded domestic fuel supply, and a long term goal of increasing United States nuclear generating capacity from roughly 100 gigawatts to approximately 400 gigawatts by 2050. The administration has remained strongly supportive of nuclear development while taking a much less favorable stance toward wind and solar, influencing where future policy support may be concentrated. It was also noted that China remains well ahead of the United States in nuclear construction, a gap that could become more important as demand for AI infrastructure continues growing.

Fermi (FRMI) also came up as a cautionary example within this theme. The company entered the public markets with significant political support and a narrative centered on dedicated on site power generation for data centers, ultimately anchored by nuclear generation, yet its shares had continued moving lower. Previous research shared with the group had already raised concerns regarding the management team's history with earlier ventures and its record of executing on stated plans. The discussion reinforced the same lesson that surfaced during earlier conversations involving Agility Robotics (CCXI). Identifying an important industry bottleneck is only part of the investment process. Determining which management teams can successfully execute against that opportunity is ultimately what separates successful investments from unsuccessful ones.

5. Meta, Muse Spark 1.1, and the Case for the Overlooked Mag 7

The frontier model discussion opened with Meta Platforms (META) following the July 9 release of Muse Spark 1.1.

The model is the second-generation system from Meta Superintelligence Labs and represents the company's first major move into API monetization. It is a closed, metered multimodal reasoning model built for agentic tasks,

featuring strong coding performance, a one million token context window, multi-agent orchestration, and aggressive API pricing. Early reporting and benchmark results suggest Muse Spark is competitive with other frontier models, placing Meta back into the conversation after a prolonged period of trailing the leading labs.

The focus then shifted to whether that leadership translates into commercial value. Nearly all of Meta's revenue still comes from advertising, and few expected widespread consumer adoption of a Meta-branded chatbot as a standalone product. The stronger investment case centered on integration. Meta owns Instagram, Facebook, WhatsApp, Marketplace, Quest, and its Ray-Ban smart glasses ecosystem, allowing Muse Spark to be embedded across each platform, with the company also indicating it is expected to replace the existing Llama models powering AI experiences across many of those products. It was noted that Meta's understanding of human behavior and social graphs remains unmatched, making significantly better ad targeting and AI-generated advertising content a more natural application than competing directly in the consumer assistant market. Meta's AI advertising tools were also discussed, with the company reporting that roughly eight million advertisers are now using its AI creative products. A counterpoint was that inference costs rise with usage, meaning heavier engagement with AI increases compute expenses, the opposite of how traditional software businesses typically scale.

Cost per token was discussed as part of the answer. Muse Spark's pricing was viewed as notably competitive among frontier models, reinforcing the expectation that inference would become substantially cheaper over the next several years. One estimate suggested costs could even reach roughly one-tenth of current levels by 2029. Lower unit costs would then encourage greater usage, reflecting the Jevons dynamic the group has discussed at prior meetings while continuing to support the broader AI infrastructure investment thesis regardless of which frontier model ultimately wins.

Valuation was also discussed. Meta trades as a cash-generating business at a relatively reasonable earnings multiple compared with most of the rest of the Magnificent Seven, continues producing enormous operating cash flow despite rapidly expanding capital expenditures, and maintains the financial flexibility to continue investing at scale. The company recently increased its expected 2026 capital expenditure outlook even further, reinforcing

management's willingness to prioritize long-term AI investment. Meta was viewed as offering AI exposure without the more binary risk profile associated with many of the infrastructure bottleneck companies, pairing substantial AI optionality with an already highly profitable core business. Its \$14.3 billion investment for a 49% non-voting stake in Scale AI, along with aggressive compensation packages used to recruit leading researchers, has required tens of billions of dollars in investment. The transaction also brought Scale AI founder Alexandr Wang to Meta, where he became the company's first Chief AI Officer and helped reorganize its AI efforts into Meta Superintelligence Labs. If integrating those capabilities across Meta's existing platforms ultimately creates hundreds of billions of dollars in value, the economics still work even if execution proves uneven.

The focus also shifted to Mark Zuckerberg as an operator. The group has consistently emphasized the idea that CEO quality compounds throughout every aspect of a business, and Zuckerberg was cited as one of the clearest examples of a founder willing to do whatever is necessary to win, whether through internal development, copying successful products, or strategic acquisitions. The Instagram and WhatsApp acquisitions, Meta's successful navigation through multiple platform transitions, and Zuckerberg's decision to reject Yahoo's acquisition offer in the company's early years were all cited as evidence, including the account in Zero to One describing his conviction during those discussions.

Wearables also became part of the discussion. Snap's (SNAP) recently announced smart glasses were revisited from a prior meeting, and those who had used both Apple's (AAPL) and Meta's devices compared their experiences. The general view was that Apple's ecosystem and hardware foundation provide a credible long-term path in the category, particularly if devices become lighter and more affordable, since consumer adoption patterns for Apple hardware differ meaningfully from those for Meta hardware. The discussion also noted the wide range of pricing emerging across the category, from Meta's consumer-focused Ray-Ban offerings to Snap's higher-priced Specs and Apple's

premium Vision Pro. The conversation ultimately expanded into the view that several of the more overlooked members of the Magnificent Seven, particularly Meta, Apple, and Amazon (AMZN), may currently offer more attractive opportunities than some of the names receiving the greatest investor interest because their AI upside is layered onto highly profitable existing businesses with substantial cash flows.

6. Moonshot AI, Kimi K3, and China's Position in the Race

One of the more significant developments over the previous two weeks was Chinese startup Moonshot AI's release of Kimi K3 on July 16.

Independent evaluations placed the model closely behind Claude Fable 5 and GPT 5.6, while outperforming several rivals on specific tasks including front end coding and long context work. Moonshot was notably candid in acknowledging that overall performance still trailed the most capable proprietary systems. K3 was announced as an open weight model, with the weights scheduled for release later in July, while the hosted version became available immediately. Demand proved overwhelming. On July 19, after roughly forty eight hours of exceptionally strong usage, Moonshot paused new consumer subscriptions, citing GPU capacity constraints. Existing subscribers remained unaffected while the company announced plans to reopen new memberships in stages and separate its general Kimi Membership from a dedicated Kimi Code Membership to better allocate compute resources.

That operational detail carried the most investment relevance. A Chinese lab releasing a highly capable model and immediately running into a compute wall is another data point supporting the infrastructure thesis. Regardless of which country or company produces the best model, the binding constraint remains the same.

The market reaction also reinforced that discussion. News surrounding Kimi K3 contributed to renewed debate over whether China's frontier models are closing the performance gap more quickly than previously expected. During the same period, several leading U.S. semiconductor and AI infrastructure companies sold off despite reporting or maintaining strong underlying fundamentals, illustrating how shifts in competitive expectations can create short term volatility even as long term infrastructure demand remains intact.

The ownership angle was reviewed as well. Alibaba (BABA) initially acquired a 36% stake in Moonshot AI. While the company's current ownership percentage has not been publicly disclosed, market estimates generally place Alibaba's remaining stake in the 15% to 20% range following subsequent funding rounds. Alibaba shares had gained roughly 17% during July amid a series of artificial intelligence catalysts, including continued progress with its Qwen models, reports surrounding Apple Intelligence in China, and enthusiasm surrounding its Moonshot investment. The group noted that Alibaba remained below \$300 billion in market capitalization, a fraction of Meta's roughly \$1.6 trillion valuation, while operating one of the world's largest e-commerce and cloud businesses, developing competitive frontier models, maintaining its Moonshot investment, and retaining favorable standing with Beijing. For those wanting exposure to China's AI ecosystem, Alibaba was viewed as one of the more straightforward

investments, with concerns surrounding transparency and financial reporting viewed as more relevant to smaller Chinese companies than to globally prominent technology leaders.

The conversation then shifted to talent flows, producing one of the more interesting discussions of the evening. Moonshot founder Yang Zhilin completed his undergraduate studies at Tsinghua University before earning his PhD at Carnegie Mellon, where he also worked with organizations including Google Brain and Meta. After returning to China, he contributed to Huawei's PanGu model and the Beijing Academy of Artificial Intelligence before founding Moonshot AI. His career reflects a pattern that has become more visible in recent years. Historically, a substantial share of STEM PhD talent feeding frontier laboratories came from China and remained in the United States. That flow has begun shifting as visa policies, the difficulty of starting companies as non-U.S. citizens, and expanding entrepreneurial opportunities in China alter the incentives. Building a company in one's home country can also be significantly easier, with a meaningful cost advantage created by earning revenue in U.S. dollars while operating

elsewhere.

The conversation also recognized that talent flows move in both directions. Someone who spends two decades building expertise inside American institutions before returning home brings substantially more accumulated knowledge than someone who leaves immediately after graduation, making the long term effects far more nuanced than a simple brain drain narrative. The discussion also touched on a noticeable shift in sentiment among younger professionals in China, with confidence in domestic economic opportunities appearing grounded in visible technological progress over relatively short periods of time, influencing where many choose to build their careers.

It was also noted that Chinese citizens do not have access to leading U.S. frontier models, meaning domestic Chinese models serve a potential user base larger than the entire U.S. population. Combined with retained technical talent, that captive demand gives China multiple avenues to remain competitive as the AI race continues evolving. Public commentary surrounding Kimi K3 also reflected a wider range of views, with some investors suggesting that more capable open models could place additional pressure on proprietary model providers by compressing inference economics while benefiting the overall AI ecosystem. The discussion ultimately favored viewing artificial intelligence as an emerging technological rivalry between the United States and China, with substantial value creation remaining possible on both sides as continued competition accelerates innovation across the industry.

7. Education, Math Academy, and What AI Cannot Replace

Education became one of the evening's most thought provoking topics after the group discussed a widely shared account of a third grader scoring a 5 on the AP BC Calculus exam after training through Math Academy.

Math Academy is built around a hand curated knowledge graph of study topics spanning elementary mathematics through advanced university level coursework. The adaptive online platform, launched in its current form in 2021 and priced around \$49 per month, diagnoses what a student knows, selects personalized tasks, and schedules reviews using spaced repetition, effectively emulating many of the decisions an expert tutor would make. Justin Skycak, the company's Director of Analytics, has described the philosophy as optimizing learning efficiency much like optimizing returns in financial markets. The underlying research traces back to Benjamin Bloom's 1984 2 Sigma Problem, which found that students receiving one on one mastery tutoring performed roughly two standard deviations above those taught in traditional classrooms. The platform is not built around generative AI today, instead relying on statistical models and prerequisite mapping, though the framework appears well suited to incorporating frontier AI capabilities over time.

Competition is also beginning to emerge. Khan Academy's Khanmigo has expanded across dozens of schools in Newark, New Jersey, reaching roughly 29,000 students, while Miami Dade has rolled out Gemini for Education to approximately 100,000 high school students. Prior deep dive work on Alpha School was also revisited, where an interesting observation was that much of the reported learning acceleration appeared tied to a carefully designed incentive system that rewarded productive habits, lesson completion, and strong academic performance alongside the AI itself.

The conversation then extended well beyond the investment implications, with one concern centered on the widening gap between public and private education. Chicago Public Schools was discussed as one example, with budget problems, repeated labor disputes, leadership turnover, and governance issues consuming institutional focus. Private schools are generally more agile, can invest more quickly, and often have parent communities that actively push them toward adopting new technology. If AI capability continues compounding, the gap between institutions able to adapt quickly and those that cannot could become substantially larger than in previous educational cycles. At the same time, AI tutoring is becoming inexpensive and widely accessible, potentially allowing motivated students at any income level to learn independently, at their own pace, and in whatever format suits them best.

Concerns also surfaced about the erosion of foundational skills. The discussion mentioned reports of a computer

science course where most students failed the final examination after

relying on AI throughout homework assignments and projects. A generational concern also emerged that children growing up with instant answers may never develop the habit of wrestling with difficult problems long enough to build deep understanding.

Another perspective focused on professions such as law and medicine. Reading legal cases is not primarily about accumulating knowledge but preparing to make difficult judgment calls, and the same principle applies to medicine, where training exists so practitioners can respond appropriately when complications arise. Arithmetic may eventually follow the path of skills that calculators made largely optional, and that may be perfectly acceptable. Judgment, however, cannot be outsourced. The concern is that AI could allow students to bypass the competency development that good judgment depends upon, since the two are far more closely connected than they first appear.

The conversation then turned toward the importance of socialization. Technology may dramatically expand intellectual capability while doing far less to develop interpersonal judgment, such as the ability to read a room and navigate complex social situations. As more interactions move online and become mediated through technology, interpersonal communication may become a scarcer and therefore more valuable skill. The discussion illustrated the point with the example of a hospital executive who was sent home from a new assignment after attempting to force organizational change without first understanding the local culture, a problem no amount of data could have solved. Research on handwriting and knowledge retention was also raised as another example of a human process that technology does not appear to fully replace.

The discussion also touched on product and investing opportunities emerging from these trends. While several major AI platforms now offer supervised access for younger users, none currently provide a purpose built educational experience designed specifically for children from the ground up. Building one well could meaningfully influence public opinion regarding AI in education. Duolingo (DUOL) was highlighted as one company well positioned within that landscape. Software valuations have come down considerably, the company continues generating strong cash flow, trades around a \$6 billion market capitalization, and remains among the leaders in applying AI to educational content creation. The shares remain well below prior highs, although the investment thesis continues evolving as investors weigh the company's AI advantages against the possibility that general purpose AI could also reshape the language learning market. Even so, it remains one of the clearest publicly traded ways to gain exposure to the long term transformation taking place in education.

8. IREN, the Compute Deal, and the RSU Letter

IREN Limited (IREN) returned to the agenda following two developments.

The first was the company's announcement that morning of approximately \$2.8 billion in new multi-year AI cloud services agreements. Reported counterparties included Perplexity, Figure AI, Together AI, Fluidstack, Fireworks AI, Fal AI, Hume AI, and one additional unnamed leading AI developer, while Microsoft (MSFT) and NVIDIA (NVDA) remained significant existing customers. Although not nearly the size of the company's earlier Microsoft agreement, the new contracts provided further evidence that demand for IREN's AI cloud capacity remains strong. The agreements increased the company's year-end 2026 AI Cloud annualized run rate revenue target from \$3.7 billion to more than \$4 billion, with customers prepaying roughly 45% of the associated GPU capital expenditures under contracts averaging about four years in duration. Management also noted that demand from hyperscalers, enterprises, and AI developers continues to exceed both available and planned capacity. Shares responded by rebounding toward the \$40 range after trading in the low to mid-\$30s.

The second was independent board chair David Bartholomew's July 8 letter to shareholders addressing the co-CEO restricted stock unit awards discussed at the prior meeting. The letter defended the grants on retention grounds, arguing that the awards reflected the founders' role in creating substantial shareholder value and were necessary to

retain them for the company's next stage of growth. It also noted that the awards vest over four years, include additional post-vesting holding requirements, prohibit further equity grants before fiscal 2031, and carry no guaranteed value because their ultimate worth depends entirely on IREN's future share price. Public reaction to the letter was mixed, though skepticism remained widespread. Within the group, the prevailing view was that the letter reinforced the original governance concerns rather than resolving them. While few questioned the founders' importance to the business, the size and timing of the awards were viewed as unusual for a company still in the early stages of executing its long-term AI infrastructure strategy. The grants were seen as compensation more commonly associated with companies that have already established a longer record of execution, making the board's emphasis on retention an insufficient response to those concerns. Even so, the underlying investment thesis remained tied primarily to the company's execution, with strong demand for AI compute capable of outweighing governance concerns if management continues delivering on its growth plans. The continued focus on executive compensation likely contributed to the weakness in the shares leading into the compute announcement.

The framing from the prior meeting remained intact. Governance concerns do not automatically invalidate the underlying investment thesis given IREN's power portfolio, expanding data center footprint, and exposure to AI infrastructure demand. They do warrant continued evaluation of

management's capital allocation decisions and the board's approach to governance as the company continues executing its long-term strategy.

9. Lightning Round

Truth Social and the paid feed: Trump Media (DJT) has discussed offering a subscription product providing faster access to the President's posts, with reporting suggesting institutional API pricing up to \$100,000 per month for high speed access and roughly \$60,000 per month under a three year commitment. The economics appeared difficult to justify from the outset. Trump Media generates roughly \$3.7 million in trailing twelve month revenue against a market capitalization of roughly \$2.3 billion to \$2.7 billion, and even a generous total addressable market in the range of \$30 million to \$200 million would still imply a demanding valuation for a business with no demonstrated ability to monetize its platform. Firms such as Citadel Securities, Jane Street, and other established quantitative trading firms already devote enormous resources to capturing and arbitraging even the smallest latency advantages, making it difficult to see how a durable informational edge could exist in practice. The general view was that this is a branding exercise rather than a business, and that offering a three year contract tied to a sitting president whose term ends before the contract expires raises obvious questions about the product's long term value proposition.

Railroads and equipment rental: Canadian Pacific (CP) and Norfolk Southern (NSC) were discussed as attractive dividend and margin driven positions with a secondary AI angle. Data center construction requires moving enormous quantities of heavy equipment, and rail is frequently the only practical way to get it to site. This also tied back to the proposed Union Pacific (UNP) acquisition of Norfolk Southern discussed at the July 6 meeting, though the investment thesis centered primarily on long term cash flow rather than merger speculation. The same second order logic extends to equipment rental businesses, including United Rentals (URI) and the smaller rental name (EQPT) previously raised in the group. These businesses are unlikely to deliver outsized returns, but they offer solid cash flow supported by a durable multi year trend.

Memory and the Korean market: SK Hynix (SKHY) completed its U.S. ADR offering, raising roughly \$26.5 billion at approximately a \$1 trillion valuation, and the stock has since given back its initial pop. The more interesting discussion concerned the Korean market itself, which has been selling off hard. Retail participation runs roughly 60% to 70% of Korean market activity versus 20% to 30% in the U.S., while more than one million retail accounts have reportedly received margin calls during the recent selloff. Korean retail also concentrates heavily in U.S. stocks, including IREN and other high beta favorites. At the same time, SK Hynix and Samsung

represent a large share of the Korean market, making the country's performance closely tied to the memory sector. Selling pressure in Korea also weighed on U.S. memory companies such as SanDisk (SNDK) and Micron (MU). Memory fundamentals were still viewed as solid, with expectations for a sharp recovery once forced selling pressure subsides, while recognizing that volatility may remain elevated. Korea was also viewed as a useful case study for how markets behave when retail participation, leverage, and social trading play an outsized role in price discovery.

eVTOL and aerial robotics: Archer Aviation (ACHR) and Joby Aviation (JOBY) are both in the final stages of FAA testing. Archer rose nearly 20% on the day of the meeting following an Anduril partnership unveiling an autonomous VTOL platform, effectively repositioning the company as a defense play. Archer also holds an Olympic sponsorship tied to LA28 with plans to move attendees from hotels to venues, an Uber partnership, and government contracts. The company plans to unveil safety critical autonomy and air traffic control capabilities later in 2026 tied to its partnership with Palantir (PLTR). The sector wide story is that flying taxi companies are quietly repositioning as aerial robotics companies, a shift that could ultimately support higher valuations if the market begins viewing them through a defense and autonomy lens. The group treated this with appropriate skepticism. These remain essentially pre revenue companies with heavy regulatory dependency, contracts recognized far into the future, and valuations that require ignoring current fundamentals entirely. There is such a thing as being too early, and eVTOL likely still sits in that zone, though the defense angle could accelerate commercialization given the incentives created by drone warfare.

Rare earths: USA Rare Earth (USAR) was revisited as the leading domestic company pursuing a vertically integrated mine to magnet strategy. The company remains pre revenue, and China controls roughly 70% to 90% of the processing chain, which is precisely the strategic argument for building domestic capacity. Cameco (CCJ) was discussed as an adjacent uranium exposure alongside broader discussion of mining companies that could benefit as Western supply chains diversify away from China. For those comfortable with the risk profile, USAR was viewed as a speculative position that could respond sharply to geopolitical developments while carrying genuine risk of going to zero.

AMC and the Nolan effect: AMC Entertainment (AMC) surged roughly 27% following reports of the highest quarterly revenue and adjusted EBITDA in the company's 106 year history. Investor enthusiasm was further supported by Christopher Nolan's *The Odyssey*, which pushed audiences toward IMAX (IMAX) screens, with only 41 true 70mm venues operating globally, including one at a local Cinemark (CNK) location. Those who had seen the film were positive on it, and Polymarket around the meeting date priced *The Odyssey* near 48% to win Best Picture, which several found notable given that awards contenders typically release later in the year. The

valuation discussion was less encouraging. AMC carries debt and lease obligations approaching \$8 billion against a market capitalization well below that level, while continuing to operate with negative book value. The company has diluted shareholders aggressively for years. Exhibitors generally retain roughly 40% to 50% of ticket revenue over the full theatrical run, with concessions driving a disproportionate share of profitability. The view remained that AMC works better as a trading vehicle around seasonal catalysts than as a long term investment.

Peptides and Hims and Hers: At the time of the meeting on July 20, the FDA's Pharmacy Compounding Advisory Committee was scheduled to meet on July 23 and 24 to consider several peptides for inclusion on the 503A bulks list, with the committee providing advisory recommendations rather than final authority. The current administration's health leadership had expressed support for expanding peptide access while reshaping the committee, making the meeting a closely watched catalyst. Hims and Hers (HIMS) was viewed as one of the best positioned public companies to benefit if additional peptide compounding ultimately receives regulatory support. GLP-1 compounding was viewed as less likely to remain available following the resolution of nationwide shortage conditions, though HIMS had already begun pivoting through branded partnerships and other initiatives. Revenue projections discussed during the meeting, based on a widely circulated public estimate from a social media commentator, suggested peptide related revenue could eventually reach \$10 billion to \$19 billion annually by 2030. Those projections were viewed with considerable skepticism given both the source and the long time horizon. Scenario analysis during the

meeting suggested a partial approval scenario appeared more likely than broad approval across every peptide under review, leaving the overall setup constructive while recognizing that the regulatory process would likely extend well beyond the committee meeting.

Closing Thoughts

The July 20 meeting reinforced several themes that have surfaced repeatedly throughout recent sessions while adding new perspective as the investment landscape continues evolving.

The first is that artificial intelligence remains far larger than the models themselves. Whether the discussion centered on Bloom Energy's fuel cells, Oracle's Project Jupiter, power generation, IREN's expanding compute business, or Moonshot AI's rapid capacity constraints following the release of Kimi K3, the same conclusion repeatedly emerged. Compute demand continues growing far faster than the infrastructure required to support it, leaving power, data centers, networking, and physical supply chains among the most important determinants of how quickly the industry can expand.

The second is that execution continues separating successful investments from attractive narratives. The meeting repeatedly highlighted the importance of management quality, capital allocation, incentives, and governance. Strong industries alone do not guarantee strong investments, while exceptional operators can often create value even in highly competitive markets. That principle surfaced across discussions ranging from Meta and IREN to Bloom Energy, Fermi, and several companies examined during the lightning round.

Another recurring theme was that competition within artificial intelligence appears to be expanding rather than narrowing. Meta's renewed push at the frontier, Moonshot AI's rapid progress in China, Alibaba's positioning within that ecosystem, and the continued evolution of open models all suggest that leadership is becoming more competitive. At the same time, companies with established distribution, existing customer relationships, and durable cash- generating businesses may ultimately be better positioned to translate AI capabilities into long- term shareholder value than those relying solely on model leadership.

The evening also reinforced that technological progress does not eliminate the importance of human judgment. The education discussion highlighted both the extraordinary opportunities created by AI-assisted learning and the risks of outsourcing critical thinking, while conversations surrounding geopolitics, market structure, governance, and investing all returned to the same underlying idea: better information does not automatically produce better decisions. Experience, judgment, incentives, and thoughtful analysis remain difficult to replace.

Finally, the meeting served as another reminder that many of today's most compelling investment opportunities no longer fit neatly within a single sector. Artificial intelligence, energy infrastructure, semiconductors, education, healthcare, geopolitics, manufacturing, and capital markets continue influencing one another in ways that would have seemed far less connected only a few years ago. Understanding those relationships, rather than viewing each development in isolation, remains one of the group's primary objectives.

For those who were unable to attend or would like to revisit the discussion, a recording of the July 20 meeting is available on YouTube (<https://www.youtube.com/watch?v=ISW4ALAfHPQ>).

We look forward to reconvening on Monday, August 3, 2026, to continue building on many of the themes explored throughout the evening.

Attendance and Acknowledgements

A sincere thank you to everyone who attended and contributed to another outstanding discussion.

Attending members, listed alphabetically by last name:

Afeef Akhtar
Michelle Maleski
Diana Ascencio (media)
Omar Pathan
Quinn Basta (recording, transcription, photography)
Miguel Salgado
Connor Darrow
Mario Sanchez (videography, media)
John Donners
Victor Sanchez (founder, media)
Mohammed Haq (recording, transcription, attendance)
Moses Shreim
Brian Jung
Eric Simpson (founder, moderation, summary, AI systems)
Sandeep Kaza
Jeffrey Tantoc
Suryansh Khatikar
Jordan Wilson
James Lewan
Abhishikth Yarlagadda

DISCLAIMER

Nothing discussed during this meeting or contained within this summary constitutes financial advice or a recommendation to buy or sell any security. All discussion is intended solely for educational and informational purposes.