THE INVESTING GROUP

July 6, 2026 Meeting Summary

July 6, 2026 | Chicago | By TIG

We held the latest meeting of The Investing Group on Monday, July 6, 2026, at the Union League Club of Chicago in Room 816.

Introduction

We held the latest meeting of The Investing Group on Monday, July 6, 2026, at the Union League Club of Chicago in Room 816.

The meeting welcomed several new attendees while continuing the collaborative discussion format that has become a hallmark of the group. Topics ranged from artificial intelligence, semiconductors, robotics, healthcare, and transportation to geopolitics, market structure, consumer technology, and individual investment ideas. Although the subjects varied widely, they consistently reinforced many of the long-term themes that have emerged throughout recent meetings.

One recurring observation was that execution continues separating long-term winners from the rest of the field. Whether discussing AI infrastructure, enterprise software, healthcare innovation, robotics, or data centers, the conversation repeatedly returned to the importance of capital allocation, management incentives, vertical integration, and disciplined execution alongside technological leadership. Another theme centered on looking beyond the most visible companies to identify the suppliers, infrastructure providers, and enabling technologies supporting broader secular trends. The meeting also reinforced how artificial intelligence continues expanding beyond software into industries such as healthcare, manufacturing, logistics, transportation, and national security, creating investment opportunities across an increasingly interconnected economy.

As always, the objective was not to predict short-term market movements but to better understand the long-term trends, competitive dynamics, and structural changes shaping the investment landscape.

1. Small-Cap Speculation, Market Structure, and the SDOT Trading Frenzy

The meeting began by examining one of the most remarkable speculative trading stories of the past several weeks, Sadot Group (SDOT), which provided another example of how today's market structure can produce extraordinary price swings in low-capitalization companies.

The stock initially surged from roughly $2 per share into the $40 range during early June before quickly falling back into the single digits. Much of that initial move followed a 1-for-20 reverse split that dramatically reduced the company's public float, illustrating how limited share availability can amplify buying pressure once momentum begins to build. The volatility continued into late June and early July as shares staged another remarkable rally, briefly climbing above $100 intraday on July 2 before giving back a substantial portion of those gains

over the following trading sessions. Rather than focusing on the company's underlying agricultural operations, the conversation centered on what the trading action revealed about today's speculative trading environment and the growing importance of market mechanics.

The financing dynamics common throughout much of the micro-cap universe also became part of the discussion.

Many smaller companies continue relying on equity issuance as a primary source of capital, creating incentives that do not always align with long-term shareholder value. Reverse splits, name changes, secondary offerings, acquisitions financed largely with stock, and repeated capital raises remain common across portions of the market, allowing companies to repeatedly access public capital even when the underlying business changes very little. While these characteristics often introduce meaningful long-term risks for investors, they also help create the sharp price movements that attract momentum-oriented traders.

The recent elimination of the Pattern Day Trader rule was also revisited as another development that could influence speculative activity. Although it remains far too early to evaluate its full impact, lowering barriers to active trading may encourage additional participation in the areas of the market already known for extreme volatility. With implementation expected to occur gradually across brokerages, it will likely take time before the full impact of the new framework can be properly assessed.

Rather than viewing SDOT as an isolated event, the group considered it another example of structural changes taking place across today's markets. Liquidity, float dynamics, retail participation, and momentum can often outweigh traditional valuation metrics for extended periods, particularly among companies with exceptionally small public floats. Understanding those market mechanics has therefore become just as important as evaluating the underlying business when analyzing highly speculative securities.

The discussion reinforced a recurring theme from previous meetings. In today's market, stock prices do not always reflect business quality. Examining incentives, financing structures, dilution risk, float characteristics, and investor psychology often provides a clearer explanation for why speculative companies experience such dramatic trading behavior.

2. Iran, Oil Markets, and Why Markets Continue Looking Beyond the Headlines

The group revisited the evolving situation involving Iran amid the ongoing state funeral for Supreme Leader Ali Khamenei, continued uncertainty surrounding the country's new leadership, and renewed diplomatic efforts aimed at preserving the fragile ceasefire and keeping the Strait of Hormuz open.

Discussion began with the latest developments surrounding negotiations between U.S. and Iranian officials following the June memorandum of understanding intended to end the conflict and restore commercial shipping through the region. While geopolitical headlines continued arriving almost daily, it was observed that financial markets had become far less responsive than they were during the initial stages of the conflict several months earlier.

Focus then shifted toward global energy markets. During the height of the conflict, front-month crude oil futures surged well above $100 per barrel for several months, with the benchmark peaking near $120 in April as investors feared prolonged disruptions to one of the world's most important shipping routes. By the time of the meeting, however, prices had fallen back below $70 as expectations shifted toward normalization. Improving shipping activity through the Strait of Hormuz, strategic petroleum reserve releases, stronger production from North and South American producers, and weaker global demand all contributed to the market's changing outlook, although it was also noted that shipping conditions remained fragile and inventories would likely take time to fully recover.

The conversation also briefly touched on China's role within the evolving energy picture. More broadly, slowing Chinese oil demand, the country's rapid buildout of renewable energy, expanding electric vehicle adoption, and other structural shifts in energy consumption continue shaping long-term expectations for global oil demand. Together with ongoing geopolitical developments, these trends illustrate the range of factors influencing energy markets over both the short and long term.

The group also revisited a point discussed during previous meetings: markets discount future expectations rather than

current events. As of the meeting, although negotiations surrounding sanctions, shipping access, regional security, and potential shipping fees remained unresolved, investors appeared to believe that the most economically disruptive phase of the conflict had already passed. Unless developments materially altered long-term expectations, geopolitical headlines gradually lost much of their ability to influence financial markets.

The conversation also examined how changing market composition may influence the relationship between geopolitical events and equity performance. It was noted that today's equity market is considerably more concentrated in technology companies than in previous decades. Many of the largest technology firms maintain substantial cash flows and stronger balance sheets than the industrial and manufacturing companies that once dominated major indices, potentially making the broader market somewhat less sensitive to temporary increases in inflation or energy prices.

Questions were also raised regarding broader geopolitical incentives. With U.S. midterm elections approaching, it was suggested that multiple parties may have strong incentives to

prevent the conflict from expanding significantly, reinforcing expectations that negotiations would ultimately produce greater stability.

Overall, the discussion reflected the view that geopolitical developments remain important, but markets continue distinguishing between temporary disruptions and structural changes. While uncertainty surrounding Iran had not disappeared, investors appeared to have shifted much of their focus back toward longer-term themes such as artificial intelligence, infrastructure investment, and technological innovation.

3. Anthropic, AI Security, Palantir, and the Strategic Importance of Data

Artificial intelligence once again became one of the evening's longest discussions following the restoration of Anthropic's frontier AI models after their temporary suspension by the U.S. government.

The conversation began with attendees comparing their experiences using Claude Fable 5 following its return. Those who had used the model extensively discussed differences relative to previous Claude versions, while others focused more broadly on what the temporary restrictions revealed about the strategic role frontier AI systems now occupy.

The discussion revisited the government's June decision to suspend access to Claude Fable 5 shortly after its launch, with restrictions ultimately lasting until early July. It was noted that while Fable 5 returned to general availability following the lifting of the export controls, access to Mythos 5 remained more limited through approved cybersecurity programs. The conversation framed the episode as another example of frontier AI models evolving into strategic national assets instead of conventional software products. Leopold Aschenbrenner's Situational Awareness framework was referenced as a useful lens for understanding how advanced AI capabilities may become more closely intertwined with national security, export controls, and government oversight.

Additional context surrounding the suspension was also discussed. The government's actions followed reports that researchers had demonstrated methods of bypassing portions of the model's safety safeguards to assist with identifying software vulnerabilities and generating exploitation code. Anthropic publicly disagreed with the decision, arguing that similar capabilities exist across multiple frontier AI systems and that applying the same standard broadly could significantly slow future frontier model releases. The episode reinforced the view that the government's response reflected not only the capabilities of a single model, but also the growing strategic importance of frontier AI more generally.

This naturally led into a discussion surrounding data ownership, cybersecurity, and enterprise AI. Considerable focus centered on Palantir Technologies (PLTR), whose software architecture was contrasted with the approaches taken by many frontier AI developers. Discussion emphasized Palantir's approach of allowing organizations to maintain ownership and control of their proprietary data while integrating AI capabilities, an architecture viewed as

fundamentally different from relying exclusively on third-party frontier models.

The conversation also examined Palantir's long-standing relationships with government agencies, defense organizations, and intelligence customers. Rather than focusing solely on the company's financial performance, discussion emphasized its strategic positioning, customer lock-in, and the importance of software platforms that become deeply embedded within mission-critical operations. These characteristics align closely with many of the monopoly- building principles described by Peter Thiel in Zero to One, including ecosystem control, switching costs, and durable competitive advantages.

Valuation naturally entered the discussion as well. Although Palantir continued trading at historically elevated earnings multiples, it was also noted that the stock had already experienced considerable volatility during the first half of the year despite continued fundamental strength. Traditional valuation techniques may prove less useful for businesses operating near the center of rapidly evolving technological shifts, though investors should continue expecting meaningful volatility as markets attempt to price opportunities whose long- term economic potential remains difficult to quantify.

The conversation then turned to international AI competition following Anthropic's allegations that operators affiliated with Alibaba (BABA) had attempted to distill capabilities from its frontier models through large numbers of fraudulent accounts. Alibaba disputed the allegations, but the reports nevertheless illustrated the powerful strategic incentives surrounding frontier AI development. The discussion highlighted China's continued progress across artificial intelligence, particularly through rapidly improving open-source models, while also noting that competition between the United States and China now extends well beyond semiconductors into AI models, cloud infrastructure, research talent, and cybersecurity.

Those who had firsthand experience using Alibaba's AI models also commented favorably on their capabilities, particularly in areas such as multilingual translation and open-source performance. The discussion reinforced the idea that while U.S. companies continue leading many frontier AI benchmarks, Chinese models have advanced rapidly, making long-term competition more balanced than many investors may assume.

The conversation ultimately returned to a recurring investment theme that has surfaced throughout recent meetings. Artificial intelligence is evolving into an infrastructure race as much

as a software race. Compute, cybersecurity, proprietary data, enterprise workflows, government relationships, and distribution channels all appear likely to become just as strategically important as the AI models themselves, with companies controlling those foundational layers potentially capturing a significant share of the long-term economic value created by artificial intelligence.

4. AI Infrastructure, Compute Economics, and the Scale of the Next Investment Cycle

Artificial intelligence infrastructure remained one of the evening's central themes as the discussion shifted from individual AI models toward the enormous amount of capital required to support the industry's continued expansion.

The conversation revisited the accelerating investment cycle taking place across hyperscalers, cloud providers, semiconductor manufacturers, and governments. The AI race was described as resembling a large-scale industrial buildout requiring unprecedented investment in compute, networking, memory, data centers, cooling, and power generation. Rather than viewing today's spending as a temporary surge, the conversation suggested that the current investment cycle may still be in its early stages as more capable AI models continue requiring exponentially larger amounts of computational resources. Supporting that perspective, the largest hyperscale technology companies have continued raising capital expenditure plans, with combined spending expected to reach record levels over the coming years.

One topic centered on the economics behind frontier AI models and why they remain so expensive to operate. Unlike traditional software businesses, each interaction with a frontier model consumes meaningful compute resources, making inference costs an important consideration as models become more capable. The discussion explored how reasoning models require substantially greater computational effort than standard chatbot interactions, with more advanced reasoning often consuming significantly larger numbers of tokens before generating a response. As AI agents become more widely deployed across enterprise applications, this growing computational intensity was viewed as another factor likely to support continued infrastructure investment.

This naturally led into discussion surrounding long-term capital expenditure requirements. The conversation explored the possibility that current forecasts may still underestimate how much infrastructure ultimately needs to be built if artificial general intelligence or superintelligence continues progressing as many industry leaders expect. Recent upward revisions to hyperscaler capital expenditure guidance, including additional spending on data center construction and AI infrastructure, were viewed as reinforcing that perspective. Across a wide range of potential

outcomes, semiconductors, networking equipment, memory, cooling systems, and electrical infrastructure remain among the industry's most important bottlenecks.

The group also acknowledged the primary downside risk to this thesis. If AI capabilities fail to improve as quickly as expected or if enterprise adoption slows materially, today's extraordinary infrastructure spending could eventually prove excessive. Elevated capital expenditures have already placed pressure on free cash flow for some technology companies, leading investors to debate how quickly these investments will generate attractive returns. At present, however, the discussion noted little evidence that demand for compute is weakening. Instead, successive announcements involving new data centers, cloud deployments, and infrastructure partnerships continue reinforcing expectations that investment remains constrained more by available supply than by customer demand.

Overall, the discussion reinforced one of the group's longest-running themes: artificial intelligence remains as much an infrastructure story as a software story. While AI models attract much of the public attention, compute, semiconductors, networking, memory, power generation, cooling, and data centers continue serving as the foundational assets supporting the industry's long-term growth.

5. OpenAI's Jalapeno Chip, Vertical Integration, and the Future of AI Hardware

The discussion then shifted toward OpenAI's recently unveiled Jalapeno inference processor, developed in partnership with Broadcom (AVGO). While the collaboration between the two companies had been announced previously, the June unveiling of the chip itself marked another step in OpenAI's effort to build a more vertically integrated AI platform.

Rather than viewing the announcement as an immediate competitive threat to NVIDIA (NVDA), the conversation focused on what it may signal about the industry's longer-term direction. Jalapeno is not a general-purpose GPU but a custom inference accelerator built specifically for large language model workloads. Although it was designed to support OpenAI's inference infrastructure, the architecture is intended to be flexible enough to work across different LLMs. The announcement nevertheless illustrates how leading AI developers are investing more heavily in specialized hardware optimized for their own software ecosystems.

The discussion also explored the growing importance of vertical integration. As AI companies continue expanding, many of the industry's largest organizations may seek greater control across multiple layers of the technology stack, including chips, networking, cloud infrastructure, software, and AI models. Companies capable of integrating these systems internally may ultimately develop stronger competitive advantages while reducing reliance on third-party suppliers.

Customized AI hardware also highlights the evolving competitive landscape surrounding AI acceleration. While NVIDIA's position remains substantial, inference processors, custom silicon, and alternative accelerator architectures continue becoming more important as AI workloads diversify. Broadcom, Google's Tensor Processing Units (TPUs), Amazon's (AMZN) Trainium platform, Cerebras, and other emerging hardware developers illustrate the range of approaches being pursued to accelerate AI workloads rather than directly replace traditional GPUs. Broadcom also stands out because its growing custom silicon business positions the company to benefit regardless of which leading AI developer ultimately succeeds with its own accelerator strategy.

One aspect of Jalapeno that stood out was the reported pace of development. The chip progressed from early design work to fabrication readiness in roughly nine months, substantially faster than many traditional semiconductor development cycles. OpenAI has also stated that its own AI models assisted portions of the engineering process, offering an early example of how AI itself may begin accelerating future chip development.

The economics of custom AI hardware also remain an important part of the investment thesis. While detailed performance data has not yet been released publicly, OpenAI has indicated that Jalapeno is designed to improve inference efficiency, and reports have suggested meaningful reductions in inference costs may be possible if the platform performs as expected. If successful, proprietary hardware could become an important component of improving long-term unit economics as frontier AI models continue scaling.

At the same time, the competitive landscape remains challenging. Building advanced semiconductor platforms requires significant engineering expertise, capital investment, and sustained execution across multiple product generations. Even with faster development cycles, creating competitive alternatives to established GPU ecosystems remains an exceptionally difficult undertaking given the software, manufacturing, and ecosystem advantages already built by incumbent providers.

The discussion also turned toward Elon Musk's broader ecosystem strategy. Tesla (TSLA), xAI, energy infrastructure, robotics, and autonomous vehicles were cited as examples of how vertical integration may become a valuable competitive advantage throughout the AI industry. Rather than viewing these businesses independently, they were discussed as components of a larger ecosystem capable of sharing technology, infrastructure, and data across multiple business lines.

A wider investment theme also emerged. Rather than focusing exclusively on today's market leaders, the conversation emphasized identifying businesses that are gradually expanding across multiple layers of the AI value chain. Companies capable of combining software, hardware,

cloud infrastructure, networking, and proprietary AI systems may ultimately possess more durable competitive advantages than businesses operating within only a single segment of the ecosystem.

While it remains too early to judge Jalapeno's long-term commercial impact, the announcement may provide an early indication of where AI infrastructure is headed. As more leading AI developers pursue custom silicon alongside proprietary software and cloud infrastructure, vertical integration could play a larger role in shaping future competition.

6. AI Drug Discovery, AlphaFold, and the Future of Healthcare Innovation

Healthcare once again featured prominently throughout the evening as the conversation explored several recent examples of artificial intelligence being applied to pharmaceutical research and personalized medicine.

The discussion began with a widely shared claim from computational biologist Douglas Yao that went viral on X roughly one week before the meeting, where he reported using AI-assisted workflows from a laboratory built in his home garage to develop a potential Alzheimer's treatment known as PAC-832 (Douglas Yao's X post). According to his public posts, AI was used throughout the discovery process to help design thousands of candidate molecules,

automate portions of laboratory screening through AI-programmed robotics, and accelerate early-stage research. The announcement generated substantial engagement across the platform, although the work itself remains self-reported, has not undergone independent peer review, and is still in the early stages of laboratory testing.

The conversation then examined the substantial gap separating encouraging preclinical findings from successful human therapies. Alzheimer's disease has historically produced one of the lowest clinical success rates in pharmaceutical research, with the overwhelming majority of experimental treatments failing before reaching patients. Although AI may substantially accelerate molecular discovery and candidate identification, successful commercialization still requires years of laboratory validation, clinical testing, regulatory review, and large-scale manufacturing before new therapies become widely available.

Rather than focusing solely on one specific drug candidate, the discussion expanded to the broader implications of AI-assisted research. If artificial intelligence continues improving scientific productivity, even modest gains in research efficiency could meaningfully shorten development timelines across many areas of medicine.

This naturally led into another widely discussed example involving machine learning engineer Paul Conyngham, who used AI tools, including AlphaFold, while collaborating with academic

researchers to help develop an experimental personalized mRNA treatment for his dog Rosie, which reportedly reduced the tumor by roughly 75 percent. Although the reported results were encouraging, the case remained a single experimental example conducted within a veterinary research setting under university oversight and illustrated how AI can accelerate portions of the discovery process while still relying on scientific and institutional expertise.

The group also discussed AlphaFold, developed by Alphabet's (GOOGL) Google DeepMind, and the growing role of computational biology. Its ability to predict protein structures has significantly expanded the capabilities available to researchers, and the work behind AlphaFold was recognized with the 2024 Nobel Prize in Chemistry. At the same time, it was emphasized that protein structure alone does not fully determine biological function, and translating computational predictions into effective human therapies still requires extensive laboratory validation and clinical testing. The continued development of Isomorphic Labs, Alphabet's drug discovery company spun out of DeepMind, also illustrates how AI-driven research is being paired with the clinical development, regulatory, and commercialization capabilities of established pharmaceutical companies through partnerships with several major drug manufacturers.

The conversation also examined how AI could reshape the competitive dynamics of drug discovery. While smaller teams and independent researchers may become capable of identifying promising therapeutic candidates using AI-assisted workflows, large pharmaceutical companies continue possessing substantial advantages in clinical development, manufacturing, regulatory expertise, commercialization, and global distribution. As a result, many successful discoveries originating outside traditional pharmaceutical organizations could ultimately become partnership or acquisition opportunities for larger healthcare companies capable of bringing new therapies through the full approval process.

The discussion then expanded into longevity, preventative medicine, and healthcare optimization. Artificial intelligence was viewed as a tool capable of accelerating research across genomics, diagnostics, personalized medicine, and disease prevention. While many of these applications remain years away from widespread adoption, they illustrate how AI could meaningfully increase the pace of medical innovation over the coming decades.

Overall, the conversation emphasized maintaining a balanced perspective. Headlines surrounding AI-generated drug discoveries naturally attract significant public interest, yet scientific progress still depends upon rigorous experimentation, independent validation, and successful clinical outcomes. Even so, the examples discussed suggest that artificial intelligence has already begun reshaping biomedical research in ways likely to become more significant over time.

7. Social Media, Gambling, and the Economics of Digital Addiction

The discussion then shifted toward several consumer-facing industries shaped by digital engagement, behavioral incentives, and platform economics.

One topic centered on TikTok's ongoing social media addiction litigation. The conversation extended beyond the legal proceedings themselves to examine platform design, user engagement, and the economic incentives driving many social media businesses. The discussion occurred against the backdrop of several recent legal developments. Earlier in the year, a California jury found Meta Platforms (META) and Alphabet's (GOOGL) YouTube liable in one bellwether social media addiction case, while TikTok reached multiple pretrial settlements, including another on June 30, reflecting how litigation surrounding platform design and user engagement continues advancing.

Discussion then turned to how short-form video platforms optimize user engagement through recommendation algorithms and continuous content feeds. Comparisons were drawn between these engagement systems and behavioral techniques historically associated with casinos and slot machines, where variable rewards encourage users to continue interacting with a platform for extended periods of time.

The conversation also examined how growing public scrutiny may influence future regulation. Although individual legal settlements may have relatively modest financial effects on companies with substantial resources, the discovery process surrounding these lawsuits could reveal additional information regarding platform design and user engagement strategies. With thousands of individual cases and hundreds of school district lawsuits still pending, the legal precedents established over the coming years may ultimately prove more significant than the immediate financial costs associated with individual settlements.

The conversation naturally expanded into sports betting following the federal indictment of former NBA players Malik Beasley and Ed Davis on conspiracy and bribery charges related to an alleged player prop betting scheme. The case was viewed alongside other recent gambling- related investigations as another example of the regulatory and integrity challenges accompanying the rapid expansion of legalized sports wagering throughout the United States.

Rather than viewing gambling as an isolated industry, the discussion connected it to a wider pattern involving prediction markets, online betting, financial speculation, and other gamified digital experiences. It was observed that many platforms now compete for the same limited resource: user engagement. Whether through financial markets, sports betting, social media, or prediction markets, behavioral incentives continue influencing how people allocate both time and capital.

The discussion also explored possible second-order investment implications. As gambling, social media, and other digitally addictive behaviors continue expanding, demand for mental health services, addiction treatment, behavioral healthcare, and preventative wellness may also grow over time. Adjacent healthcare markets were discussed as potential long-term beneficiaries of these societal trends.

The meeting closed the topic by recognizing that regulation has historically tended to follow technological adoption rather than precede it. As with previous technological innovations, governments will likely continue balancing consumer protection with innovation while adapting legal frameworks to technologies evolving faster than traditional regulatory systems.

8. Robotics, Physical AI, and the Next Layer of Infrastructure

The conversation then returned to robotics following several recent developments involving humanoid hardware, robotic components, and industrial automation.

One topic centered on Proception, a robotic hand startup founded by former Tesla (TSLA) Optimus technical lead Jay Li. Shortly before the meeting, the company announced an $11 million seed funding round, introduced its

ProHand 1.0 robotic hand, and settled litigation involving Tesla. Although the news itself generated interest, the discussion focused more on what it revealed about the current state of humanoid robotics. Much of the conversation centered on the view that while meaningful progress has been made in mobility, perception, and artificial intelligence, dexterous robotic hands remain one of the industry's most difficult engineering challenges.

The discussion emphasized that humanoid robotics is not simply a question of intelligence. Consistently manipulating real-world objects with speed, precision, and reliability remains one of the primary obstacles to broader commercial deployment, with fine motor control, tactile sensing, durability, and manufacturing costs continuing to limit what today's humanoid systems can accomplish outside carefully controlled environments. The conversation also revisited observations from the recent Automate conference in Chicago, where several attendees spent time exploring the latest developments in industrial automation and robotics. Seeing a wide range of humanoid robots operating in person reinforced that while mobility and AI capabilities have advanced considerably, dexterous hands and reliable object manipulation remain among the industry's most significant engineering challenges. As the robotics ecosystem continues evolving, specialized components such as robotic hands may ultimately become a supplier market of their own rather than being developed entirely in-house by every humanoid manufacturer. Many commercial deployments today also remain concentrated in warehouse

automation, manufacturing, logistics, machine vision, and other specialized industrial environments as the technology continues maturing.

The discussion then shifted toward the investment implications of these trends. Rather than attempting to identify the eventual winner among humanoid robot manufacturers, the conversation explored companies supplying the broader robotics ecosystem. Vision systems, industrial sensors, actuators, bearings, motion control components, cameras, and specialized semiconductor companies were all highlighted as potential second-order beneficiaries regardless of which humanoid platform ultimately achieves widespread adoption.

Companies such as Cognex (CGNX), Teradyne (TER), Ambarella (AMBA), along with Japanese industrial leaders including Fanuc, Keyence, Yaskawa, THK, and Nidec, were referenced as examples of businesses already supplying important technologies across industrial automation. Rather than depending entirely on the success of a single robotics company, these suppliers may benefit from continued automation across a wide range of industries.

The conversation also acknowledged the speculative nature of many publicly traded robotics companies. Recent announcements surrounding Agility Robotics' planned SPAC transaction (CCXI), along with a growing pipeline of proposed humanoid robotics listings in Asia, illustrated both the strong investor enthusiasm surrounding physical AI and the challenges of valuing companies that remain in the early stages of commercialization. While some companies have begun deploying commercial systems with enterprise customers, long-term investment outcomes will ultimately depend on execution, scaling production, and achieving sustainable economics rather than promotional narratives.

The discussion also explored the likely pace of adoption. Although many attendees remain highly optimistic about robotics over the coming decade, several suggested that widespread commercialization may still be several years away. Current technological progress appears encouraging, but deployment is likely to occur gradually rather than all at once. As a result, enabling technologies may continue offering attractive risk-adjusted opportunities while the competitive landscape among humanoid platform developers continues evolving.

The discussion returned to a theme that has surfaced repeatedly throughout recent meetings. Artificial intelligence appears to be extending beyond software and into the physical economy. While robotics may still be in the early stages of its adoption cycle, continued advances in AI models, semiconductors, sensors, batteries, and industrial automation suggest that physical AI could become one of the next major waves emerging from today's infrastructure buildout.

9. IREN, AI Infrastructure Governance, and Capital Allocation

The conversation then shifted toward IREN Limited (IREN), which generated one of the evening's more detailed discussions surrounding management quality, capital allocation, and corporate governance.

Discussion began with the company's recently announced multi-year global partnership with the Golden State Warriors. The agreement, reportedly worth more than $50 million annually and described as the richest sponsorship deal in North American team sports, will place the IREN logo on Warriors jerseys beginning with the 2026-27 NBA season while also extending across the Golden State Valkyries, Santa Cruz Warriors, and Chase Center. While the partnership substantially increased IREN's visibility, questions were raised regarding whether a large consumer-facing sponsorship represents the most effective use of shareholder capital for a business whose primary customers are enterprise AI and cloud infrastructure companies. It was also noted that management has stated the sponsorship is intended to strengthen IREN's presence within the Bay Area AI ecosystem, where many potential customers and startups are located.

The discussion then turned toward the company's newly announced restricted stock unit awards for co-founders and co-Chief Executive Officers Daniel and William Roberts. Approved on June 30, the compensation package granted approximately 9.1 million restricted stock units to each executive, or roughly 18.2 million units combined, representing about 5% of the company's outstanding ordinary shares. The awards generated considerable discussion because of both their size and structure. Unlike many long-term executive compensation plans tied to clearly defined operational or financial performance targets, the awards vest in equal annual installments over four years based primarily on continued employment, with each vested tranche remaining subject to an additional two-year holding period. Neither executive is eligible to receive another equity grant until fiscal 2031. The structure raised questions regarding incentive alignment between management and shareholders despite having been unanimously approved by the company's independent directors following consultation with an independent compensation adviser.

While the awards were generally viewed as an unfavorable development from a governance perspective, the discussion stopped short of viewing them as sufficient reason on their own to abandon the investment thesis. Instead, it focused on the importance of continually evaluating management incentives as new information becomes available. Successful investing requires remaining objective while avoiding emotional attachment to previous assumptions. Even companies operating within attractive industries must still demonstrate disciplined capital allocation and shareholder-friendly governance over time.

The conversation also revisited IREN's strategic position within AI infrastructure. Despite the recent governance concerns, the company's long-term investment thesis remained largely intact. Its expanding power portfolio, growing data center footprint, and exposure to hyperscaler demand continue to represent significant competitive strengths if management executes successfully.

This naturally led into a discussion of sentiment and investment opportunity. Investor sentiment surrounding IREN had weakened considerably following both company-specific governance announcements and broader pressure across AI infrastructure providers after concerns emerged that Meta Platforms (META) could expand the resale of excess AI compute capacity. By early July, IREN shares had declined substantially from their late-May highs despite continued strength in the long-term outlook for AI infrastructure. The conversation noted that periods of negative sentiment can often warrant closer examination when long-term industry fundamentals remain largely intact, while recognizing that improving sentiment alone does not resolve governance concerns.

The discussion also compared IREN with several competing infrastructure providers, including Nebius Group (NBIS), CoreWeave (CRWV), Cipher Mining (CIFR), and other emerging AI compute businesses. Although each company possesses different strengths and weaknesses, the comparison reinforced that identifying an attractive industry is only one part of the investment process. Execution quality, capital allocation, management incentives, shareholder alignment, and corporate governance can ultimately become just as important as exposure to favorable

long-term trends.

The conversation closed by reinforcing an idea that extended well beyond one individual company. Attractive secular growth themes can create compelling opportunities, but they do not eliminate the need to continually evaluate management decisions and incentive structures as new information becomes available. Maintaining an objective investment process remains essential, particularly in rapidly evolving industries where substantial amounts of capital continue flowing into new infrastructure projects.

10. Lightning Round: Transportation, Memory, Prediction Markets, and Second-Order Opportunities

The meeting concluded with a wide-ranging lightning round that touched on several individual companies and investment themes while reinforcing many of the ideas that had surfaced throughout the evening.

One of the more interesting discussions centered on the proposed merger between Union Pacific (UNP) and Norfolk Southern (NSC). The conversation explored how combining two of

North America's largest railroad operators could reshape freight transportation, logistics, and competitive dynamics across the industry, extending well beyond the transaction itself. Discussion included the ongoing regulatory review by the Surface Transportation Board, with a key supplemental filing deadline approaching later in July, as well as the long-term implications of creating the first truly coast-to-coast rail network if the transaction ultimately receives approval.

Discussion also turned toward several second-order businesses that could benefit depending on the outcome, including logistics providers, intermodal transportation companies, and firms positioned to capitalize on greater freight efficiency. Hub Group (HUBG) was highlighted as one potential beneficiary given its long-standing intermodal relationships with both Union Pacific and Norfolk Southern, with improved network efficiency and expanded single-line service potentially creating new freight opportunities if the merger proceeds. The discussion reinforced a recurring theme from previous meetings that significant industry developments often create attractive opportunities several layers beneath the headline company.

Prediction markets also returned as a topic of discussion. Intercontinental Exchange (ICE) generated interest following its investment and data partnership with Polymarket as institutional demand for prediction market data continues expanding. The conversation also explored the growing role of prediction markets as information markets capable of aggregating expectations across a wide range of economic, political, and technological topics. While the regulatory framework continues evolving, companies providing the underlying exchange infrastructure, market data, and distribution networks could benefit alongside the platforms themselves if adoption continues expanding.

Another second-order AI infrastructure opportunity emerged through discussion of Clean Harbors (CLH). Beyond semiconductor manufacturers and cloud providers, attendees also examined how companies specializing in hazardous waste disposal, industrial recycling, and environmental services may benefit from continued data center expansion. Data centers require the ongoing replacement and disposal of batteries, cooling systems, electronic components, and other specialized industrial materials, creating recurring demand for businesses capable of managing those assets throughout their operational life cycles. It was noted that this remains a forward-looking investment thesis rather than a primary driver of the company's current financial performance, though AI infrastructure could eventually become another source of long-term demand alongside its existing environmental services business.

Oracle (ORCL) generated another detailed discussion following its significant share price decline despite continuing to report strong operating results. Concerns were raised regarding the company's substantial capital expenditures, negative free cash flow, and expanding AI infrastructure commitments, while the recent weakness was also viewed as

evidence that

market sentiment may have become overly pessimistic. The conversation centered on whether hyperscalers and cloud providers may experience temporary pressure on profitability as they continue investing aggressively in AI infrastructure before realizing the corresponding revenue growth. Oracle's record backlog was also noted as evidence supporting the view that long-term customer demand remains robust despite near-term financing requirements. Discussion surrounding Oracle also reinforced that temporary financial pressure does not necessarily invalidate a long-term investment thesis if underlying demand continues strengthening.

Micron Technology (MU) also returned as an important topic following another quarter of exceptionally strong financial results. The conversation explored whether the memory industry has fundamentally changed as AI infrastructure demand accelerates or whether traditional semiconductor cycles will eventually reassert themselves. Demand for high-bandwidth memory was discussed as a potential structural shift, with artificial intelligence requiring substantially greater quantities of advanced memory than earlier computing workloads and much of Micron's near-term HBM production already committed under customer agreements. Even so, semiconductor markets have historically experienced periods of overinvestment followed by slower demand growth, making disciplined expectations important when evaluating long-term prospects.

While no one can predict exactly how future semiconductor cycles will evolve, the discussion noted that current evidence continues pointing toward exceptionally strong demand supported by AI infrastructure investment, hyperscaler spending, and increasingly capable frontier models.

The lightning round reinforced one of the group's longest-running themes: some of the most compelling investment opportunities often emerge one or two layers beneath the most visible narratives. Infrastructure suppliers, logistics companies, industrial service providers, data distributors, specialized component manufacturers, and other enabling businesses frequently receive far less attention than headline technology companies despite occupying strategically important positions within rapidly expanding ecosystems.

Closing Thoughts

The July 6 meeting reinforced how quickly today's investment landscape continues evolving and how interconnected many of its most important themes have become. Although the discussion moved across artificial intelligence, healthcare, robotics, semiconductors, transportation, geopolitics, consumer technology, market structure, and individual companies, each topic ultimately connected back to many of the same long-term forces reshaping the global economy.

Artificial intelligence remained the central thread throughout the evening, but the conversation consistently extended well beyond AI models themselves. Compute infrastructure,

semiconductors, custom silicon, cloud platforms, cybersecurity, enterprise software, data ownership, robotics, biotechnology, transportation, logistics, and industrial automation were all discussed as interconnected components of a much larger technological transformation. Together, these developments continue reinforcing the view that AI is evolving into a broad industrial buildout requiring enormous investment across multiple layers of the economy.

The meeting also highlighted that identifying attractive long-term industries represents only part of successful investing. Throughout the discussion, execution repeatedly emerged as one of the defining factors separating long-term winners from the rest of the field. Capital allocation, management incentives, corporate governance, and disciplined leadership remained recurring considerations whether evaluating AI infrastructure companies, enterprise software, healthcare innovators, or emerging technology businesses. Long-term secular tailwinds create opportunities, but sustained value creation ultimately depends on execution.

Another theme that surfaced repeatedly was the importance of looking beyond the most visible headlines. Many of

the most compelling ideas discussed throughout the evening centered on the companies supplying critical infrastructure, enabling technologies, specialized components, industrial services, and supporting ecosystems beneath larger trends. From AI data centers, custom chips, and high-bandwidth memory to robotic components, rail infrastructure, and environmental services, opportunities often appeared one or two layers removed from the companies receiving the greatest public attention.

Healthcare also continued standing out as one of the group's highest-conviction long-term themes. Recent examples involving AI-assisted drug discovery, AlphaFold, computational biology, diagnostics, and personalized medicine illustrated both the tremendous potential and the scientific discipline still required to translate technological breakthroughs into meaningful patient outcomes. While progress is unlikely to be linear, advances across these fields continue suggesting that artificial intelligence may become an increasingly important catalyst for medical innovation over the coming decade.

Perhaps the most consistent lesson throughout the meeting was the importance of maintaining an interdisciplinary perspective. Understanding technology alone is no longer sufficient. Long- term investing increasingly requires connecting developments across engineering, economics, healthcare, geopolitics, infrastructure, regulation, finance, and human behavior while remaining willing to reassess assumptions as new information becomes available.

As always, the objective of The Investing Group remains unchanged: to challenge assumptions, exchange research, evaluate emerging trends from multiple perspectives, and deepen our understanding of an investment landscape that continues changing at an extraordinary pace.

We look forward to reconvening on Monday, July 20, 2026, at the Union League Club of Chicago as these themes continue evolving and new opportunities emerge across global markets. For those who were unable to attend or would like to revisit the discussion, a recording of the July 6 meeting is available on YouTube (https://www.youtube.com/watch?v=I6ObTIqBzXU).

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

Mario Sanchez (videography, media)

Diana Ascencio (media)

Victor Sanchez (founder, media)

Quinn Basta (recording, transcription, photography)

Sneha Shrivastav

Toby Chan

Moses Shriem

Seamus Cullinan

Eric Simpson (founder, moderation, summary, AI systems)

John Donners

Raza Siddiqui

Mohammed Haq (recording, transcription, attendance)

Daniel Valentino

Matthew Hilgart

Caleb Valdovinos

John Hockberger (founder, media)

Gabriel Zheng

Brian Jung

Jose Zuniga

Sandeep Kaza

Bill O'Brien

Suryansh Khatikar

Al Pakrosnis (AI systems)

Arthur Koslowski

Roberto Salgado