

DEAR BARON OPPORTUNITY FUND SHAREHOLDER:

PERFORMANCE

During the third quarter, Baron Opportunity Fund® (the Fund) rose 4.04% (Institutional Shares), outperforming the Russell 3000 Growth Index (the Benchmark), which gained 3.42%, but trailing the S&P 500 Index, which advanced 5.89%. For the first nine months of 2024, the Fund posted solid gains, increasing 25.31%, slightly outperforming the Benchmark, which rose 24.00%, and beating the S&P 500 Index, which appreciated 22.08%.

Table I.
Performance†

Annualized for periods ended September 30, 2024

	Baron Opportunity Fund Retail Shares ^{1,2}	Baron Opportunity Fund Institutional Shares ^{1,2,3}	Russell 3000 Growth Index ¹	S&P 500 Index ¹
Three Months ⁴	3.98%	4.04%	3.42%	5.89%
Nine Months ⁴	25.06%	25.31%	24.00%	22.08%
One Year	43.85%	44.24%	41.47%	36.35%
Three Years	3.22%	3.49%	11.31%	11.91%
Five Years	20.82%	21.15%	19.09%	15.98%
Ten Years	16.96%	17.27%	16.04%	13.38%
Fifteen Years	15.85%	16.16%	16.17%	14.15%
Since Inception (February 29, 2000)	9.67%	9.86%	7.48%	8.04%

REVIEW & OUTLOOK

U.S. equities closed higher for a fourth consecutive quarter, as generally upbeat economic reports supporting the soft-landing narrative and the Federal Reserve’s long-awaited dovish pivot spurred a meaningful rotation out of large-cap, *Magnificent Seven* stocks and into value/cyclicals/small caps. Beneath the surface, the market advance was anything but uneventful, with the S&P 500 Index falling by nearly 9% in late July/early



MICHAEL A. LIPPERT
PORTFOLIO MANAGER

Retail Shares: BIOPX
Institutional Shares: BIOIX
R6 Shares: BIOUX

August and by more than 4% in early September, before quickly recouping losses on both occasions. Shifting investor sentiment was reflected in the movement of the VIX Index, which briefly spiked above 65, its highest level since the onset of the COVID-19 pandemic and a mark hit only a few times in its 35-year history. The Magnificent Seven, which dominated market performance during the first half of the year, moved in opposite directions during the third quarter, as double-digit gains from **Tesla, Inc.**, **Meta Platforms, Inc.**, and **Apple Inc.** were mostly offset by declines from Alphabet, **Amazon.com, Inc.**, **Microsoft Corporation**, and **NVIDIA Corporation**, with the group posting a modest gain of 1.7% for the period. Sector leadership was centered in rate-sensitive sectors expected to benefit from lower interest rates, namely Utilities, Real Estate (REITs), Industrials (building products), Financials (regional banks and insurers), and Consumer Discretionary (homebuilders). Information Technology (IT) and

Performance listed in the above table is net of annual operating expenses. Annual expense ratio for the Retail Shares and Institutional Shares as of September 30, 2023 was 1.32% and 1.06%, respectively. The performance data quoted represents past performance. Past performance is no guarantee of future results. The investment return and principal value of an investment will fluctuate; an investor’s shares, when redeemed, may be worth more or less than their original cost. The Adviser may waive or reimburse certain Fund expenses pursuant to a contract expiring on August 29, 2035, unless renewed for another 11-year term and the Fund’s transfer agency expenses may be reduced by expense offsets from an unaffiliated transfer agent, without which performance would have been lower. Current performance may be lower or higher than the performance data quoted. For performance information current to the most recent month end, visit BaronCapitalGroup.com or call 1-800-99-BARON.

† The Fund’s historical performance was impacted by gains from IPOs and there is no guarantee that these results can be repeated or that the Fund’s level of participation in IPOs will be the same in the future.

1 The **Russell 3000® Growth Index** measures the performance of the broad growth segment of the U.S. equity universe. The **S&P 500 Index** measures the performance of 500 widely held large-cap U.S. companies. All rights in the FTSE Russell Index (the “Index”) vest in the relevant LSE Group company which owns the Index. Russell® is a trademark of the relevant LSE Group company and is used by any other LSE Group company under license. Neither LSE Group nor its licensors accept any liability for any errors or omissions in the indexes or data and no party may rely on any indexes or data contained in this communication. The Fund includes reinvestment of dividends, net of withholding taxes, while the Russell 3000® Growth Index and S&P 500 Index include reinvestment of dividends before taxes. Reinvestment of dividends positively impacts the performance results. The indexes are unmanaged. Index performance is not Fund performance. Investors cannot invest directly in an index.

2 The performance data in the table does not reflect the deduction of taxes that a shareholder would pay on Fund distributions or redemption of Fund shares.

3 Performance for the Institutional Shares prior to May 29, 2009 is based on the performance of the Retail Shares, which have a distribution fee. The Institutional Shares do not have a distribution fee. If the annual returns for the Institutional Shares prior to May 29, 2009 did not reflect this fee, the returns would be higher.

4 Not annualized.



Baron Opportunity Fund

Communication Services underperformed due to declines from certain members of the Magnificent Seven, specifically Microsoft, NVIDIA, and Alphabet.

Despite this backdrop for the quarter, Fund performance was solid, with a strong quarterly absolute gain of just over 4%, and modest relative outperformance versus the Benchmark. This was due to stock picking, as the Fund's selection and interaction effect added nearly 100 basis points of relative gains. Relative outperformance stemmed from our biotechnology investments in **argenx SE** (antibody immunology platform) and **Arcellx, Inc.** (cell therapy for multiple myeloma); our application software investments in **Guidewire Software, Inc.** (property and casualty (P&C) insurance platform) and **Samsara Inc.** (industrial digitization pioneer); electric vehicle (EV) and autonomous driving and robotics leader, **Tesla, Inc.**; and consumer digital plays **Duolingo, Inc.** (language learning innovator), **Spotify Technology S.A.** (streaming audio leader), and **Shopify Inc.** (e-commerce platform). Not owning Alphabet was another positive relative contributor in the quarter. On the flip side, notable relative detractors included our auto-semiconductor holdings, **indie Semiconductor, Inc.** and **Mobileye Global Inc.** (see below); cybersecurity leader **CrowdStrike Holdings, Inc.**, due to the well-publicized outage that occurred in July (see below); semiconductor lithography equipment leader, **ASML Holding N.V.**, which retreated on concerns regarding enhanced U.S. government trade restrictions; and the Fund's underweight in Apple, which rose this quarter in anticipation of the launch of Apple Intelligence to come in early 2025, Apple's initial foray into generative AI (Gen AI).

AI Infrastructure/Semiconductor Demand

Following up on the AI discussion in our last letter,^[1] our research over the past three months, including a recent Baron research trip up and down Silicon Valley, has confirmed that we remain in an AI arms race. This arms race can be described as a contest to develop the one foundational model^[2] to "rule them all," to borrow a line from Tolkien, and be the first to achieve artificial general intelligence, or AGI.^[3] The racecourse has been built on scaling laws, and the race obstacles are data, compute, and models. AI scaling laws describe the predictable trend of improvement in model performance with increases in three key factors: model size (number of parameters), data size (amount of training data), and compute instances (amount of computation used). Here is a simple breakdown:

- **More data:** As we train models on larger datasets, they generally perform better. More data provides a richer set of examples for the model to learn from, helping it capture patterns and nuances.
- **More Compute:** Increasing computational power (more GPUs or faster processors) allows for training larger models or training existing models more thoroughly. With more compute, models can process data faster and experiment with more parameters, leading to better performance.
- **Larger Models:** Increasing the size of the model, i.e., the weights or parameters, helps the model gain more capacity or brainpower to learn from larger datasets and capture more intricate patterns.
- **Better Answers:** Larger models provide better answers to user questions.

Companies that don't enter the AI race are left watching from the sidelines, missing out on the biggest technology and industrial transformation and economic jackpot in a generation. The entry fee of capital and operating investments is high. But the winners' podium will likely hold more than just the first, second, and third places, but probably less than ten. None of the large hyperscalers nor consumer internet players, not to mention a handful of sovereigns, can risk not having runners in this race.

To be sure, just the other day, Jensen Huang, NVIDIA's CEO, stated in a CNBC interview that "the demand for Blackwell [NVIDIA's next-generation platform] is insane. Everybody wants to have the most and everybody wants to be first." Larry Ellison, the CEO and founder of Oracle, publicly told the story of how he and Elon Musk were literally "begging Jensen for GPUs" over dinner with him. Elon's xAI recently opened a new data center, called Colossus, in Memphis, Tennessee, which was described by an xAI employee as "the largest data center on the planet." On our research trip, we met with the CEO of a major AI infrastructure player, who stated, in very blunt terms, we are "building...big a** networks for them," explaining that each of the "big 5" – Amazon, Google, Microsoft, Meta, and Oracle – are planning to build data centers at the scale of 100,000 AI accelerator chips (i.e., GPUs).

Here are a few quotes from hyperscaler CEOs on the AI infrastructure arms race:

- Amazon CEO Andy Jassy: "While we're investing a significant amount in the AI space and in infrastructure, we would like to have more capacity than we already have today. I mean, we have a lot of demand right now, and I think it's going to be a very, very large business for us."
- Meta CEO Mark Zuckerberg: "The downside of being behind is that you're out of position for like the most important technology for the next 10 to 15 years...I'd rather risk building capacity before it is needed rather than too late, given the long lead times for spinning up new infra projects...The amount of compute needed to train Llama 4 will likely be almost 10 times more than what we used to train Llama 3 and future models will continue to grow beyond that."
- Alphabet's Sundar Pichai: "When we go through a curve like this, the risk of underinvesting is dramatically greater than the risk of overinvesting... these are infrastructure which are widely useful for us."

While scaling laws have governed the mounting investments in model training infrastructure, they are now set to drive an inflection in inference compute requirements. Just a few weeks ago, OpenAI released a new series of "chain of thought" models, nicknamed Strawberry and released as the o1 series, designed to spend more time thinking and reasoning before they provide an inference response (measured by the number of tokens). These models reason through complex tasks and solve harder problems than previous models in science, coding, and math. This is a new type of inference architecture, where the models iterate on the answer after the initial prompt or query. In simple terms, the first answer becomes the new prompt and the

[1] As well as the Baron Insights piece, "AI Hype and the Death of Software," dated September 2024, which can be found on the Firm's website.

[2] A foundational model is a machine learning or deep learning model that is trained on broad data sets such that it can be applied across a wide range of use cases.

[3] AGI refers to the hypothetical intelligence of a machine that possesses the ability to understand or learn any intellectual task that a human being can.

model cycles through iterations as it “thinks” until the best answer is achieved. Inference innovations and evolutions like this require significantly more compute for inferencing, another growth driver for the AI accelerator manufacturers, like NVIDIA and **Broadcom Inc.** One of our team asked Jensen about this and he replied that “this is a massive inflection...all of a sudden, the model has to be inferred many more times for each token... depending on the quality of the answer.” A recent Morgan Stanley note stated: “inference is starting to solve much more complex problems which will require a richer mix of hardware...the notion that a more thoughtful, task-oriented inference would cause an exponential jump in inference complexity strikes us as an important new avenue for growth.”

The Data Layer

In our discussion of software in our last letter,^[4] we stated that “the enterprise software winners will have to be better at delivering AI services and features than build-your-own AI tools” and that “the winners will be the ones that have a well-established product development culture of innovation and iteration; differentiated proprietary, industry, and customer data; distribution advantages with large customer bases, successful go-to-market efforts, and key partners; well-designed workflows where AI improves the user interface, intelligent predictions/recommendations, and automation; and established always-on connectivity and feedback from their customers; among other things.” In this piece, we’d like to focus on the importance of the “the data layer” of the AI stack and in building differentiated AI applications.

While much attention today is often placed on the semiconductor or infrastructure layer of the AI stack, as we did above, our research has confirmed that it is the data layer – proprietary data and robust data processing – that forms the foundation for advanced AI applications. To build transformative, reliable AI applications, companies will need to build strategies around their data. Without a robust and scalable data infrastructure, AI models and algorithms lack essential raw material to generate meaningful insights, predictions, and automation. An advantageous data asset is not just about volume, it’s about how the data is managed and processed – it’s about quality (is the data accurate, complete, clean, and reliable?), structure (how is the data organized, stored, and manipulated?), accessibility (can models access it for inference and training?), orchestration (how is the data organized and combined with data from multiple sources?), and governance (is the data compliant and secure?). Proprietary datasets are becoming key differentiators. Organizations with unique access to high-quality, domain-specific data are positioning themselves to build smaller, use-case-specific models that deliver superior performance, enabling them to solve more complex problems. For example, Samsara is accelerating its market share capture in the telematics and commercial fleet safety sector because of the 10 trillion data points it collects from the more than 70 billion miles driven annually. The growing data asset – which is cleaned, structured, and stored in one connected operation cloud – fuels more AI-powered insights that help reduce accidents, emissions, and insurance premiums, driving more value than competitive solutions. **Datadog, Inc.** monitors the IT infrastructure and application activity for 28,700 customers in every industry, collecting trillions of data points per hour from a very diverse set of application types and IT environments. It then applies AI techniques like root cause analysis,

natural language querying, and autonomous investigations to help customers solve IT problems faster than competing solutions. Looking forward, we believe companies that have a firm grasp on data infrastructure, combined with unique access to proprietary datasets, will be essential drivers of value in the AI ecosystem.

At the end of the third quarter, the Fund had significant investments in the semiconductor and software industry verticals. As we define software,^[5] we had a 26.6% weight in the industry group, 800 basis points overweight our Benchmark. For semiconductors, we had a 21.2% weight in the industry group, 706 basis points overweight our Benchmark.

Below is a partial list of the secular megatrends we focus on:

- Cloud computing
- Software-as-a-service (SaaS)
- AI
- Mobile
- Semiconductors
- Digital media/entertainment
- Targeted digital advertising
- E-commerce
- Genetic medicine/genomics
- Minimally invasive surgical procedures
- Cybersecurity
- EVs/autonomous driving
- Electronic payments
- Robotics

We continue to run a high-conviction portfolio with an emphasis on the secular trends cited and listed. Among others, during the third quarter we initiated or added to the following positions:

- Semiconductors: **Broadcom Inc.** and **indie Semiconductor, Inc.**
- Software: **Atlassian Corporation Plc, Cloudflare, Inc., Cadence Design Systems, Inc., PAR Technology Corporation, GitLab Inc., and HubSpot, Inc.**
- Digital learning: **Duolingo, Inc.**
- Data centers: **GDS Holdings Limited**

Table II.
Top contributors to performance for the quarter ended September 30, 2024

	Contribution to Return (%)
Tesla, Inc.	1.08
Meta Platforms, Inc.	0.73
Guidewire Software, Inc.	0.64
argenx SE	0.60
Arcellx, Inc.	0.48

Tesla, Inc. designs, manufactures, and sells EVs, related software and components, and solar and energy storage products. Tesla shares contributed to performance during the quarter, reflecting increased investor confidence and optimism in Tesla’s AI initiatives, stabilization in the company’s industrial operations, including strong growth in its energy segment, and the anticipated launch of new vehicle models in the first half of 2025. After years of industry-wide investments in autonomous vehicles,

^[4] See also the Baron Insights piece mentioned in footnote 1.

^[5] We view both Dayforce, Inc. and PAR Technology Corporation as software companies, although GICS characterizes them in different industry verticals.

Baron Opportunity Fund

advancements in AI technology have accelerated the development of autonomous driving technology. Tesla deployed its AI-based Full Self Driving (FSD) solution last year and has demonstrated rapid improvements in driving performance. It has articulated a goal of achieving nearly a 20-fold improvement in miles driven between critical disengagements – soon exceeding 10,000 miles – over a two-month period this fall.

AI relies on vast amounts of high-quality data and computational power, and we believe Tesla possesses distinct assets that will serve as a strong foundation for its AI initiatives. Since 2016, every Tesla vehicle produced has been outfitted with cameras and essential hardware, resulting in millions of connected cars globally that gather data from billions of miles driven each year by the Tesla fleet. This rich and unique dataset is invaluable for FSD training. Tesla is also differentiating with its AI training compute factory. Tesla finished 2023 with close to 15,000 NVIDIA H100 chip equivalents in training computation power. By the second quarter of 2024, it doubled this capacity. In the third quarter, the company activated its advanced training data center in Texas, which should allow the company to harness up to 90,000 H100 equivalent compute power by the end of 2024 – six times the compute capacity it had at the beginning of the year and by far the world's largest autonomous driving training cluster. Unlike any other automotive company, Tesla is investing billions of profits generated by its automotive segment in its AI initiatives aiming to capture material share in lucrative markets of autonomy and robotics.

The rapid improvements in AI models and their ability to generalize machine learning efficiently, is also contributing to the developments of humanoids. Tesla is leveraging its multi-year, multi-billion-dollar investments in technologies such as inference and training compute, actuators, batteries, inverters, and scaled production capabilities – developed initially for its automotive business – to kickstart its Optimus humanoid-robot business. The company aims to deploy thousands of robots in its production facilities by the end of 2025, with plans to sell them to customers by 2026. We believe that the humanoid robot market opportunity could one day be even larger than the combined automotive and robotaxi opportunities.

On October 10, the company hosted its much-anticipated AI and robotaxi event, titled "We, Robot." This event provided additional insights into the company's initiatives and may further strengthen investor confidence in Tesla's market position.

Shares of **Meta Platforms, Inc.**, the world's largest social network, were up this quarter, due to impressive top-line growth of 22% year-over-year and solid forward guidance. Despite its large scale, Meta continues to outgrow the broader digital advertising industry, with better AI-driven content recommendations increasing engagement in products like Instagram Reels, and AI improving ad targeting and conversion rates. Our industry checks have continued to validate advertiser adoption and satisfaction, with improvements in Reels monetization, as well as strong adoption of Advantage+, Meta's AI-driven service to allocate advertiser budgets across its content surfaces.

Meta continues to innovate in Gen AI, with a leading AI research lab and the best open-source models to date. We believe CEO Mark Zuckerberg's open-source approach will encourage a broader developer ecosystem and standardization based on Meta, which will be beneficial for the company even if Meta doesn't directly monetize model usage over the near term. In a blog this summer, titled "Open Source AI Is the Path Forward," Zuckerberg laid out his case:

- Today, several tech companies are developing leading closed models. But open source is quickly closing the gap. Last year, Llama 2 was only

comparable to an older generation of models behind the frontier. This year, Llama 3 is competitive with the most advanced models and leading in some areas. Starting next year, we expect future Llama models to become the most advanced in the industry. But even before that, Llama is already leading on openness, modifiability, and cost efficiency.

- Today we're taking the next steps towards open-source AI becoming the industry standard. We're releasing Llama 3.1 405B, the first frontier-level open-source AI model...In addition to having significantly better cost/performance relative to closed models, the fact that the 405B model is open will make it the best choice for fine-tuning and distilling smaller models...Now you'll be able to take the most advanced Llama models, continue training them with your own data and then distill them down to a model of your optimal size – without us or anyone else seeing your data...Many organizations don't want to depend on models they cannot run and control themselves. They don't want closed model providers to be able to change their model, alter their terms of use, or even stop serving them entirely. They also don't want to get locked into a single cloud that has exclusive rights to a model... Many organizations handle sensitive data that they need to secure and can't send to closed models over cloud APIs. Other organizations simply don't trust the closed model providers with their data. Open source addresses these issues by enabling [organizations] to run the models wherever [they] want.
- Meta's business model is about building the best experiences and services for people. To do this, we must ensure that we always have access to the best technology, and that we're not locking into a competitor's closed ecosystem where they can restrict what we build. One of my formative experiences has been building our services constrained by what Apple will let us build on their platforms. Between the way they tax developers, the arbitrary rules they apply, and all the product innovations they block from shipping, it's clear that Meta and many other companies would be freed up to build much better services for people if we could build the best versions of our products and competitors were not able to constrain what we could build. On a philosophical level, this is a major reason why I believe so strongly in building open ecosystems in AI and AR/VR for the next generation of computing.

Beyond foundational research and development, we are beginning to see Meta's core apps incorporate Gen AI in the user experience, with Meta AI reaching over 500 million monthly active users today. Meta is also the mega-cap technology company most focused on profitability and efficient innovation. Long term, we believe Meta will utilize its leadership in mobile advertising, massive user base, innovative culture, leading Gen AI research and potential distribution, and technological scale to perform, with further monetization opportunities ahead.

Shares of P&C insurance software vendor **Guidewire Software, Inc.** contributed to performance for the quarter. After a multi-year period, we believe the company's cloud transition has now successfully "crossed the chasm," as I have sometimes referred to the shift for software vendors from on-premises systems to cloud architectures. For the company's fiscal fourth quarter, Guidewire reported an acceleration in annual recurring revenue (ARR) growth to 14%, cloud subscription revenue growth of 37%, subscription gross margins that improved by over 1,000 basis points, and cash flow from operations that beat Street estimates by over \$50 million (\$192 million versus \$142 million). We believe that cloud will be the sole

path forward for Guidewire and its customers, with ARR benefiting from new customer wins and migrations of the existing customer base to InsuranceSuite Cloud. We are also encouraged that Guidewire is running well ahead of its 2025 mid-term margin targets and delivering the inflection in free cash flow generation that we long projected. We believe that Guidewire will be the critical software vendor for the \$2 trillion global P&C insurance industry, capturing 30% to 50% of its \$15 billion to \$30 billion total addressable market and generating margins above 40%.

Table III.
Top detractors from performance for the quarter ended September 30, 2024

	Contribution to Return (%)
indie Semiconductor, Inc.	-0.70
Microsoft Corporation	-0.58
Mobileye Global Inc.	-0.40
CrowdStrike Holdings, Inc.	-0.39
NVIDIA Corporation	-0.36

Indie Semiconductor, Inc. is a fabless designer, developer, and marketer of automotive semiconductors for advanced driver assistance systems (ADAS) and connected car, user experience, and electrification applications. Indie's stock fell during the quarter as it guided 2024 revenue growth below Street expectations, as auto production is expected to be incrementally worse, and excess inventory in the automotive supply chains of its customers has delayed indie's new chips from ramping into high-volume programs. Despite these near-term headwinds, indie is outperforming peers who are seeing significant year-over-year sales declines. It has not suffered a program cancellation for any intact car programs, and it continues to win new sockets in future car platforms, positioning the company for strong growth over the medium and long term, supported by its \$6.3 billion design win backlog, of which \$4.6 billion is in ADAS applications. Indie has several large-volume programs set to ramp beginning in early 2025, including a marquee radar-related win, the biggest program in the company's history, which we believe will drive a return to outsized growth in 2025 (indie doubled revenue each year from 2021 through 2023). We believe indie can continue to significantly outpace the broader industry and will approach \$1 billion in revenue by the end of this decade, all supported by its contracted visibility, and its stock will re-rate as rapid growth returns.

Microsoft Corporation is the world's largest software and cloud computing company. Microsoft was traditionally known for its Windows and Office products, but over the last five years it has built a \$147 billion run-rate cloud business, including its Azure cloud infrastructure service and its Office 365 and Dynamics 365 cloud-delivered applications. Shares gave back some gains from strong performance over the first half of this year. For the fourth quarter of fiscal year 2024, Microsoft reported a strong quarter with total revenue growing 16%, in line with the Street; Microsoft Cloud up 22%; Azure up 30%; 43% operating income margins; and 36% free cash flow margins. Core Azure growth came in one point shy of expectations, however, due to a soft European market and continued constraints on AI compute capacity. In the same vein, while Microsoft reiterated its fiscal 2025 targets of double-digit top-line and operating income growth, quarterly guidance called for Azure growth to slow a bit before accelerating in the back half of the fiscal year, as capital expenditures increase, yielding an expansion of AI compute capacity. We believe this investment is a leading indicator for growth, with more than half of the spend related to durable land and data center build outs, which should monetize over the next 15-plus years. We remain confident that

Microsoft is one of the best-positioned companies across the overlapping software, cloud computing, and AI landscapes, and we remain investors.

Shares of **Mobileye Global Inc.**, a provider of ADAS and autonomous driving technologies for the automotive industry, detracted from performance. Excess inventory among key customers proved to be a headwind to growth earlier in the year. Mobileye also experienced a significant decline in market share in China as local original equipment manufacturers shifted to domestic suppliers and in-house technology. Although this was an embedded risk, it materialized faster than expected and included market share losses at key customers. While we maintain our belief in the size and strategic importance of the autonomous vehicle market (see Tesla above), we decided to exit our Mobileye position during the quarter and book a tax loss.

CrowdStrike Holdings is a leading software company focused on the cyber security market. The company's products help organizations prevent malicious actors from gaining access to, attacking, and shutting down a company's hardware systems, applications, and databases. Following a well-publicized and disruptive global outage in July – tied to a software glitch in connection with CrowdStrike's configuration update for Microsoft operating systems – shares fell sharply. Since this time, we have met with management three times, including three days at the company's recent user conference, speaking to over 50 customers, key partners, and new prospects. It is clear the company will see increased scrutiny on deals, including the implementation of a new incentive program offering customers short-term discounts, free products, and billing flexibility, which will negatively impact bookings, revenue, and cash flow over the near term. However, we came away from our diligence believing CrowdStrike can retain its premium brand and likely emerges stronger over the next year and beyond.

PORTFOLIO STRUCTURE

Table IV.
Top 10 holdings as of September 30, 2024

	Quarter End Market Cap (\$ billions)	Quarter End Investment Value (\$ millions)	Percent of Net Assets (%)
Microsoft Corporation	2017	165.4	12.2
NVIDIA Corporation	2018	156.4	11.5
Amazon.com, Inc.	2014	90.5	6.7
Meta Platforms, Inc.	2022	70.4	5.2
Apple Inc.	2024	63.1	4.7
Broadcom Inc.	2024	62.1	4.6
Tesla, Inc.	2014	61.5	4.5
Space Exploration Technologies Corp.	2020	36.1	2.7
Gartner, Inc.	2007	35.3	2.6
argenx SE	2017	35.1	2.6

We invest in secular growth and innovative businesses across all market capitalizations, with the bulk of the portfolio landing in the large-cap zone. Morningstar categorizes the Fund as U.S. Large Growth. As of the end of the third quarter, the largest market cap holding in the Fund was \$3.5 trillion and the smallest was \$500 million. The median market cap of the Fund was \$41.5 billion, and the weighted average market cap was \$1.3 trillion.

To end the quarter, the Fund had \$1.4 billion of assets under management. We had investments in 43 unique companies. The Fund's top 10 positions accounted for 57.2% of net assets.

Baron Opportunity Fund

RECENT ACTIVITY

Table V.
Top net purchases for the quarter ended September 30, 2024

	Quarter End Market Cap (\$ billions)	Net Amount Purchased (\$ millions)
Broadcom Inc.	805.7	14.2
Atlassian Corporation Plc	41.5	12.2
Duolingo, Inc.	12.3	11.1
GDS Holdings Limited	4.0	10.0
PAR Technology Corporation	1.9	7.9

We continued to build our position in **Broadcom Inc.**, a global technology leader that designs, develops, and supplies a broad range of semiconductor and infrastructure software solutions.^[1] As AI continues to proliferate, we believe hyperscalers – such as Meta, Microsoft Azure, Google Cloud Compute, and Amazon Web Services, to name just a few – will increasingly deploy custom accelerator chips for their AI workloads as they can be more cost-effective and energy-efficient than using NVIDIA's general-purpose GPUs. Broadcom has a leading position partnering with hyperscalers to develop these custom chips, with its AI customer accelerator business up 3.5-times year-over-year in its most recently reported quarter, and a goal of at least \$8 billion in custom accelerator revenues for this fiscal year. Additionally, VMware continues to perform better than expected as Broadcom is implementing its product simplification and subscription revenue model strategy. Further, its non-AI related semiconductor business, which tends to be more cyclical, is in the early stages of a recovery. Combined, all these factors will drive strong revenue and earnings growth over the next several years.

Atlassian Corporation Plc is leading software company focused on the collaboration and productivity markets. The company initially focused on serving software engineers (over 20 million worldwide), but its newer products, features, and use cases address a much larger set of users, including business teams (IT, marketing, project management, and human resources) engaged in product development (over 100 million worldwide) and the much broader group of knowledge workers (over 1 billion globally). The company recently completed a complex three-year migration of its on-premises customers to the cloud. While the complexity of this migration process weighed on shares (down about 50% since 2022), we believe the company is now better positioned to cross-sell new products and expand usage in the existing customer base. We recently met with several senior executives at the company's offices across multiple cities and attended Atlassian's annual user conference. At the latter we spoke with dozens of customers and partners and observed a notable improvement in the company's innovation cadence and customer receptivity to these new innovations versus last year's conference, including: (i) enhancements to Jira Service Management that simplifies incident management (identifying and remediating customer IT service issues); (ii) Jira Product Discover, that helps teams organize and prioritize projects that align with corporate goals; (iii) new AI capabilities enabling the use of natural language queries to quickly access reliable and relevant information about your organization and tasks; and (iv) Rovo, a search tool that helps teams find information across different applications, learn and iterate on knowledge work, and act faster with virtual agents. While the macro backdrop and related IT budget environment remains fluid, we have greater confidence that growth is bottoming, the company is delivering more products it can sell, and that growth is poised to reaccelerate driving free cash flow margins to expand.

Given our increased conviction in their guidance and mid-term targets, and the dislocated valuation (revenue multiple is trending near a 10-year low), we initiated a position in the stock during the quarter.

Duolingo, Inc. is the world's leading language learning app with over 100 million monthly active users, known for its effective gamification and high engagement. After monitoring the company over the past year and a half, we developed conviction to buy the stock for a few reasons. The company has maintained premium levels of user growth (daily average user growth of over 50%) and revenue growth (40%-plus), executed well against their product roadmap, gained early traction with new functionality, and maintained impressive 40%-plus incremental margins. We view the founder-led management team as best in class, technically capable (CEO and CTO both earned PhDs in machine learning from Carnegie Mellon University), and product focused. We initiated a position in the quarter as the share price fell to what we deemed attractive levels from a long-term valuation perspective, coupled with material catalysts on the horizon, particularly the broader launch of AI functionality (branded "Max") that enables users to have real-time conversations with AI based characters and a substantial improvement of the company's Advanced English offering. We believe that these two initiatives take Duolingo from more of a hobby app to a company that can address the broader market of 1.8 billion people learning English today. As these products roll-out in the coming quarters, we believe their adoption should drive the realization of higher pricing, faster revenue growth, lower churn, and continued margin improvement. We also believe there is additional optionality in newer products such as math and music, which are earlier in their product evolution.

In the most recent quarter, we re-initiated a position in **GDS Holdings Limited**. GDS is a pan-Asia data center operator with 1.5 gigawatts of power capacity across approximately 100 data centers in and around "tier one" cities in mainland China (GDS Holdings or GDSH), as well as 1.0 gigawatts of power capacity in its rapidly growing Asia ex-China business (GDS International or GDSI). GDS develops and leases data center space (on a power reservation basis) to the top global technology companies such as Alibaba, Tencent, ByteDance, Microsoft, Google, and Oracle under long-term, contracted arrangements. We recently met with CEO/founder William Huang and CFO Daniel Newman in our offices and believe the best days for the company are ahead of it due to durable secular tailwinds in cloud adoption (early innings in Asia, which are lagging the U.S. and rest of the world), continued growth in data, increasing demand from AI applications, and global constraints on power availability yielding sustained pricing power in light of low available supply amid continued strong demand. On a sum-of-the-parts basis, we see a path for the business to be worth \$45 to \$55 per share in two to three years versus approximately \$20 at recent market price. For GDSI, based largely on contracted customer commitments, we see cash flow growing from less than \$50 million today to over \$500 million over the next three years, with the opportunity to ramp towards \$1 billion a few years after that. We value GDSI at \$15 per share over the near term and \$25 per share over the next four to five years. Regarding GDSH, we believe the mainland China data center business is at the doorstep of a growth inflection and see cash flow growing from about \$700 million today to \$1 billion over the next three years based on lease-up of its available power capacity. We value GDSH at \$30 to \$40 per share over the near term and remain encouraged that there will be several catalysts to further enhance value (including a structure to place certain stabilized data center assets into a listed REIT vehicle).

^[1] Please see our first quarter 2024 letter for our initial write-up of Broadcom.

During the quarter we purchased shares of **PAR Technology Corporation**, a leading software, hardware, and service provider to the foodservice industry. The restaurant industry has historically underinvested in technology, and PAR is building an all-in-one software platform for enterprise restaurants to run the most critical portions of their technology stacks. Within the past year, PAR completed the final steps of its multi-year organic and inorganic product journey championed by CEO Savneet Singh, selling off its non-core legacy government services business and acquiring online ordering, international, and convenience-store related products to complement its existing point-of-sale, loyalty, back office, and payments processing products. We believe PAR will grow its software revenues over 20% per year for the next several years, driven by wins with new restaurant brands and cross-sell of its product portfolio to existing customers, and faces benign competition in the industry, with legacy providers offering less-modern products and more modern competitors having limited success and track records with enterprise restaurant chains (over 50 locations). We project the company will report its first quarter of adjusted EBITDA profitability in the third quarter of 2024 and expect it to deliver strong operating leverage from revenue growth going forward. As PAR scales into vertical software-like margins in the coming years, driving a ramp in bottom line profitability and cash flow, we believe the stock should appreciate meaningfully.

Table VI.
Top net sales for the quarter ended September 30, 2024

	Quarter End Market Cap or Market Cap When Sold (\$ billions)	Net Amount Sold (\$ millions)
Microsoft Corporation	3,198.4	18.6
NVIDIA Corporation	2,978.9	17.5
Advanced Micro Devices, Inc.	265.6	15.6
Marvell Technology, Inc.	55.4	11.4
Lam Research Corporation	95.3	7.7

Given the stellar returns of their stocks over the last couple of years, particularly **NVIDIA Corporation**, and the weights they grew to in the portfolio, we trimmed NVIDIA and **Microsoft Corporation** during the period. As we articulated above, our views regarding AI and the leadership of these two companies have not changed. On an absolute basis, NVIDIA and Microsoft remain the top two positions in the portfolio – as of this writing NVIDIA is our largest position and Microsoft is second – and both remain material overweights versus the Benchmark.

In the semiconductor space, we also trimmed **Advanced Micro Devices, Inc.** and exited **Marvell Technology, Inc.** and **Lam Research Corporation**. We shifted some of the capital from these sales and the NVIDIA trim to increase our weight in Broadcom, with the aim of owning meaningful positions in the leading general-purpose AI accelerator manufacturer, NVIDIA, and the leading customer accelerator partner to the hyperscalers, Broadcom. We also made slight adds to **Taiwan Semiconductor Manufacturing Company Limited** and **ASML Holding N.V.**, when they traded down amidst the volatility of the period.

In the software space, we shifted capital from our Microsoft trim into several non-mega names, including **Atlassian Corporation Plc**, **PAR Technology Corporation**, **Cloudflare, Inc.**, **Cadence Design Systems, Inc.**, **GitLab Inc.**, **HubSpot, Inc.**, and **Samsara Inc.**, as listed above and described herein.

I remain confident in and committed to the strategy of the Fund: durable growth based on powerful, long-term, innovation-driven secular growth trends. We continue to believe that non-cyclical, durable, and resilient growth should be part of investors' portfolios and that our strategy will deliver solid long-term returns for our shareholders.

Sincerely,

Michael A. Lippert
Portfolio Manager

Investors should consider the investment objectives, risks, and charges and expenses of the investment carefully before investing. The prospectus and summary prospectus contain this and other information about the Funds. You may obtain them from the Funds' distributor, Baron Capital, Inc., by calling 1-800-99-BARON or visiting BaronCapitalGroup.com. Please read them carefully before investing.

Risks: Securities issued by small and medium sized companies may be thinly traded and may be more difficult to sell during market downturns. Companies propelled by innovation, including technology advances and new business models, may present the risk of rapid change and product obsolescence, and their success may be difficult to predict for the long term. Even though the Fund is diversified, it may establish significant positions where the Adviser has the greatest conviction. This could increase volatility of the Fund's returns.

The Fund may not achieve its objectives. Portfolio holdings are subject to change. Current and future portfolio holdings are subject to risk.

The discussions of the companies herein are not intended as advice to any person regarding the advisability of investing in any particular security. The views expressed in this report reflect those of the respective portfolio managers only through the end of the period stated in this report. The portfolio manager's views are not intended as recommendations or investment advice to any person reading this report and are subject to change at any time based on market and other conditions and Baron has no obligation to update them.

This report does not constitute an offer to sell or a solicitation of any offer to buy securities of Baron Opportunity Fund by anyone in any jurisdiction where it would be unlawful under the laws of that jurisdiction to make such offer or solicitation.

Free Cash Flow Yield is a financial solvency ratio that compares the free cash flow per share a company is expected to earn against its market value per share.

BAMCO, Inc. is an investment adviser registered with the U.S. Securities and Exchange Commission (SEC). Baron Capital, Inc. is a broker-dealer registered with the SEC and member of the Financial Industry Regulatory Authority, Inc. (FINRA).