September 2017

Factor investing the ATP way

The Danish Labour Market Supplementary Pension Fund (ATP) manages pension assets of €100bn on behalf of the Danish population. In 2015, ATP revised its investment portfolio strategy, resulting in a new portfolio allocation approach based on risk factors rather than asset classes. The fundamental idea is that assets are perceived on the basis of their exposure to a well-defined set of common factors (risk sources), such as equity, interest rate and illiquidity factors. The labelling, i.e. the asset class, is not important. The factor approach offers a rigorous framework for composing the desired risk profile and a method for constructing a total portfolio with this profile.

Authors

Kasper Ahrndt Lorenzen
Chief Investment Officer (CIO)
Email: kas@atp.dk

Kasper Ahrndt Lorenzen (MSc Economics) is CIO at ATP.

Søren Fiig Jarner
Vice President
Email: sj@atp.dk

Søren Fiig Jarner (PhD) is Vice President for Quantitative Research at ATP and Affiliated Professor of actuarial science and statistics at the University of Copenhagen.

The ATP Group, Kongens Vænge 8, DK-3400 Hillerød.

Many people have contributed to the development of ATP's investment strategy. We would like to credit, in particular, Mads Gosvig and Henrik Gade Jepsen for the key role they played in the design of the fundamental ideas and principles.

Introduction

ATP is a statutory, savings-based pension scheme for wage earners and recipients of transfer income in Denmark. The scheme was founded in 1964 and it has today 5 million members and net assets in excess of €100bn. ATP's primary objective is to provide a supplement to the (old-age) state pension. Half of today's pensioners have no other pension income than state pension and ATP. Also on longer horizons, state pension and ATP will continue to be the main source of income for low-income pensioners. Given ATP's role as provider of basic financial security, ATP has maintained that all pension entitlements are guaranteed and lifelong (small pensions are however paid out as a lump sum at retirement). At the top level, asset management is divided into two main activities: hedging and investment. The purpose of the hedging portfolio is to ensure that ATP will always be able to honour the issued pension guarantees. This is achieved by investments in bonds and other interest-bearing instruments with fixed payments.

The purpose of the investment portfolio is to generate a return in excess of the return already guaranteed. While the hedging portfolio is exposed to interest rate markets only, the investment portfolio has a broad market exposure to, for example, listed and unlisted equities, government and credit bonds, commodities, infrastructure, real estate, as well as direct investments in specific companies. Investment portfolio returns are used to increase members' pensions – in addition to the level originally guaranteed – and to finance higher provisions resulting from increasing life expectancy.

The separation into hedging and investment activities arose in the wake of the introduction of market value accounting in Denmark at the turn of the millennium.[1] Market value accounting marked the beginning of a development, led by the Danish Financial Supervisory Authority (FSA), with increasing focus on managing and measuring the underlying risks rather than quantitative limits on individual asset classes. In the same period, ATP's approach to investment portfolio construction mirrored this development: from allocation to asset classes to allocation with focus on the underlying factors. The latest step was taken in 2015, when ATP fully adopted factor investing. The transition was a 'quiet revolution' with far-reaching implications for objectives, portfolio construction, performance evaluation and risk management. In this paper, we describe ATP's implementation of factor investing. We start with a brief introduction to ATP's business model and risk management framework, which is necessary for understanding the chosen implementation.

1 The impact of market value accounting on ATP's business model is described in Jepsen (2006).

ATP's pension product and business model

A typical wage earner pays approx. €435 to ATP per year. Guaranteed pension rights are acquired for 80% of the contribution, while the remaining 20% enters collectively owned free reserves (bonus potential). The guarantee is calculated on the basis of ATP's annually updated forecast of the member's life expectancy and the current interest rate level. When the bonus potential is sufficiently large, a part of the funds is used for indexation of the guaranteed pensions (bonus). The pension, including bonus, is paid out as a life-long, monthly benefit when the member reaches his or her state pension age. In 2016, ATP received total contributions of €1.3bn and paid out €2.1bn in pension benefits.

Year-end 2016, ATP had made provisions of €88bn for guaranteed pensions, while the bonus potential was close to €14bn. The provision of €88bn corresponds to the present value of future guaranteed benefits, see also the comments in Fact box 1. By default these funds belong to the hedging portfolio, the sole purpose of which is to ensure the return necessary to meet the guarantees, regardless of the future interest rate level.

The hedging portfolio could, in principle, invest all €88bn in long-term bonds to ensure the necessary return. In practice, however, only about half of the return comes from long-term bonds. The other half is generated by a large number of interest rate swaps entered by ATP and various financial counterparties. Under these contracts, ATP receives future fixed payments in exchange for variable payments determined by the short-term interest rate. In brief, this corresponds to owning a bond financed by a variable rate loan.

Given the structure of the hedging portfolio, ATP only needs to 'bind' half of the €88bn in long-term bonds to hedge the guaranteed pensions. The other half is available to the hedging portfolio in the form of liquidity ('money'). The hedging portfolio may choose to place those funds in short-term bonds, which give a return corresponding to the variable payments on the swap, or it may make some of the funds available for the investment portfolio.

Chart 1 illustrates the structure of ATP's balance sheet year-end 2016. The investment portfolio is approx. €28bn, i.e. the value of the bonus potential of €14bn and a 'loan' from the hedging portfolio of another €14bn. The €14bn is invested broadly in equities, government and credit bonds, commodities, infrastructure, real estate and other assets. The return from the investment portfolio accrues in the bonus potential, after the hedging portfolio has received interest on the loan and after tax has been paid. From the point of view of the hedging portfolio, it makes no difference whether its liquid assets are at its own disposal or whether they have been lent to the investment portfolio. What matters is that the hedging portfolio receives a (short-term) market interest rate on its liquid assets.

The investment portfolio's task is to deliver as high a return as

CHART 1: Stylised balance sheet of ATP year-end 2016

Note: *The bonus potential (BP) amounts to €14bn, while a provision of €88bn has been made for guaranteed pensions (GP). The investment portfolio (IP) has 'borrowed' €14bn from the hedging portfolio (HP), to be invested together with the bonus potential.*

possible for a given amount of overall *risk*. The risk is limited by the size of the bonus potential, since losses must never cause the bonus potential to fall below a certain level. The most effective use of risk is typically achieved by investing in both very risky assets ('equities') and in – a comparatively larger amount of – less risky assets ('bonds'). Hence the invested *capital* will typically be substantially higher than the bonus potential. Access to liquidity from the hedging portfolio helps ATP to maintain a well-diversified investment portfolio with efficient use of risk. Hence, in contrast to many other investors, ATP's primary investment constraint is the size of the risk budget rather than the amount of invested capital.

The separation into hedging and investment assets corresponds to Robert Merton's classic 'separation theorem', see Merton (1992). According to this result, the optimal portfolio for an investor consists of a combination of the risk-free asset (hedge) and a diversified portfolio (investments). The relative weights of the two portfolios depend on the investor's risk appetite and the size of the (risk-adjusted) returns obtainable in the financial markets.

Investment and risk management principles

Long-term value creation in ATP consists partly of the return embedded in the guarantees, partly of the return generated by the investment portfolio. Investment returns accrue in the collective bonus potential and are subsequently transferred to the individual members' pensions in the form of bonus. In addition, since the turn of the millennium a substantial share of the investment returns have been used to finance increased provisions as a result of increasing life expectancy.

ATP must at any time be able to honour the pension guarantees. Hence, ATP must always hold assets at least matching the value of guaranteed pensions. The bonus potential is the difference between the assets and the value of guaranteed

pensions, and if it comes under pressure, ATP will have to reduce the investment risk to protect the guarantees. In order to ensure the ability to take investment risk (needed for high long-term value creation) it is crucial to protect the bonus potential against large losses.

ATP's investment activities and overall risk management rest on four fundamental principles with the aim of protecting the bonus potential: hedging of pension guarantees, risk diversification, protection against tail risks and dynamic adjustment of the level of risk. The principles are shown in Chart 2 and described in more detail below.

Hedging of pension guarantees
All pension guarantees are hedged in full by a separate hedging portfolio. The purpose is partly to ensure the return necessary to meet the guarantees, partly to protect the bonus potential. The hedging portfolio and the value of the guaranteed pensions have the same interest rate sensitivity (after tax), meaning that they move in 'sync', when interest rates change. The result is immunisation of the bonus potential to interest rate movements.

Risk diversification
The investment portfolio is broadly invested, both across markets and across assets in each market. This gives exposure to many different risk premia and reduces the risk of simultaneous losses. Investments are allocated on the basis of risk, not on the basis of invested capital.[2] Risk is allocated across several different markets so that no single risk source dominates. The aim is to have a high *risk-adjusted* return on the total portfolio.

Protection against tail risks
Under 'normal' market conditions, investment in many assets in a given market provides good protection against simultaneous losses. For example, an equity portfolio with many stocks entails a significantly lower risk of loss than a portfolio with only one stock. However, this type of diversification offers no protection against systematic risks, such as the general drop in stock markets that occurred during the financial crisis, or surging inflation that undermines the purchasing power of pensions. Financial insurance contracts ('options') can be used for specific protection against such 'tail risks' against payment of a premium. More indirect protection can be achieved by investing in specific assets, e.g. gold, which, historically, has yielded high returns during crises, or by implementing dynamic strategies, where allocation (of risk) to, for example, equities depends on the latest developments. ATP uses all these methods – to a greater or lesser extent – to reduce tail risks (risks with small probability, but major implications).

Dynamic adjustment of risk level
The bonus potential protects the guarantees by absorbing losses in the investment portfolio. The bonus potential must, however, also cover risks from life expectancy developments, hedging of the guarantees and operational errors. ATP uses a proprietary model to calculate the overall risk from all material quantifiable risks, and the level of risk in the investment portfolio is dynamically adjusted such that the total risk of major loss of bonus potential is under control.

FACT BOX 1: Accounting rules and discount curve

ATP shifted to market value accounting in 2002 and developed its business model accordingly, as outlined. Under the former accounting rules, provisions were calculated using solely Danish market rates (the 30-year interest rate was used for discounting all liabilities with maturities from 30 years and longer). Under this regime, the natural asset to ensure the guaranteed return was a bond, owned either directly or indirectly via an interest rate swap. This is the paradigm we have used as starting point for our description of ATP's business model.

In 2016, the European solvency rules for life and pension insurance companies came into effect, i.e. the Solvency II rules. Under these rules, the discounting rates used for the longest liabilities are no longer pure market rates, but combinations of market rates and a centrally fixed long-term interest rate of currently 4.2% – the Ultimate Forward Rate (UFR). ATP is not subject to Solvency II, but chose to make a similar change to its discount curve at the end of 2013. In ATP, liabilities with maturities of 40 years or longer are all discounted with 3% irrespective of market interest rates. The background for the change was that, in reality, there is no liquid financial market for very long-term payments to be used as a basis for objective pricing. Instead, the 3% reflects the return which ATP, with a high degree of certainty, expects to achieve on an investment portfolio if held for a long period.

Factor investing

Factor investing is based on the observation that financial assets are exposed to a limited set of common factors, and the dynamics of these factors explain a large part of the return in the various asset classes. This perception of the financial universe leads to a unified framework for risk profile composition, investment portfolio construction, risk management, and performance evaluation. Consequently, factor investing is not only a portfolio construction method, but rather a framework for approaching the entire investment process. Andrew Ang's highly readable book Ang (2014) gives a comprehensive account of factor investing as an investment approach. Here, we will focus on the three most important aspects of the factor-based approach from the perspective of ATP.

2 Risk of investment assets is measured by expected shortfall (ES), which is a measure of the probability of losses and their size. ATP uses a model-based ES measure on a 3-month horizon.

Allocation to factors

We start with an example illustrating the difference between allocation to factors and the more traditional allocation to asset classes. The starting point for a traditional (simple) portfolio construction approach is to have a certain allocation to equities and a certain allocation to bonds, or, slightly more sophisticated, equities with a certain amount of *risk* and bonds with a certain amount of *risk*. As long as you invest only in listed equities and government bonds, this approach is sufficient, but if you also want to include, for example, credit bonds (e.g., loans to corporations) in your portfolio, you encounter problems. At first sight, a credit bond seems to be a bond, but the price very much depends on the financial standing of the issuing corporation. In reality, a credit bond is a hybrid between a stock and a government bond.

In the case at hand, the problem could be solved either by adding the rule that, for example, half of the bond risk should come from government bonds and the other half from credit bonds,[3] or by introducing an additional category to the portfolio, so that it now comprises equities, government bonds and credit bonds. The problem is that for each new asset class you have to either impose more rules specifying how the risk should be distributed within the existing categories or add more categories to your portfolio. Both approaches increase the complexity making it more difficult to understand and manage the characteristics of the total portfolio.

In contrast, factor investing focuses on the amount of *equity risk* and the amount of *interest rate risk* in the portfolio. Equities hold primarily equity risk, and government bonds hold primarily interest rate risk. A credit bond, on the other hand, holds equity risk and interest rate risk in approximately equal measures. For a factor investor, it is, in principle, of no importance whether the exposure to, for example, equity risk originates from equities or credit bonds; portfolio characteristics are (primarily) determined by the aggregate exposure to equity and interest rate risk, respectively, and not by the type of assets used to achieve the exposure. There is no preconceived view on the desired allocation to each asset. Rather, assets are chosen to provide the desired risk exposure at the best price.

The two sources of risk are referred to as *equity factor* and *interest rate factor*, respectively, and in factor investing lingo, a portfolio is *exposed* to a collection of factors. In practice, there will be more than two factors, but the ideal is to have a small number of factors which together describe all the assets held. The most important consequence of the factor approach is that it provides a clear and concise picture of the characteristics of a given portfolio – regardless of the number of underlying assets and their combination.

Bad states

A key aspect of factor investing is the identification of particularly adverse scenarios. We refer to these states of the world as *bad states* (Ang (2014) uses the term *bad times* to the same effect). The bad scenarios depend on what you want to achieve as an investor and on the restrictions you face. The chosen factors must be able to describe the bad scenarios, and the factor allocation must reflect how much return you are willing to give up mitigating the effect of bad states. This trade-off between return and risk reflects your risk preference. Most investors will view large losses as bad states, but there may also be other less obvious bad states.

CHART 2: The four fundamental investment and risk management principles for ATP

ATP's objective is to provide pension to ensure basic financial security for a large proportion of Danish pensioners. Since the guaranteed pensions are nominal, the value for the recipients very much depends on ATP's ability to increase pensions via bonus to keep up with inflation. Consequently, a scenario with high inflation *without* ATP being able to increase pensions sufficiently is a bad state.

Inflation is therefore included in ATP's factor universe, as high inflation is a significant risk to the members of ATP. Inflation exposure is included in the portfolio construction on an equal footing with exposure to equities and interest rates. However, in contrast to the equity and interest rate exposure, the purpose is not to generate regular returns, but to create an extraordinarily high return in the event of high inflation.

Choice of factors

The number of factors is a balance between, on the one hand, a succinct universe that offers good intuition about the interdependence of the factors, and on the other hand a more complex, but also more accurate description of the individual assets. Generally, parsimony outweighs complexity when it comes to the overall allocation, while it can be useful to consider more specific sub-factors when implementing the general factors. Having few factors also facilitates clearer communication and focuses the discussion of the overall investment profile.

All factor investors agree that few factors 'spanning' the relevant universe in an intuitive way is ideal. However, the concrete choice of factors will often depend on investor-spe-

3 In that case it would be natural to also increase the bond share of the overall risk to compensate for the fact that bond risk is now only partly interest rate risk.

cific circumstances. In view of ATP's size and dynamic risk management, it is of key importance for ATP that the factor exposure can be implemented on a large scale using liquid instruments. This in turn impacts the choice of factors. Smaller investors or investors with other business models will not necessarily choose the same factors as ATP.

The three liquid market factors

ATP uses three liquid market factors and a fourth 'illiquidity factor' to describe the overall allocation. The number highlights the desire for clarity, while the choice of the specific factors reflects the markets and risks to which ATP wishes to be exposed. This section discusses the three liquid market factors, while the fourth – more complex – factor will be described in the next section.

The equity and interest rate factors

In the Western world, the equity and interest rate markets in the respective countries are the two primary *liquid* sources of return (and risk). These markets also determine to a wide extent the development of a range of other markets, such as credit bonds and private equity. Moreover, they impact – to a lesser extent – the market for commercial real estate. Furthermore, large secondary markets for derivatives are closely linked to the equity and interest rate markets.

For an investor such as ATP, desiring and requiring substantial liquid exposure, the equity factor and the interest rate factor therefore clearly belong to the set of fundamental factors. As a result of global trade and other relations between countries, the various equity and interest rate markets have many commonalities. Crises and recoveries are global phenomena, although their severeness and timing may vary across countries. Therefore, deciding the total amount of equity and interest rate risk, respectively, is of more importance than the choice of the markets in which to obtain this exposure. The factor approach helps prioritising these decisions.

In practice, ATP invests in both domestic and foreign equity and interest rate markets. Exposure is obtained via both equities and bonds (*cash assets*) and *futures* on equity and bond indices. In line with the factor approach, ATP does not, in advance, have any preferences as to whether the exposure is obtained in cash assets or futures. That depends on other considerations, e.g., liquidity demands. Currency risk is minimised by hedging all exposures denoted in currencies other than Danish kroner or euros. The rationale being that currency exposures in general do not yield a *systematic* return. In other words, absence of hedging consumes risk budget with no expected reward.

The inflation factor

According to the factor approach, factors are selected on the basis of their ability to describe or 'span' the relevant markets, making the equity and interest rate factors must-haves. In addition, the factors must be able to represent the investor's bad states, such that the allocation can take into account preferences for bad state protection. As previously identified, high inflation without corresponding indexation of pensions is the primary bad state for the members of ATP. The third liquid market factor in ATP's factor universe is therefore the inflation factor.

Inflation exposure can be obtained in many ways. It can be obtained either directly in the form of exposure to commodities such as oil, gold and other metals, or partly through exposure to, for example, real estate or infrastructure, where rental income or selling prices are often inflation-linked. This type of exposure is known as *linear* exposure, as revenue rises and falls in line with inflation. In some situations, non-linear or *convex* inflation exposure is wanted, e.g. a high return when inflation is high, without a corresponding low return when inflation is low. Such exposure can be obtained by purchasing options, and ATP has an extensive option programme of this type.

In addition to high inflation, large losses of bonus potential are also considered to be bad states. Losses reduce ATP's ability to grant bonus today, and they also reduce ATP's risk capacity and thereby the future value creation. The equity and interest rate factors constitute the two main sources of return and risk for ATP as well as the two main sources of potential losses (of bonus potential). Protection against losses from these factors can be achieved by buying equity puts or swaptions, and such insurance is ascribed to the inflation factor. High inflation, big drops in equity prices and substantial interest rate increases are typically consequences of the economy generally being in a bad (macro) state. The inflation factor is used to represent exposures that protect against general bad states in the economy. Since (high) inflation is the most important of several possible bad states in the economy, the factor has been named accordingly, although it covers 'bad state' exposures in a broader sense.

This means that the inflation factor contains both the linear exposure to inflation and various types of insurance strategies. The inflation exposure is expected to provide protection against bad states as well as diversification of the returns from the equity and interest rate factors, but apart from this no further systematic contributions to the total return over time is expected.

Risk parity

Once the relevant factors have been identified, the desired exposure to each of them must be decided. The decision falls naturally in two parts: the total factor exposure and the relative exposure to the individual factors. The size of the bonus potential and ATP's other risks (primarily longevity risk) induce an upper limit to investment risk, and the total factor exposure is therefore only partially an investment decision. Here, we focus instead on the relative factor allocation.

In factor investing the focus is on the exposure to the under-

lying factors, rather than on their 'labelling' in the form of asset classes. Exposure is measured by the associated risk, i.e. by the size and probability of loss rather than by the capital committed. Allocation thus becomes a question of distribution of *risk* to the various factors.

ATP's portfolio construction follows the *risk parity* approach, which prescribes that equal amounts of risk are allocated to the fundamental return factors. In the case of ATP, this means equal amounts of equity and interest rate risk. There is both empirical and theoretical evidence of risk parity portfolios generally outperforming capital-weighted portfolios, see Dalio, Prince and Jensen (2015) and Asness, Frazzini and Pedersen (2012). One of the theoretical arguments is that many investors either cannot or do not want to implement a risk parity portfolio, as this requires a lot of capital or use of derivatives (to obtain sufficient interest rate risk). The observed excess return of risk parity portfolios is 'reward' to investors without these limitations. Another argument in favour of risk parity is that it leads to more robust portfolios compared with portfolios based on classical mean-variance optimisation.

Chart 3 shows the cumulative excess return of a traditional portfolio with 60% invested in equities and 40% in bonds and from a risk-weighted portfolio with equity, interest rate and inflation risk in the proportion 40/40/20.[4] The latter portfolio is very close to ATP's long-term reference for the liquid market portfolio. Both portfolios are rebalanced every month to an (annualised) volatility of DKK 1. The two strategies are thereby equally risky, and the accumulated excess return divided by the length of the period (in years) gives an estimate of the Sharpe Ratio (SR) over the period.[5]

The risk-weighted portfolio has a high SR of more than 90%, while the 60/40 portfolio has an SR just below 70%. The cumulative returns moved in tandem until the turn of the millennium, but hereafter the risk-weighted portfolio performed better, especially during the financial crises around 2001 and 2008. Note that *risk* remains constant in Chart 3, while the principal (the amount borrowed/invested) varies from month to month. The chart thus illustrates the difference in *risk-adjusted* returns for the two strategies. Alternatively, we could have shown the cumulative return from continually reinvesting an initial principal of DKK 1 (total return index) without taking into account the different risks

CHART 3: Accumulated excess return of a risk-weighted portfolio and a portfolio with 60% equities and 40% bonds

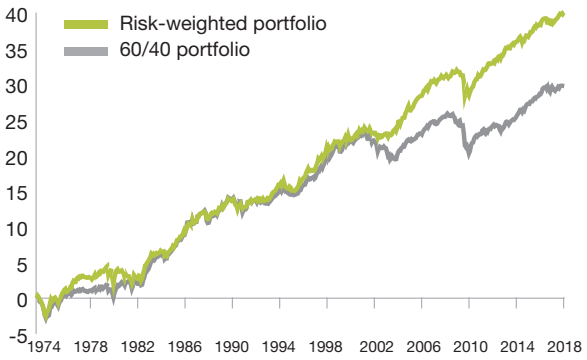

Note: *The chart shows the cumulative excess return of a risk-weighted portfolio and a portfolio with 60% equities and 40% bonds. Both portfolios have a constant (annualised) volatility of DKK 1.*

in the strategies. The risk-adjusted comparison was chosen as the best reflection of the returns obtainable by ATP as an investor restricted by risk, rather than capital.

The fourth factor

ATP assumes a large part of its investment risk in the traditional liquid markets. However, from both a diversification and return perspective it is desirable to also have a considerable exposure outside the traditional liquid markets. ATP obtains this exposure by investing in illiquid assets and by implementing alternative liquid risk premium strategies. These investments are analogous in terms of risk, and the fourth factor in ATP's factor universe is chosen to represent the distinctive risks associated with these types of investments.

Illiquid assets

Direct investments in real estate, private equity and loans to individual companies all require highly specialised investment skills, as prices and terms are negotiated directly between the parties. The investments typically have 5-10 years duration and they are *illiquid* in the sense that in practice it is not possible to retrieve the invested capital prematurely – at least not without significant losses. On the other hand, illiquid assets can provide access to attractive, alternative sources of return not obtainable from the liquid markets.

In contrast to liquid assets, there is typically no unique market price for illiquid assets. The assets are, by definition, only rarely sold, and the market price of specific buildings or companies can often be difficult to assess. The value of illiquid assets is therefore often based on model calculations, so-called *mark-to-model*. In normal market conditions, valuation by mark-to-model causes a more stable evolution of returns on illiquid assets than on liquid assets. Many investors buy illiquid assets for partially borrowed money

4 Own calculations based on returns from Bloomberg and Datastream. The return on equity is a weighted return on S&P 500, EURO STOXX 50, TOPIX and US Credit; the interest rate return is a weighted return on 10-year US and European bonds (constant maturity index); the inflation factor is a weighted return on oil, gold and other industrial metals. All returns are in Danish kroner (DKK). Volatilities are estimated on a rolling window of realised returns, and these estimates are used in the monthly rebalancing of the portfolio to the desired risk profile.

5 SR is a portfolio performance measure developed by William Sharpe. SR measures a portfolio's average excess return (return less the risk-free rate) to its volatility.

(leverage), and they depend on the loans not being terminated. Times of crisis with tightening credit markets, may therefore result in a series of forced sales among investors of illiquid assets and consequent large drops in market value. These drops are typically far larger than in the liquid markets, and far larger than would be expected on the basis of the normal, stable returns on the assets.

An investor in illiquid assets gives up investment flexibility. In contrast to an investment in, for example, listed equities, an illiquid investment is locked for a potentially long period after the initial commitment. In this period, the investor cannot sell or otherwise change the investment, and this lack of flexibility (or optionality) entails shadow costs for the investor. Francis Longstaff has analysed illiquidity in a series of papers and shown that the lost flexibility corresponds to the investor having sold an option; see, for example, Longstaff (1995) for an accessible and illustrative example of the option approach to illiquidity. Based on this insight, ATP represents the illiquidity risk of an asset as the risk of a sold put option, and this additional risk is attributed to the fourth factor.[6] For example, the risk on private equity is modelled as a combination of equity risk and risk from a sold put option (on the stock). In this way the loss is 'magnified', as a drop in stock prices generates both a direct loss and a loss from the sold put option, since its value increases. In stylised terms, this modelling corresponds to the market dynamics with 'aggravated losses' on illiquid assets outlined above.

Alternative liquid risk premia

Besides the three liquid market factors, the financial markets at large offer additional systematic and academically well-established sources of return, which can be harvested by active strategies. These may be, for example, curve positions, in which bonds with different maturities are traded against each other, currency positions or market-neutral equity strategies. Collectively, these strategies are known as *alternative risk premia* (ARP) strategies, and they are used by a growing number of investors worldwide. The strategies offer returns that are uncorrelated with the three liquid market factors.

Antti Ilmanen provides an extensive overview of ARP strategies in his commendable book Ilmanen (2011). One strategy could, for example, consist in buying low-risk equities and shorting the general equity market. There is both empirical and theoretical evidence supporting that such a strategy yields a systematic return over time, due to different investor preferences and restrictions, among other factors. Note that the strategy yields neither a gain nor a loss as a result of general market movements.

ARP strategies are typically implemented in highly liquid markets, but in terms of risk they are closely related to illiquid investments. They are complex and require specia-

CHART 4: ATP's long-term reference for allocation of risk in the investment portfolio (Balanced Beta)

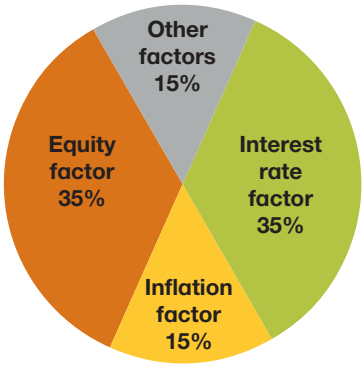

6 The option approach provides a rigorous framework for pricing of illiquidity. However, it is not always apparent which option and underlying asset(s) best represent the illiquidity risk of a specific asset.

CHART 5: Illustration of the factor exposure of typical liquid and illiquid asset classes

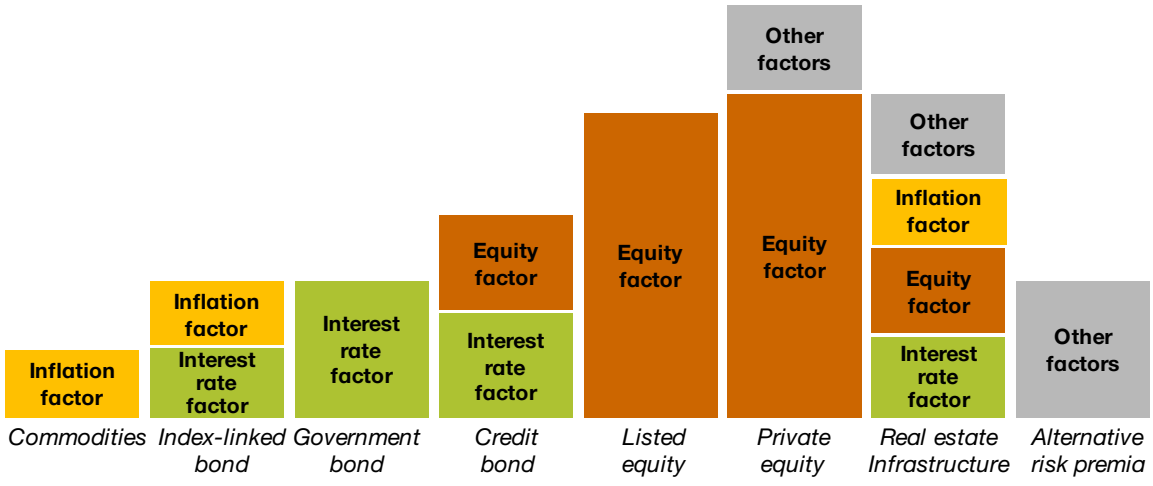

lised investment skills to implement, they tie up capital for a period of time, and they often rely on leverage. Due to these common traits, the risk profile of ARP strategies resembles illiquidity risk. For the same reason, the risk from both illiquid investments and ARP strategies is attributed to the fourth factor.

The total investment portfolio

The long-term reference for ATP's total investment portfolio is defined in terms of the distribution of risk to the equity factor, interest rate factor, inflation factor and other factors (the fourth factor). In accordance with the risk parity mind-set, an equal amount of risk is allocated to the two primary market factors (equity and interest rate). These two factors together account for 70% of the risk, while the remaining 30% is allocated equally between the inflation factor and other factors, see Chart 4. Almost a third of the risk is thus allocated to protection against bad states ('inflation') and to sources of risk other than the three fundamental market factors. The investment portfolio is also known as the factor portfolio or the *Balanced Beta* portfolio due to the balanced factor exposure.

As mentioned in the introduction, ATP completed the adoption of factor investing in 2015. However, the balanced approach to investing with focus on risk rather than capital has characterised ATP's investment philosophy for more than a decade. ATP has long been a balanced, multi-*asset* investor with a broad exposure to many asset classes, so in this light, the recent transition to being a balanced, multi-*factor* investor merely represents a refinement and systematisation of existing ideas and principles.

With the long-term risk allocation set, the next step is to find assets with the desired aggregate exposure to the four main factors. Certain assets provide pure exposure to one of the factors, primarily listed equities and government bonds, but most assets are exposed to two or more of the factors. Chart 5 illustrates the factor exposure of typical asset classes.

Illiquid assets, such as real estate and infrastructure investments, play a special role in the total portfolio. A priori, they fit in well, as they often have a balanced factor exposure. Commercial real estate is, for example, exposed to both the interest rate and inflation factors via rental income and to the equity factor, as a property's value depends on the general level of economic activity. This is illustrated in Chart 5, where only real estate and infrastructure are exposed to all the factors. In 2016, approx. 20% of the equity risk came from assets other than listed equity and private equity.[7]

There are persuasive diversification and return arguments for having significant exposure to illiquid assets. However, illiquid assets tie up capital and reduce investment flexibility, including the possibility of adjusting the investment risk in the event of loss of bonus potential. This however is not the largest risk associated with illiquid assets.

As previously described, ATP's business model is based on the investment portfolio being able to borrow funds from the hedging portfolio, enabling investments exceeding the bonus potential. However, the hedging portfolio's lending capacity is highly sensitive to the interest rate level. If interest rates increase, the value of swap contracts declines, and ATP has to post collateral. A potential problem can arise if illiquid assets comprise a very large portion of total assets, as illiquid assets cannot serve as collateral. The change of the discount curve in 2013 considerably reduced the interest rate sensitivity of the guarantees and hence of the hedging portfolio, and it reduced the potential liquidity problem in the event of an interest rate increase correspondingly.

Implementation in practice

Each of ATP's investment assets is analysed and its risk is decomposed into contributions from the four factors. For example, the entire risk of a future on the US equity index S&P 500 is attributed to the equity factor, and the same is the case for a credit default swap (CDS) on a basket of large US companies. Similarly, the entire risk from a German government bond is attributed to the interest rate factor, while the risk from a commercial property is composed of risks from all four factors: inflation and interest rate risk from rental income, equity risk from the value of the property as well as illiquidity risk, see Fact box 2 for more details.

The total factor exposure and the exposure to the individual factors are calculated daily and serve as a central investment and risk management tool. The long-term reference charts the course, but in addition there are a number of limits on currency risk, risk from ARP strategies, concentration risk on individual companies, regional risk, short and medium-term illiquidity profiles etc. Together, these limits ensure a well-diversified portfolio in many different dimensions.

The factor approach provides a common framework for evaluation of assets. Returns from different assets can be compared via their exposure to the common set of factors. By specifying a required rate of return for each factor individually, an asset's factor profile can be converted to a required rate of return on the given asset. The common yardstick allows comparison of individual assets within and across asset classes and it allows the total investment portfolio to be evaluated against the factor benchmark, i.e. the long-term reference. Using the same methodology at all levels of the portfolio ensures uniform and consistent comparisons.

The factor approach offers many advantages regarding consistency and overview, but it also entails operational challenges. The factor profile of all positions must be determined, and methods must be developed for determining the best possible factor representation. As there is not a unique

7 The average risk allocation of the investment portfolio in 2016 is shown in ATP (2016) p. 33. The equity factor accounted for 49% of the risk, with 40 percentage points coming from listed and private equity, and 9 percentage points coming from other assets, e.g. real estate and infrastructure.

decomposition of a given asset, it is also necessary to continuously evaluate the chosen methodology and adjust it if changes in the asset's factor profile are detected. The factors are intended to provide an overview, but they also introduce an extra dimension and more complex reporting.

FACT BOX 2: Factor exposure of real estate

The risk from real estate investment entails exposure to all four risk factors. The (expected) future rental income entails exposure to both the interest rate factor and the inflation factor. For example, the interest rate sensitivity of a property with a market value of €100m may be €10m, and the inflation sensitivity may be €5m. In practise, this is represented as exposure to nominal swap rates and break-even inflation (BEI), respectively, in the proper currency.

The part of the property's value which cannot be attributed to future rental income is perceived as equity exposure. For example, if the value of future rental income is €80m, the remaining €20m is treated as equity exposure. In addition, there is illiquidity risk corresponding to a sold put option with the same principal as the equity exposure.

The decomposition of risk to exposures to the four factors is model-based. The methodology is evaluated regularly, e.g. by backtesting the returns of the factor representation against the historical returns on ATP's real estate portfolio.

Each factor has an associated risk premium, and the expected return on an asset can therefore be calculated from its factor representation. Risk premia are estimated based on historical returns on the given factor representation: Interest rate and equity risk both have a Sharpe Ratio (SR) of approx. 30%, the SR for inflation risk is close to 0%, while the SR for illiquidity risk is approx. 70%. The 'factor return' is used, among other things, for screening of investment assets. To be of interest, potential investment assets must deliver an (expected) return at or above the return warranted by their factor representation.

The fairly large risk premium on illiquidity implies that an illiquid asset must provide a significant excess return compared to a similar liquid assets. This excess return is compensation for the investment restrictions imposed on ATP by illiquidity. Conversely, the risk premium on (break-even) inflation is very small. This reflects that inflation exposure is not expected to yield a direct return, but rather to provide diversification and protection in bad states.

Moreover, it is our experience that it takes both time and effort to establish a common understanding of the factor universe in the investment organisation.

Conclusion

In this paper, we have described the key aspects of ATP's implementation of factor investing. As seen, factor investing is more of a mind-set than a specific investment strategy. Factor investing promotes focus on the underlying common factors to which assets are exposed, and it incites investors to identify and take account of their bad states in the allocation. The number of factors should be kept low, and they should ideally have an intuitive interpretation.

Factor investing also encourages investors to identify how they differ from other investors and to take these differences into account. For example, ATP differs from many other institutional investors in that the primary investment restriction is risk rather than capital. This influences, among other things, ATP's appetite and ability to hold illiquid assets.

Finally, it should be mentioned that the factor approach has also facilitated simple and streamlined communication about ATP's investment strategy. The factors and the long-term reference are easy to comprehend for both internal and external stakeholders, promoting dialogue on purpose and direction.

References

— Ang, Andrew, 2014: *Asset Management: A Systematic Approach to Factor Investing*. Oxford University Press.

— Asness, Clifford S., Andrea Frazzini and Lasse H. Pedersen, 2012: Leverage aversion and risk parity. *Financial Analysts Journal*, 68, pp. 47-59.

— ATP, 2016: The ATP Group – Annual Report 2016.

— Dalio, Ray, Bob Prince and Greg Jensen, 2015: Our thoughts about risk parity and all weather. *Bridgewater Daily Observations*.

— Ilmanen, Antti, 2011: *Expected Returns: An Investor's Guide to Harvesting Market Rewards*. Wiley.

— Jepsen, Henrik Gade, 2006: Bør pensionsselskaber anlægge et nyt verdenssyn? (Eng. 'Should pension funds adopt a new world view?') *Finans/Invest* 2/06, pp. 10-13.

— Longstaff, Francis A., 1995: How much can marketability affect security values? *Journal of Finance*, 50, pp. 1767-1774.

— Merton, Robert C., 1992: *Continuous-Time Finance*. Wiley-Blackwell.