

# Warden Capital War Cap

June 2026

---

War Cap Q2 2026 update - by Hawkins Entrekin

Warden Capital Subscribe Sign in War Cap Q2 2026 update Hawkins Entrekin Jul 21, 2026 9 Share I know I said I'd likely stop writing these big letters, but I just couldn't help myself this quarter. So much is happening, and I wanted to jot my thoughts down as much for myself as anything else. And also to crow about a big win we had. Quickly on the returns. We had a big quarter, with YTD returns up to 15.91% vs 9.3% for the S&P (my main target benchmark) and 16.96% for the USRT (again not really relevant comp these days but including for historical continuity). Since inception we are up to 236.3%, vs 181.29% for the S&P and 89.84% for the USRT. And we did it with no semiconductor exposure (and thus we have avoided the momentum bloodbath since Q2 ended), and a fairly good sized short book to boot. I will lead with the star of the show this quarter, which was Uniqure, a stock I mentioned first in late 2025 I believe. As a reminder this is a biotech that has developed a promising treatment for Huntington's disease, a relatively rare, 100% fatal neurodegenerative disease. It published results in September, the stock skyrocketed to ~\$70, then in November the FDA changed tack on what seems to have been previously agreed to guidance accepting an external control group for an accelerated approval and the stock crashed down into the \$20s (where I first bought). Then in March, the FDA fully refused to even consider Qure's accelerated application, saying Qure needed to run a full phase 3, and the stock crashed down to as low as \$9/share, below cash balances. I doubled down near these lows, believing that if the drug worked (and no one really seems to think it doesn't), that worst case you waited out the phase 3 (which Qure had cash to complete), and best case the FDA could change its mind either via personnel change, when additional data came out (they have year 4 data slated for September), or Qure could monetize via approvals abroad (which they are pursuing). In an absolute whirlwind of a few months, the former FDA leadership was canned after what they did to Qure and several other promising disease treatments, and the new FDA interim leaders have allowed Qure to proceed with a filing for accelerated approval, with the stock now trading around \$40s/share. I will be honest and admit I didn't think it would happen this quickly, but what a ride! Sadly for me I subscribe to some levels of risk/concentration management and sold some of my position after the initial pop from the lows around \$10. But we still have a large position here. I was a bit worried about competition but after Roche pulled its latest trial I am increasingly convinced that drugs which lower wild type Huntingtins protein across the whole brain or body, which essentially all the major competitors do, are problematic 1. This then makes Qure's approach of selective targeting of the striatum all the more strategic, and I think there is still significant upside from here if Qure can get full approval. All that said, Qure is not our biggest position right now, that honor goes to Gitlab, which has done quite well since we bought earlier this year. Gitlab crashed significantly from fall of 2025 to today along with basically all other software businesses on a fear that LLMs would replace them all. The basic thesis here is that as AI dramatically increases the velocity of code production, code control, review and approval (which Git sits at the center of), becomes even more critical than it ever was. The code repos also offer a very natural place to serve up coding agents - most easily visible in Microsoft pushing its Copilot agent through competitor Github. The growth

is already visible in Gitlab's usage, with code pushes up nearly 50% YoY. Gitlab doesn't get paid on usage (yet) <sup>2</sup>, but it's a really good sign when people are using your platform a lot more generally speaking. I believe Gitlab is well positioned to monetize this growth itself, or that it would present an incredibly strategic acquisition opportunity for one of the major AI labs or companies. The market cap is a mere ~\$5.5B (and \$4.2 ev after cash), a drop in the bucket against the trillion dollar plus valuations in play elsewhere. Gitlab serves a whopping 50% of the F100, and would be a really valuable distribution tool. Microsoft already bought major competitor Git Hub, so there is precedent here as well. One of the firm's trying to produce an AI model that is struggling with adoption I think would find a company like Gitlab very tempting, such as Meta or SpaceX. Stepping back, the investment book feels about as good as it has since covid times, there are many interesting firms trading at attractive valuations, and I feel fairly good about returns continuing to be strong from here despite significant economic uncertainty.

Macro Musings And speaking of uncertainty, I am growing more concerned about our two big macro risks, the reclosure of the Strait of Hormuz, and the AI boom. In order to keep this letter brief I won't spend too much time on Hormuz - but if a deal is not reached soon and the strait remains closed it looks like oil inventories could hit crisis levels in 1-3 months. I suspect/hope something will get done as Trump seems to be aware of this, although unfortunately it may involve paying off Iran or conceding Iranian control of the strait, an effective strategic defeat for the US. The AI boom is a bit less binary, and therefore there is more to discuss. I believe the AI space is getting overheated, and as the scale has continued to ramp up I am worried that the impact of a bust is approaching a level that may cause a downturn. The numbers involved here are absolutely staggering, and simply put, there does not appear to be any way that a reasonable return on the investment is earned barring the achievement of AGI within the next few years. I have written about this several times before, but to repeat, at today's level of investment (~\$700 billion/year), there needs to be ~\$1 trillion in end AI annual revenue (from the model cos / providers) by 2029, and that figure would ramp to nearly \$4 trillion by 2036 assuming capex stays constant. Here is a link to a quick tweet summary of the math if you are interested <sup>3</sup>. What is crazy to me is that many estimates have AI capex increasing significantly in 2027, and some even have it approaching \$2 trillion by 2030! Obviously as the capex grows the required revenues also grow. The tech industry is simply not prepared for their formerly asset light, high margin golden children to turn into low margin commodity capital intensive industrial companies. Almost no one in the industry has the heuristics or experience for this. I cannot tell you how many times I see supposedly knowledgeable Silicon Valley people talking about inference gross margins with no reference to capital investment or D&A - this is like talking about the marginal cost of running an oil well but ignoring the cost to drill it! Or it's like asking - 'is building apartments profitable', and someone responding 'yes, the rental operating margins are very high'. The answer is absolutely and completely irrelevant on its own without reference to a return on capital invested & the cost to build the apartments. Amusingly, the one group of tech people who viscerally understand this, the management team's of the memory companies (a business famous for its boom/bust cycles), have opted into a strategy that seems to capitalize on this knowledge by riding the wave and maximizing their own profits, damage to the overall ecosystem be damned. This is, I suspect, the optimal strategy here, although if they push too hard they may incentivize the hyperscalers to enter into the memory game as well. I rather liked my friend Andrew Walker's observation relating to Micron and its total capital investment vs Google. Micron has only spent ~\$125 billion in capex in its entire history (or at least since 1992, if ChatGPT is accurate). Or alternatively, its

book value is \$73 billion - this is perhaps a more realistic assessment of the current 'replacement cost', if you will. Google is going to spend \$185 billion this year . How hard would it be for Google to spin up a memory division? Or Amazon. Perhaps more likely is Apple or someone stands up a new player in China to secure and permanently commodify the memory supply chain so critical to their products. Obviously such a thing would not be easy, but it would not be impossible, and if memory profits were to stay elevated it may happen. It also incentivizes chip designs and model utilization that would be less reliant upon memory as well - human ingenuity is a powerful thing, and I do not like betting against it long term. More likely though I would guess the whole thing crashes down, and the memory executives know this is likely to occur before too many new entrants come into the space. They are also probably aware of China's rapid ramp of memory capacity via the growth of CXMT (now almost as large as Micron in terms of capacity!) & YMTC, and the likely terrible price wars and margin pressure to come in several years from China's entry into this market. China has successfully commoditized every other electronics sector they have entered, and have moved up the value chain into autos. I don't see why memory should be any different. How Much Longer The big question is, how much longer can this AI capex boom last? It is hard to say, but I think the canary in the coal mine will be when model co revenues slow down. Given their epic growth thus far this year, it wouldn't surprise me if this happens soon, but honestly I have no idea. It is notable that in the last few days we have seen the release of GPT 5.6, a new Meta Spark model, and Grok 4.5, and all 3 release put significant emphasis on their lower costs and price points. Not to mention a new open weights model out of China that is competitive with Fable and GPT 5.6 trained at what was supposedly a fraction of the cost. These releases with an emphasis on cost mark a pretty big change from earlier releases which were really only focused on performance, and is potentially a dangerous sign of things to come for model providers. It could even indicate that the model cos are seeing corporate users become price sensitive in real time, which would presage revenue growth beginning to slow down. Broadly speaking my read on the narrative is that frontier AI users are now actively working to rein in spend. However this may be counterbalanced or even overwhelmed by continuing diffusion throughout the broader corporate ecosystem. At some point though that diffusion completes and reality begins to set in. A little rough AI generated summary of this concept - since diffusion is rapidly completing, the key unknown is where token intensity shakes out. It will peak or at least level off at some point, the question is simply when. The adoption of AI in corporate America has been incredible - I don't think any technology has ever diffused quite so quickly. McKinsey and BCG have some surveys on this subject - McKinsey estimates 88% of employees surveyed said their company was using AI, back in 2025. BCG had a survey with 72% of respondents (all corporate employees), saying they were regular AI users, over a year ago in June of 2025. 4 We are likely approaching effective full employee diffusion, with recent revenue growth driven by more complex workflows and agent usage (aka each user spending more). This rapid diffusion has been great for AI revenues, but also means that the end of the adoption cycle is going to come much faster than previous technologies, and that the crash could be all the more violent because of how compressed the cycle is. We will have to wait for the Q2 financial leaks to see the latest. But the model cos will need to approach \$300-400B of revenues in a year, and ~\$1 trillion by 2030, to justify all the AI capex investment. They are growing quite quickly today, but would need to accelerate further to hit these lofty targets. If revenue growth slows and/or the capex spend slows down, it could well take down the entire US economy given the size of the AI sector in the public markets 5 , and how important public markets are as a percentage

of household net worth today. From Bloomberg and Barclays. Stocks are at an all time high as a percentage of HH networth. The stock market is also near a record on the cyclically adjusted PE measure, second only to the dotcom bubble. The CAPE P/E uses inflation adjusted earnings over the last 10 years, in this way it can capture potential earnings bubbles and price bubbles. Today's absolute P/E is not as high relative to prior figures, but the massive recent spike in earnings means the CAPE is much higher. And setting aside the stock market, at this point a fairly large chunk of GDP is directly & indirectly tied to data center and power investments at this point. Even the automakers are re-orienting their battery operations around datacenters! All that said, it is hard to predict with much certainty what is going to happen with the overall economy. While I am relatively confident the AI boom turns into a bust at some point, I do not know when, and I am not 100% sure it will cause a US recession, and even if it does it is very hard to know how bad that might be. Ideally it would be a shallow, concentrated one ala 2000. There is some risk though of a broader, deeper crisis though, as one of the biggest differences this cycle is the asset & debt heavy nature of much of the investment, at least in the neocloud and datacenter space. Debt distress is much more likely to cause a longer, more painful crisis than equity. However much of this issuance appears to be backstopped by the still fairly good credit hyperscalers, so I am hopeful that a downturn wouldn't lead to a full blown financial crisis (but I wouldn't rule it out, especially the longer the boom goes!).

Fin As always, thanks for reading. Despite the macro uncertainty I feel good about our portfolio, our stocks are quite cheap & may even benefit from a rotation out of momentum into value. And we have a significant short book that would hopefully offset a broader economic downturn. Stay safe out there. Thanks for reading Warden Capital! Subscribe for free to receive new posts and support my work.

Subscribe 1 To get a bit further into the weeds but not quite all the way, regular huntingtin protein, or wtHTT, plays a critical role in other parts of the brain in particular, and there is good biological reason to believe that suppressing it across the brain could cause long term problems. My view here is supported by competitive trials which show a trend of higher dosage actually causing worse results, and more advanced patients doing worse, but promising results early in time and patient disease stage. That said this opinion is loosely held, there are still many unknown variables in all of these trials. But at the end of the day Qure's AMT-130 is the only treatment which seems to show any real long term efficacy in slowing disease progression.

2 Technically this isn't quite true, there are some ancillary services Gitlab charges per task on, but the vast majority of revenues are seat based for Gitlab. Gitlab's new coding harness (which allows you to use any underlying model), does have usage based cost but its still a very small share of revenues.

3 The assumptions here are actually fairly conservative (to the benefit of semis bulls), when I dug a little deeper the math actually was much worse because the share of datacenter spend going to chips has continued to increase, which means the useful life estimate I had in there was overall too high. This means required revenues would actually be significantly larger. The main offset to this is that Google is a good sized share of the AI spend, and one could argue that their spending need not have an economic return. As Google faces an existential risk to its search business, and so their investment may simply be something they must do to preserve that business, even if it doesn't actually generate much if any incremental new profits or revenues. The useful life estimates would add ~50% to the required revenues, while Google is about ~22-25% of total capex in the space, so crudely my figures here are a net ~25% conservative. So one year of current AI capex needs ~10 years of \$400b / year in revenues to pay it back.

4 The same survey recently came out for 2026 but didn't give a comparable all employee figure, instead breaking it down by position. But its like

roughly the figure is now at 83-85% - there is not much growth left to be had here. 5 I have seen estimates as high as close to 50% of the stock market. The AI trade has drawn in pretty much the entire utility sector and a large swath of even the industrials space as companies reorient their offerings around the data center boom. 9 Share Discussion about this post Comments Restacks Top Latest Discussions No posts Ready for more? Subscribe © 2026 Hawkins Entrekin · Privacy Terms Collection notice Start your Substack Get the app Substack is the home for great culture

This site requires JavaScript to run correctly. Please turn on JavaScript or unblock scripts