

STUMBLING TOWARDS EQUILIBRIUM

After a 15-year run, “lower for longer” took its final breath in late 2021 or early 2022, as it became clear to the world’s central bankers that post-pandemic inflationary pressures would not fade so easily after all. The exact date of its demise is disputable, but a plausible candidate is 29 November, when Jerome Powell finally conceded to Congress that “it’s probably a good time to retire [the] word [transitory].” Shortly thereafter, the Fed began to aggressively raise its policy rate, as did virtually all central banks across advanced economies. Ever since, the US and Europe have seen short-term interest rates inch higher and higher. This phenomenon is old news by now, and we have discussed its implications in several of [our letters](#) in recent years.

However, we have witnessed something new during the last few months: similar to what happened with short-term interest rates about a year ago, long-term rates recently spiked to levels not seen in at least a decade. Indeed, after expanding at a relatively gentle pace throughout most of the current tightening cycle, the far end of the yield curve shifted out rapidly over the last two quarters. Back in April, 30-year treasuries traded at yields of about 3.6%. By late October, 30-year treasury yields had shot up to over 5%, their highest level in about 15 years. After retreating somewhat in recent weeks, these yields were resting at about 4.2% in early December (at the time of writing). Following a similar path, yields on 30-year gilts rose from 3.8% in April to over 5% by October, and now stand at about 4.3%.

On the one hand, this shift is consistent with views we have articulated previously. In [our May 2021 letter](#), we argued that the ultra-low long-term interest rates of the last decade were contingent on specific structural forces that emerged several decades ago and were unlikely to last much longer. Most importantly, these forces included (1) the integration of China (and ex-Soviet states) into the world economy via globalisation, beginning in the 1990s, and (2) sharp growth in the global working-age population, both in absolute terms and relative to the size of the world’s “dependent” (i.e., child and elderly) population. On the margin, these two trends expanded the supply of global savings above a level that could be productively absorbed by investment. This gave rise to a persistent savings glut that dragged interest rates down to their zero lower bound. However, the low-hanging fruit of globalisation have now been picked—for example, a country like China could only industrialise and integrate into the world economy once—suggesting slower growth for global trade in the future. At the same time, longstanding demographic trends have now gone into reverse, with China’s working age population shrinking and global old-age populations growing rapidly. Accordingly, we expected long-term interest rates to rise over the next decade, as the dynamics that had depressed them began to unwind.

Although we think the ideas described above remain broadly right, we have started to question whether the market may have now overcorrected. With respect to the fundamental dynamics that drive long-term interest rates, the world of today is actually quite similar to the world of five years ago. It is true that the world’s population is ageing, China’s working-age population is shrinking, global trade is decelerating, and Western governments are increasingly concerned with building up domestic industrial capacity. Each of these trends should put upward pressure on long-term interest rates, too. Yet these observations were all accurate five years ago, and they were not substantially less visible back then either. Investors seem to have updated their interest-rate expectations to a degree that seems potentially unjustified by the comparatively limited amount of new information they have received.

One explanation is that markets have conflated the determinants of short-term interest rates—above all, near-term inflation expectations—with the determinants of long-term interest rates. The last several years were indeed filled with surprising developments pertaining to near-term inflation

expectations. However, almost all of these developments stemmed from a one-time shock: the pandemic itself and the expansionary fiscal and monetary policies enacted in response to it. Not only have these inflationary pressures mostly subsided by now, but they were always largely independent of the structural dynamics that govern long-term rates anyways (e.g., demographic trends, long-term growth expectations, etc.). There is a lot of middle ground between “transitory” and “the next 30 years.”

All things considered, we believe that the market is directionally right about long-term rates, but for the wrong reasons. As a result, we suspect it could also be wrong about the magnitude: long-term interest rates should be higher now than they were five years ago, but probably lower than they are today. We expect that they will eventually settle somewhere in the middle.

THE CASE FOR EUROPEAN DATA CENTRES

Two years ago, [we wrote](#) about how the data centre market was evolving in ways that would create significant opportunities for real estate investors willing to innovate in the sector. A great deal has happened in real estate markets since we wrote that letter, but we believe that Europe’s data centre sector looks even more attractive today than it did then. As we discuss below, we see an exceptional investment environment now emerging on three separate fronts:

- (1) data centre demand—which was already expanding in Europe—has now started to accelerate, driven by both (a) longstanding demand drivers and (b) new AI-related requirements;
- (2) due to power constraints, new data centre supply is tightly restricted in Europe, insulating incumbent owners against competition and mitigating certain obsolescence risks; and
- (3) despite robust fundamentals, major data centre owners are selling assets as they look to recycle capital, creating opportunities to acquire well-located, operational data centres at fair prices.

Old tailwinds continue to blow, particularly in Europe

We start by looking at recent history. Over the last decade, global data centre take-up was driven by the digitalisation of consumer and commercial activities of all stripes. Shopping left brick-and-mortar stores and relocated to the internet, with global e-commerce sales growing from \$1.2 trillion in 2012 to \$6.5 trillion in 2022.¹ Entertainment moved online, led by the rise of streaming services, like Netflix and Spotify, and social media platforms, like TikTok and Instagram. Awash with consumer data, businesses started to use digital information in increasingly sophisticated ways, through targeted advertising and data analytics. Altogether, the adoption of new digital technologies resulted in a 30-fold expansion in the volume of created, stored, delivered, and consumed data between 2010 and 2020.²

Digitalisation required digital infrastructure, including a massive amount of new hardware—like fibre-optic cables, networking devices, and computer servers—which ultimately consolidated into dedicated data centre buildings. This outcome was the product of a second technological development that took place alongside wider digitalisation: the rise of cloud computing services, like Google Cloud, Amazon Web Services, and Microsoft Azure. (In short, the cloud delivers computing infrastructure and software via the internet, making it possible for users to effectively bundle their hardware and store it remotely under a shared roof. We wrote about cloud computing in depth [in our 2022 letter](#).) Coupled with the cloud, digitalisation generated sustained growth in data centre demand

¹ Statista

² Statista

throughout the last decade, particularly in the United States. Annual capital investment in hyperscale data centres tripled between 2014 and 2020, increasing from \$30 billion to \$94 billion.³ America's data centre capacity more than doubled in turn, growing from less than 7 GW in 2014 to over 17 GW by 2022.⁴

Notably, however, these two trends—(1) digitalisation and (2) the consolidation of digital infrastructure into data centres—have not yet reached maturity in Europe. As such, we think that both have considerable runway for growth over the next 10 years.

Consider three examples. [As we discussed in our last letter](#), e-commerce activity in Europe has historically lagged e-commerce activity in the US. However, we believe that Europe's e-commerce sector is now set up for rapid growth, as e-commerce firms start to invest aggressively in the continent's logistics infrastructure. The same holds for video streaming—another domain where European consumers have trailed Americans, with user penetration at just 26% in Europe versus 40% in the US. As streaming service providers increasingly focus on growth in less saturated markets, Europe has started to catch up here, too: Europe's streaming market is now growing at faster pace than America's, and the number of European users is expected to exceed the number of American users for the first time in 2024.⁵ Enterprise cloud migration is tracing a similar trajectory. Public cloud spending currently represents 15% of total enterprise IT spending in the US, compared with 12% in the UK and the Netherlands, 9% in Germany and France, and 7% in Italy and Spain.⁶ Despite these gaps, European cloud migration is now picking up: IDC expects European public cloud spending to grow from \$148 billion in 2023 to \$258 billion by 2026, at a CAGR of about 20%.

All told, KPMG estimates that Europe has just 1 MW of data centre capacity per 100,000 persons, compared with 12 MW in the US. As existing digital technologies continue to gain wider adoption across Europe, demand for European data centre capacity should expand in tandem, and increasingly close this gap.

Enter AI

Over the last 12 months, we also saw a new source of demand emerge: generative artificial intelligence ("GenAI"). Under the surface, developments in AI had been accelerating for some time, supported by a doubling of global investment in the sector between 2019 and 2021, from about \$45 billion to over \$90 billion.⁷ However, GenAI entered the public spotlight in late 2022, with the launch of OpenAI's ChatGPT chatbot. ChatGPT quickly realised the steepest adoption curve in the history of consumer technology, reaching 100 million monthly active users in just two months. (For comparison, it took Instagram 2.5 years to reach a similar number of users.) Looking ahead, Goldman Sachs now expects annual AI-related investment to rise to \$160 billion within the next two years. As investment ramps up, applications designed for specific use-cases will increasingly come online. In turn, we should expect to see AI technologies deployed across a growing number of consumer and commercial settings, with increasing intensity.

Although it's too early to know exactly how AI will play out, initial indications suggest that AI's impact on the data centre market will be positive and substantial, and that AI-related requirements will generate demand for data centres of multiple types. On the one hand, AI models require vast amounts of high-performance computing power to train. (It takes an estimated 10 GWh of power to train a large language model, like ChatGPT. This is approximately equal to the annual power consumption

³ Credit Suisse

⁴ McKinsey

⁵ Ampere

⁶ Gartner

⁷ Goldman Sachs

of 1,000 American households.) Since model-training will generally not involve real-time processing or require low-latency transmission, it will tend to take place in far-flung hyperscale data centres with hundreds of megawatts of cheap power.

On the other hand, most AI technologies will get used by consumers and workers in the places where those people live and work. This will effectively involve frequent communication between user devices (in pockets, homes, and offices) and AI models (in data centres). To keep latency low and minimise network congestion, these technologies will need to be serviced by data centres located close to population centres, at the so-called “edge.” These requirements are likely to be immense. For example, while a typical Google Search query consumes about 1 kJ of power, a typical GPT-4 query consumes between 4 and 36 kJ, depending on the complexity of the request. In short, the directional trend is already clear: AI will generate significant incremental demand for data centre capacity over the next decade, particularly at the edge.

“Buy land power—they’re not making it anymore”

Next, we turn to the supply landscape. Against a backdrop of growing demand, Europe’s data centre supply is structurally constrained. As a result, we expect new data centre requirements to widely outpace new capacity over the next several years, leading to a persistent supply/demand imbalance.

New supply is constrained for a simple reason: data centres require a great deal of power, and power is increasingly scarce in Europe. This is particularly true in Europe’s major population centres, where AI deployments will require significant capacity. For example, in urban parts of England, new data centre development has been effectively throttled by an obsolete power grid and a sclerotic national energy-provider. In a public letter released last summer, Virtus (one of the largest domestic data centre operators in the UK) said that it had forgone about £500 million of investment in Southeast England due to “a combination of lack of immediate power combined with delays in future power.” Meanwhile, the queue for connecting new-build data centres to the UK’s power grid grows longer each year, with some 70% of applicants in the last 12 months granted connection dates beyond 2028 and a quarter granted connection dates beyond 2032.⁸

Data centre developers elsewhere in Europe face similar constraints. Last year, the Netherlands imposed a nine-month ban on all new hyperscale builds, citing power supply and sustainability concerns. In Ireland, where data centres are on track to consume 70% of the country’s power by the end of the decade, the government is now flooding the industry with new regulations, in principle aimed at encouraging greener data centre projects but in practice likely to curtail development altogether. All told, it’s getting harder to build new data centres in Europe at a time when new data centre capacity is exactly what European customers require.

This dysfunction bodes well for incumbent providers. Specifically, barriers to new development should help insulate existing data centres against competition, benefitting owners across virtually all segments of the asset class. For one, due to inelastic supply, rising demand will put substantial upward pressure on rents. In fact, we have already started to see meaningful rental growth materialise: last year was one of the strongest periods in over a decade for European data centre rates. At the same time, constrained supply should also bestow data centre owners with qualitative bargaining powers over their customers. In particular, future demand will simply not fit into shiny new-build facilities. As a result, customers that require low-latency capacity will increasingly need to concede on their

⁸ TechUK

“nice-to-have” demands—e.g., for ultra-low PUEs⁹—and instead absorb available capacity wherever it still exists.

Under these circumstances, older but well located data centres should not face material obsolescence risks. Ultimately, power in the right location is the limiting factor in this market; it overrides almost all other requirements. Thus, as long as (1) power allocations close to population centres remain effectively frozen and (2) demand continues to rise, then older, well located data centres must be made to work. And since incumbent data centre owners already control these (inelastic and increasingly undersupplied) power allocations, the cost of making older facilities work should predominantly fall on their customers in the near-to-medium term. Accordingly, not only do we see limited obsolescence risk in the older segment of the edge market, but we also think that, as demand outpaces supply, occupiers should end up footing much of the capex needed to upgrade older assets, either directly or through higher rents.

An attractive entry point

As we described above, accelerating demand growth and tight supply constraints have combined to create a favourable competitive landscape for European data centre owners. Fundamentals this strong are not commonly paired with a large opportunity-set for value-minded investors. However, current market conditions are somewhat unique. While publicly traded data centre operators are now investing heavily to capitalise on accelerating demand, listed real estate multiples have contracted substantially in recent years, depressing their stock prices. By contrast, private-market cap rates have not expanded commensurately. This has made public equity markets relatively expensive places for data centre operators to raise capital. To fund new investment, these companies are instead opting to recycle existing capital by selling off portfolio assets. Consequently, we now see highly motivated sellers looking to quickly dispose of well-located European data centres at fair prices. This creates an attractive entry point for buyers.

BACK TO THE OFFICE

Back in March, [we wrote](#) that “as technology companies look to get more productive, we may see office occupancy continue to rebound at a faster pace, and to a higher equilibrium level, than the office bears had expected.” Recent data have reinforced that view. Over the last few quarters, we saw a meaningful uptick in the number of employers that have started to implement—and actually enforce—stricter office attendance policies. What’s more, this latest push back to the office is being led by the American tech-sector firms that had adopted some of the laxest in-person attendance policies in the wake of the pandemic.

For example, in August, Amazon sent emails to employees who were “not currently meeting our expectation of joining your colleagues in the office,” warning those workers that “we expect you to start coming into the office three or more days a week now.” A couple months earlier, Google notified its workers that office attendance would count in their future performance reviews. (Like Amazon, Google has adopted a three-day mandatory minimum.) Around the same time, the ride-hailing company Lyft said that most employees would need to spend three days per week in the office, barely a year after telling workers that they could work indefinitely on a fully remote basis. Salesforce now wants all of its customer-facing employees in-person four days a week, and even Zoom (which presumably has no commercial interest in leading by example here) has ordered employees who live within 50 miles of an office to come into the workplace at least twice per week. This is not quite a

⁹ “PUE” refers to power usage effectiveness, which is the ratio of the total power consumed by a data centre to the amount of power available for use by the IT equipment in that data centre. A higher PUE implies lower power efficiency.

reversion to pre-pandemic work patterns, but it hardly looks like a full-blown “remote work revolution” either.

The same directional trend is evident outside the tech sector, too, where organisations that had previously required their workers to come into the office three days per week are now pushing that requirement higher. At the beginning of the summer, Blackrock announced that its employees would need to come into the office four days a week, starting in September. JP Morgan is now instructing its managing directors to show up every day in-person, explaining in an internal memo that “they have to be visible on the floor, they must meet with clients, they need to teach and advise, and they should always be accessible for immediate feedback and impromptu meetings.”

What is driving this return to the office? Executives and managers seem to believe that remote work is simply less productive. This sentiment is reflected in a recent survey conducted by KPMG, which found that about two-thirds of CEOs at large multinational companies want their organisations to be full-time in-office by 2025. A growing pile of evidence suggests that these managers may have a point, too.

A new article by the economists Jose Maria Barrero, Nicholas Bloom, and Steven Davis surveys recent academic research on the topic, concluding that “fully remote work is associated with about 10% to 20% lower productivity than fully in-person work.” Analysing data from 2020, a different study published last year by the New York Federal Reserve found that non-commuting workers reduced their work time by about 3.5 hours per day, trading that time for about one extra hour of sleep and 2.5 extra hours of leisure. At least in practice, remote work seems to mean less work.

Finally, even workers themselves seem to acknowledge the shortcomings of remote work. In a survey conducted last year by McKinsey, employees scored hybrid work arrangements as worse than fully in-person work arrangements with respect to multiple “factors that impact ability to effectively perform work,” including ability “to learn new skills to meet changing job expectations,” “mental-health issues,” and ability “to share your full self at work.” We obtained similar results in a [poll that we recently commissioned](#) of 1,800 UK office workers. Some 59% of young workers (ages 18-24) said they worked less productively at home, while 43% of young workers reported feeling socially isolated while working from home. Evidently, many employees recognise the professional and social value of regular, in-person interaction with colleagues, bosses, and mentors.

Some amount of remote work is surely here to stay. This is a good thing: not everything that used to take place in offices needed to take place there. At the same time, offices serve important functions that cannot be perfectly replicated online. This has been borne out by the work-from-home experiment run over the last three years. As employers and workers alike increasingly come to terms with limitations of remote work, we expect to see office buildings continue to fill up.

WORKERS OF THE WORLD, UNITE?

While some workers spent the past year crawling back into the office on their bosses’ orders, other workers decided to exercise their bargaining powers by going on strike. Over the last 18 months, the US and parts of Europe experienced widespread labour disputes of a magnitude not seen in at least two decades, with hundreds of thousands of workers walking off the job to demand higher pay, better working conditions, and a mix of other

Out of character for America, as of October, the US had seen some 300 strikes this year, involving nearly half a million workers, ranging from baristas at Starbucks and Amazon drivers to machinists at a Boeing supplier and Hollywood writers. On the more strike-prone side of the Atlantic, the Netherlands saw more industrial actions in 2022 than in any year since at least 1999. And in the UK, as Figure 1 shows, the number of working days lost per month due to labour disputes has spiked this

year to a level not experienced since 1989. While this uptick looks insignificant relative to the levels of labour unrest typical in the UK pre-1990, it nonetheless suggests a meaningful departure from the status quo of the last three decades.

FIGURE 1. LABOUR STRIKES REACHED A THREE-DECADE HIGH IN THE UK THIS YEAR¹⁰



Economic conditions during the last few years created a perfect storm for labour-market unrest. Negligible (or even negative) real wage growth gave workers a strong motive to push for better pay. Tight labour markets provided an easy opportunity to strike from a position of strength. Under these circumstances, recent strikes, or mere threats to strike, have often delivered significant victories for the workers. Most famously, in July, UPS and Teamsters reached a deal that will see UPS drivers earning \$170,000 in total compensation by the end of the five-year contract.

Looking ahead, we think these labour disputes herald a broader rebalancing of power in favour of workers throughout advanced economies, particularly those at the less skilled end of the labour force. This rebalancing cuts against the trend that prevailed during the last several decades, when offshoring and automation eliminated millions of industrial jobs in the West. Increasingly, this is no longer the case. Manufacturing activity is returning to advanced economies. Labour shortages have become ubiquitous, especially in service sectors. Wage growth for lower-income workers is now robust. To the extent that these trends persist, lower-income workers in the US and Europe will likely continue to benefit.

Disclaimer:

The views contained herein are the views of Castleforge Partners Limited ("CFP") as of the date indicated and such views are subject to change without notice. CFP has no duty or obligation to update the information contained herein. This letter is being made available for educational purposes only and should not be used for any other purpose. The information contained herein does not constitute and should not be construed as an offering of advisory services or as an offer to sell or solicitation to buy any securities or related financial instruments in any jurisdiction. Accuracy of such information and has not independently verified the accuracy or completeness of such information or the assumptions on which such information is based. This letter, including the information contained herein, may not be copied, reproduced, republished or posted in whole or in part in any form without the prior written consent of CFP.

¹⁰ ONS