Liability Driven Investment

[Nematrian website page: LDI, © Nematrian 2015]

Abstract

Ultimately if an investor had no liabilities or other similar requirement to build up an asset pool there would be no need for investment management. This is the key observation underlying the trend towards *liability driven investment* ('LDI'). In essence, LDI involves understanding the nature of the investor's liabilities and designing an appropriate investment strategy to suit these liabilities. Usually, investors seek a balance between risk and reward, so this does not necessarily result in adopting an investment strategy with economic sensitivities that closely match the sensitivities inherent in the liabilities. Instead, investors may deliberately seek to mis-match relative to their liabilities, because they expect to achieve higher returns by doing so.

The aim of these pages is to explore this topic further. They are based in part on material in an Appendix to Kemp (2005).

Contents

1. Introduction

[LDI1]

Material covered

[LDI1a]

1.1 There is growing interest in the concept of *liability driven investment* ('LDI') for many types of institutional investor across many different locations. Liability driven investment is commonly adopted by life insurance companies (although they may not use this term) for liabilities which contain guarantees or underpins, and by defined benefit ('DB') pension schemes. Indeed, it can be argued that *every* institutional investor to some extent bears in mind its liabilities when formulating its investment strategy – few investors have *no* liabilities at all which they wish to honour!

1.2 In these pages we will mainly focus on LDI as applied to DB pension funds (and to a lesser extent insurance companies), although we will also from time to time refer to LDI applied to other

types of institution. Large mature DB pension schemes with a greater bond focus typically seem to be more interested in this type of investing than less mature more equity focused clients.

1.3 There are several different ways in which a liability driven investment portfolio might be structured. Perhaps the simplest involves two or three parts:

(a) A *return seeking* portfolio (typically actively managed), invested in those asset classes and vehicles the investor believes have the highest likelihood of outperforming without carrying undue risk.

(b) A *protection* portfolio involving investment in physical securities, often bonds, chosen in broad terms to have economic characteristics that mimic those of the relevant liabilities. For example, if the liabilities are partly fixed in monetary terms and partly linked to inflation, i.e. to movements in a Consumer Price Index (CPI), then it might incorporate some fixed interest and some inflation-linked bonds.

(c) A *swaps overlay* (or more generally, a *derivatives overlay*) portfolio. This would typically consist of one or more swap contracts (or other similar derivatives) that involve the pension fund giving up one set of future cash flows (e.g. ones like those arising from the portfolios in (a) or (b)) and receiving in return another set of future cash flows (e.g. ones more closely matching the relevant liabilities). Precisely how these swaps might be structured can vary. For example, there might be one swap that pays away to the swap counterparty cash flows akin to those arising from the portfolio in (a) or (b) in return for interest payments on some notional principal linked to prevailing cash rates (e.g. LIBOR). There might then be a second swap that paid away this LIBOR cash flow in return for a cash flow that more closely matched the pension fund's expected liability outgo. Or there might be several swaps on each side that handled different parts of the cash flow (e.g. differentiating by term or by liability type). Or, all of the cash flows might be wrapped up in a single overarching swap.

1.4 Usually, investors seek a balance between risk and reward, hence the existence of the return seeking element as per (a). The relative size of the return seeking portfolio versus the protection portfolio/swaps overlay (or to be more precise the relative sizes of the exposures within them) will depend on where within the spectrum of possible alternatives the investor wishes to pitch this balance. Investors may, for example, believe that there is an intrinsic likelihood of long-term outperformance associated with a particular type of asset, e.g. equities (e.g. because a capitalist economy 'ought' over the longer term to reward entrepreneurs and hence to reward equity holders more than those who bear less risk of loss from business failure). However, investors may not be willing (or may not be allowed the latitude under regulation) to invest their entire asset base in return seeking assets, even in a diversified portfolio of such assets, because of the risk that the asset performance might be worse than expected. If the investor is subject to mark-to-market regulatory principles (and Kemp (2009) argues that there are good reasons why governments should want most investors to be subject to such rules, if suitably defined, even if this is not always the current position) then 'performance' here may not relate merely to the behaviour of the assets between purchase and maturity but also to their market value movements in the meantime. Investors not interested at all in LDI can be viewed as a special case of the above, where their entire portfolio is held in (a).

Similarities with other techniques

[LDI1b]

1.5 The concepts behind the protection portfolio and the swaps overlay are similar to the actuarial theory of matching. Indeed, if the liabilities are short enough and the trustees want a passively managed low risk approach then 1.3(c) might become superfluous and 1.3(b) might be merely involve a more traditional cash flow matched portfolio using, say, government bonds.

1.6 If the current incarnation of LDI can be said to be novel relative to what has gone before, the core 'new' idea is the use of swaps or other similar derivatives, within 1.3(c). They are used because the liabilities are typically of too long duration (or, in the case of an insurer hedging guarantees, too complex) to be matched merely using physical bonds. So, we need a 'synthetic' method of artificially lengthening the duration of the assets if, for example, we do not want to be exposed to the risk that very long dated yields will fall more than we expect. This is possible using interest rate swaps, particularly long-dated ones. Using swaps also gives the investor a wider range of underlying bonds in which it can invest; indeed the swaps be overlaid on top of the return seeking portfolio as well as the (physical) protection portfolio, potentially allowing investors to invest a higher proportion of their overall portfolio in return-seeking assets than would otherwise be the case (from which they would presumably expect to receive a suitable reward).

1.7 The same overall concept is still applicable for the swaps overlay component if the liabilities are CPI linked (or contain other more complex inflation-linked characteristics such as Limited Price Indexation, 'LPI') instead of merely involving long dated fixed monetary amounts. Such characteristics are common amongst the liabilities of many types of defined benefit pension scheme. The only real difference of substance is that the cash flows that the swaps pay to the pension fund need to include these features if they are to hedge against such risks, i.e. they need to involve the investment banks selling 'inflation' (here used as a shorthand for cash flows with characteristics sensitive to future inflation rates) to the pension fund. Of course, swap counterparties (typically banks) will typically want to hedge their exposures. So they will be on the lookout for other market participants (e.g. utility companies or infrastructure projects) prepared to sell them inflation. The two sides do not need to be in identical form (e.g. one might strictly increase in line with the RPI, the other might be more LPI in nature). The 'art' of good derivatives intermediation is to be able to access both sides of the flow, make a good return between the two and keep the inevitable residual mismatches well controlled and hedged (and to charge an appropriate spread for carrying this risk).

1.8 We referred in Section 1.4 to a balance between risk and reward. Investors may also have views on the extent to which different types of risk are more or less 'expensive' to hedge, where here 'expensive' means extent of potential reward foregone. For example, they may believe that 'extraneous' currency exposures within a typical asset portfolio will typically be less well rewarded over the longer term than a bias towards equities (e.g. because they believe that a successful capitalist economy should reward the latter but there is no particular reason to believe that it will reward the former). This may influence the extent to which particular risks are hedged and therefore on the relative sizes of different types of risk exposures catered for in the protection portfolio or swaps overlay.

2. The 'core' element of such a structure

[LDI2]

2.1 In some respects LDI can be thought of as a variation of the core-satellite investment paradigm that was common some years ago. This involved subdividing the portfolio into a low-cost (often passively managed) *core* element that provided the bulk of the exposures the investor wanted, together with a higher-cost, actively managed *satellite* element that would hopefully deliver better returns (sufficiently better to justify the extra expenses).

2.2 An important advantage of using an overlay element structure within the structure is that it divorces the managing of the 'core' (physical) asset base from the 'bespoke-ness' needed to achieve a close match to the investor's own liabilities. The 'core' can then be managed in a practical manner, e.g. along the lines of a manager's standardised investment process against some relatively standard benchmark, offering potential economies of scale.

2.3 The precise structure of the core element can still express investor preferences, but these preferences can now primarily refer to the assets in isolation, rather having simultaneously also to cater for the precise shape of the liabilities. For example, the core element might eschew gilts in favour of a greater proportion of less well-rated credits. This might be because the yield spread of such bonds over gilts is believed by the trustees to over-compensate the holder for the likely future default loss experience on such bonds on the grounds of liquidity criteria, but see Kemp (2009).

2.4 It also ceases to be necessary for the core component to be exclusively bond orientated, even merely the protection element of the 'core'. Instead, the core could make use of *portable alpha*. Nowadays swaps come in a very wide variety of forms. It is now possible to swap almost any sort of return stream, property-like, equity-like, bond-like, cash-like or inflation-like, into any other sort of return stream, embedding into the swap, if we so wished, caps, floors and other option-like characteristics. So, if an investor has confidence in a given active manager's skill at adding value then this skill can be in any asset class we like with the added value *ported* onto a liability orientated benchmark merely by swapping the return on the relevant active manager's benchmark into the return on the benchmark set by reference to the liabilities.

2.5 But whether such refinements are likely to be appreciated by most sets of pension fund trustees is less clear. A few asset managers do offer portable alpha products, but take-up to date has been relatively limited, perhaps because of the difficulties involved in educating trustees in the concepts involved (or in being sure that there is no leakage of value by the porting process). Also, one can argue that the swap contracts might be more keenly priced if they are swapping similar sorts of return streams. So, all other things being equal, if our desired cash flows are akin to fixed or inflation-linked bonds (just rather longer than is easily available in the physical market place) then starting with similar sorts of cash flows may be preferable.

3. The swaps element of such a structure

[LDI3]

Clarification of who does what

[LDI3a]

3.1 Divorcing the core physical portfolio from the derivatives overlay helps to clarify who is responsible for what decisions. The following parties are involved and would typically have the following responsibilities if LDI is being applied to a UK defined benefit pension scheme:

(a) *Trustees*: Carry ultimate legal responsibility for the fund. They would be responsible for choosing who manages the core element and the swaps overlay. In the above structure, they would also be responsible for instructing the investment manager when to execute

exactly what swap transaction (although in practice there would have been prior liaison with the investment manager in choosing how best to frame these instructions).

(b) *Scheme actuary*: Would normally prepare any required liability cash flow projections, and update them as necessary at regular intervals. See below for what such projections might contain.

(c) *Investment consultant*: Would normally advise the trustees on overall investment strategy, on fund manager selection and on how to monitor the fund manager and measure the manager's performance. Together with the actuary would advise on exactly what liabilities to match (e.g. should it include pensions in payment, deferred pensions and/or actives' liabilities?).

(d) *Fund manager*: Likely to be responsible for managing the underlying bond portfolio and for actual implementation of the swap transactions. The role in relation to the swaps overlay could perhaps best be classified as 'execution only' in the sense that the fund manager would probably help draft up any instructions formally given to it by the trustees and/or investment consultant, but otherwise the swap portfolio would be 'non-discretionary'. This would be in contrast to the core physical portfolio (which would most typically involve discretionary active management). The fund manager would most likely provide education to the trustees, views on transaction timing and valuations of the individual swaps. The fund manager would also most likely arrange for the collateralisation of the swap portfolios.

(e) *Investment bank*: Would be the trustees' actual swap counterparty, i.e. the entity whose balance sheet would honour the contractual obligations in any given swap transaction. In principle, trustees (or their consultants) could deal directly with such banks (subject to any overriding requirement on the trustees to avoid 'day-to-day' investment activity if they are not FSA regulated). But in practice, banks' derivatives desks are remunerated on a transaction-orientated basis. This is not obviously conducive to acting in the best interests of the trustees. It is most likely that the trustees would delegate choice of swap counterparty to their fund manager, who would make the choice by reference to the usual sorts of 'Best Execution' criteria that apply to fund manager dealing activity (subject to any overriding criteria set by the trustees such as a credit rating requirement). There could be several such banks, as the fund manager in principle needs to apply Best Execution criteria each time new swap transactions take place.

Practical considerations

[LDI3b]

3.2 In practice, there is likely to be close liaison between the actuary/investment consultant and the fund manager when preparing suitable liability projections and hence a proposed structure. The fund manager might also typically work with a few well-chosen investment banks who can help to identify what derivatives are most likely to meet the client's requirements.

3.3 There needs to be such interaction because overly exact cash flow matching might result in an overly complex (and therefore expensive) structure, bearing in mind the inherent approximations involved in liability projections (and the inherent approximations involved in modelling how the actively managed core portfolio might behave). There are also minimum amounts below which it is impractical to effect swap contracts, which depend in part on how non-standard the swap is. An exact hedge of all of the risks embedded in the liabilities may be prohibitive or even impossible (e.g.

liability driven 'investment' has rarely to date attempted to include scheme-specific longevity protection). Experience suggests that complicated overlay structures may initially be discussed with trustees and their consultants, but typically only relatively simple structures seem to be used in practice.

3.4 At regular intervals (say yearly) the client (in conjunction with its actuary/investment consultant) would probably revise its cash flow projections and, after discussion with the fund manager, would instruct the fund manager to alter the structure of the swaps within the swap portfolio. Again this would be done subject to the usual Best Execution rules, perhaps if necessary novating or cancelling previous swap transactions with new ones (to avoid building up large numbers of swap transactions that largely cancel each other out and which might be burdensome to administer).

3.5 This flurry of activity at outset contrasts with what happens the rest of the time. The fund manager does incur some ongoing costs, most notably the costs of sorting out the collateralisation of the swaps, as well as ongoing reporting/valuation. These costs are typically smaller than the costs of actively managing a portfolio, and might be absorbed within an all-in fee covering both arrangements. It would be possible for the fund manager of the swaps overlay to be different to the fund manager of the underlying physical bonds (just as a scheme's tactical asset allocation manager does not need to manage any of the underlying assets). However this may make collateralisation procedures more complicated.

4. Mitigating credit risk within swap contracts using collateralisation

[LDI4]

4.1 Normally the pension scheme would want the swap counterparty to collateralise the swap contract. The aim is to reduce the exposure that the pension fund has to the risk of default of the bank involved. The aim is to have moved some suitable form of collateral from the bank to the pension fund whenever such a default might be costly to the pension fund. This involves marking to market the swap (by definition this is the estimated cost of effecting a similar sort of swap with another counterparty), and whenever this builds up to be materially positive as far as the pension fund is concerned, for additional collateral to be 'posted' by the bank to the fund. If the mark to market then declines, some of the collateral would be released and returned back to the counterparty.

4.2 The counterparty might of course also require the swap to be collateralised for the same reason but in reverse. Over the last few years, many life insurers entering into over-the-counter derivative transactions have discovered that they may be deemed less credit-worthy than their counterparties, although this may have become less true of late. Underfunded pension funds may face the same learning curve!

4.3 For most transactions of any size, it is now common for collateral flows to occur quite frequently, even daily (although there will typically be minimum thresholds and a minimum build-up of exposure, typically dependent on credit rating, before any flow occurs). It may be possible to pledge securities held within the underlying portfolio. Or, it may be necessary to hold some cash buffer within the swap portfolio itself to meet such calls. If instead the bank is posting collateral to the scheme then it too needs looking after, since it may need to be returned at some stage.

4.4 Typically, the asset manager would negotiate collateralisation arrangements on behalf of its client via a Credit Support Annexe within its wider negotiation of the master International Swap

Dealers Association (ISDA) legal documentation governing the overall relationship between the client and its bank counterparty. Normally the client would legally be one of the two parties to swap, with the asset manager merely acting as its agent. The pension fund might therefore want its own lawyers to review or negotiate these contracts. But in practice, the investment manager is likely to have greater negotiating clout with the bank, given other relationships it may have. The investment manager may therefore adopt umbrella documentation relating to all of its clients that wish to transact with the relevant counterparty. Where the client has multiple swap transactions with the same counterparty it is normal to have them all netted off within the relevant ISDA and Credit Support Annexe. Otherwise one party can find that in the event of the other party defaulting it owes money to the defaulted party on one transaction but cannot recover what it is owed on another.

5. Monitoring such a structure

[LDI5]

Main elements

[LDI5a]

5.1 There are three key elements to the above structure that might need monitoring (other than the usual monitoring that would be carried out even if no LDI approach was being adopted on the return seeking portfolio):

(a) *The (actively managed) underlying bond portfolio*. This would be assessed as usual for the asset management product in question. For example, if it involved management of a credit portfolio against a market index then performance and risk measurement and attribution analyses versus the benchmark in question might be reported as per the asset manager's/pension fund's usual reporting cycle.

(b) *The (passive) swaps overlay*. This might for simplicity also be reported upon to a similar frequency, although most attention would be focused on those occasions when the swap positions needed to be altered.

(c) *The effectiveness of the choice of swaps overlay structure in relation to the scheme's liabilities*. Various approximations will have been interposed between the precise liability model available from the actuary and the precise structure of the swap portfolio. The swap portfolio being 'execution-only' in nature, this element of the decision-making is actually one that lies with the trustees, albeit only after taking advice from other parties.

5.2 The key additional requirement is to construct some sort of liability benchmark (or index) that reflects in a market-orientated way the nature of the liabilities. Constructing such a benchmark may also directly guide the choice of swaps to hold within the overlay portfolio.

5.3 The most obvious way to proceed is first to develop some cash flow projections, differentiating between ones with different sorts of economic sensitivities (particularly those where the sensitivities have option-like characteristics, such as LPI). For example, the liability flows might be differentiated by year of projected payment into those that involve:

(a) Fixed monetary sums, e.g. those arising from benefits not subject to any increases.

(b) Fully RPI inflation-linked sums, e.g. benefits subject to full RPI linked increases.

(c) Sums that increase on a year by year basis on some more complicated measure driven by inflation at that time, e.g. LPI-type increases in payment. For these sorts of liabilities, the expected outgo during a given future year can still be derived from a single expected amount at outset, together with the history of RPI increases since then. If different ceilings, say 2.5% and 5% pa caps, apply then these flows should in principle be differentiated, as swaps to match them exactly would also differ.

(d) Cash flows governed by more complex increase formulae dependent on multi-year investment or economic conditions. At least in principle, benefits linked to LPI in deferment fit into this category. The big difference between these sorts of cash flows and the sorts referred to in (b) or (c) are that they in principle require multi-dimensional matrices to specify as they depend jointly on date of withdrawal, assumed date of retirement, assumed date of payment and (for those already deferred pensions at outset) on how large RPI increases were prior to the start of the projection relative to the caps and floors present in individual members' benefits. As with (b) and (c) they also depend on RPI increases post the start date of the projection.

Swap portfolio structure

[LDI5b]

5.4 The choice of numeraire (e.g. whether the cash flows are in nominal or real terms, or if they are expressed using some present value metric) is not particularly important as long as the cash flow analysis ultimately precisely specifies the assumed cash flows. This explains which cash flows as per 5.3(d) are more problematic – they require lots more detail to specify precisely. It may be possible to develop suitable approximations that simplify them into a form that was more easily specifiable. It might also in practice be possible to simplify away liabilities of the form described in 5.3(c). It is also worth noting that the cash flows are not deterministic in nature. If the numbers of members involved is quite small then the random incidence of individual deaths will introduce uncertainty. For more sizeable schemes, the unpredictable nature of future changes in general levels of longevity is likely to be more significant (as is whether the mortality table in question is suitable for the actual type of individuals represented by the scheme membership).

5.5 Once the liabilities have been expressed in a suitably simplified form it becomes possible to structure swaps that capture the main characteristics of these cash flows. Liabilities that are fixed in nominal terms would be matched using swaps that generate fixed cash flows whilst those that are RPI-linked would utilise inflation swaps. LPI-linked liabilities can be catered for in a similar fashion although often their costs seem high to clients. This seems to be because clients worry less than the market as a whole does about the possibility of inflation becoming negative.

5.6 Performance (and risk) measurement and attribution of the swaps portfolio can then also be carried out by reference to the simplified cash flows, discounted (probably) at swap rates, versus mark to market movements in the value of the swaps.

5.7 There is a link between liability driven investment and fair valuation principles. The actuary will typically have placed some value on the liability cash flows. Assuming that the liability cash flow projections are truly correct (and ignoring some of the niceties surrounding credit risk on cash

deposits etc.), we might ask how we can tell if this sum would actually be sufficient to provide all of the projected cash flows. This depends on whether the actuary's valuation is bigger or smaller than the fair value of the liabilities derivable from the mark to market value of the swaps. It is not sufficient merely to compare the return on the liability driven portfolio with the movement in value placed on these liabilities by the actuary. The movement needs to be unbundled into its various parts, including potentially a part relating to the difference between the fair valuation and the actuary's valuation.

5.8 Even the above analysis involves simplifications. For example, there is an implicit assumption in the above that the fund's mortality experience can be well predicted at outset. But merely differentiating between nominal, real and LPI-linked increases provides no protection against unexpected improvements in mortality. There may be future discretionary benefit improvements. Active members' liabilities are particularly difficult to project reliably in this context given their sensitivity to uncertain future member-specific salary increases. For a full picture one would in principle differentiate between each such risk, as per Section 4. In practice this is likely to be challenging, although at least thinking about such matters may help to highlight what sorts of risks a liability driven investment portfolio does or does not hedge against.

6. Alternative approaches

[LDI6]

6.1 The overlay approach described in Section 5 clearly demarcates who is responsible for what. But trustees might prefer merely to set their investment manager a liability driven benchmark akin to the one described above, and say "get on with it", with the investment manager free to use whatever instruments it likes (including swaps and other derivatives) and whenever it likes, to match the liabilities or preferably to add value versus them.

6.2 Key requirements for such an approach are for the trustees and their consultants to carefully craft an appropriate liability driven benchmark as above, for the fund manager to have good systems for measuring at all times how far its portfolio deviates from this benchmark and for it to be very clear exactly what is expected of the fund manager. The bespoke nature of such a service is likely to make it practical only for larger accounts. It is worth noting that if the fund manager cannot practically hedge a particular part of the liability benchmark then there will be a 'random' element to his performance. The fund manager may stress this whenever he thinks it has worked to his disadvantage, and the trustees may do the opposite whenever they think it has worked in the fund manager's favour. Unfortunately, there is almost certain to be disagreement about which is the case, unless the whole arrangement is very carefully managed. An advantage of the swaps overlay approach described above is that it airs and manages these potential disagreements at outset, via the discussions needed around the formulation of the swaps overlay.

6.3 The trustees may deliberately want to adopt a strategy that deviates from the most precise liability driven benchmark. In these circumstances, a clear liability driven benchmark might still be defined but then deliberately modified to focus on what the trustees want.

6.4 For example, the trustees may feel that banks might be quoting excessive prices for buying cash flows that embed option-like inflation characteristics such as those implicit in LPI linked benefits. Yet they may still want some hedging of such risks. They might then ask the fund manager to hedge these risks in a more approximate way, using dynamic hedging, to avoid ceding this supposed profit margin to the bank. This could perhaps most easily be achieved by giving the investment manager a benchmark that changes in a dynamic fashion as the underlying economic

parameters change. The aim would be to mimic the economic sensitivity of the fair value of the option-like characteristics insofar as far as these depend on the parameters in question. A perfect hedging algorithm, were one to exist, would of course also depend on volatility, which would require the use of more complicated derivatives (but this would then defeat the point of seeking to avoid the use of such derivatives because they are believed to offer poor value-for-money).

6.5 Some modification to the swaps overlay approach may be needed for smaller schemes. A single swap might be easier to have 'segregated' in this context than a whole bond portfolio, but there are still implicit lower limits on the sizes at which they become practical. A better alternative may be to create specially tailored long duration pooled bond funds. Several investment managers appear to be designing such products. In real life, a portfolio of pension liabilities typically gets shorter over time, so any pooled approach is unlikely to match any particular scheme's liabilities as well as a more bespoke approach.

7. Other comments

[LDI7]

7.1 Liability Driven Investment is closely allied with Asset Liability Modelling (ALM, aka Asset Liability Management). ALM can be thought of as involving:

(a) Modelling of how assets and liabilities might interact, and then

(b) Managing the assets, liabilities, or both, so that the nature of this interaction is favourable to the entity.

Typically in both (a) and (b) there is some trade-off between risk and reward, suitably defined, for the sorts of reasons highlighted in 1.4.

7.2 Interestingly, what might be described as 'state-of-the-art' ALM seems to differ between insurance/pensions and banking/investment management:

In ***insurance/pensions*** the focus is principally on:

(a) Projecting assets and liabilities into the future, and from them also deriving potential future behaviour of related features, e.g. solvency, free asset ratio; and

(b) The projections are usually stochastic in nature, i.e. they involve projection of results under many different scenarios.

In contrast, in ***banking/investment management*** the focus is more often on:

(c) Identifying the extent of the current mismatch between assets and liabilities; and

(d) Quantifying this mismatch via Value-at-Risk, tracking error or other similar risk metrics.

7.3 This apparent difference in focus is explored further in Kemp (2005b) and Kemp (2009). Part of it is linked to implicit or explicit assumptions about how rapidly it might be possible for the firms/entities to take remedial action. The apparently different focus also in part reflects presentational preferences within the relevant client types.

References

[LDIRefs]

Kemp, M.H.D. (2005). Risk Management in a Fair Valuation World. British Actuarial Journal, 11, No 4, pp. 595-712

Kemp, M.H.D. (2009). *Market consistency: Model calibration in imperfect markets*. John Wiley & Sons [for futher information on this book please see Market Consistency]