Turnbull Consulting & Analytics Ltd.

A Valuation, Capital and Matching Adjustment Methodology for Equity Release Mortgages and their Securitisations

Craig Turnbull FIA

6th July 2020

Table of Contents

1 Introduction

The Purpose of this Paper

This paper is intended as an actuarial practitioner's guide to the implementation of valuation and capital methods for Equity Release Mortgage (ERM) portfolios and their securitisations. The paper will not attempt to offer new scholarly research that fundamentally advances theoretical knowledge about the nature and behaviour of this complex asset class. Nor will it offer any empirical analysis that may be useful in calibrating quantitative assumptions about the various important variables that impact on ERM valuations and the assessments of their capital requirements. A number of other papers have been published over the last couple of years that have advanced our theoretical understanding of the asset class and that offer useful empirical analysis[1]. To explain the purpose of this paper and how it fits into current actuarial thought on the treatment of ERMs, it may be useful to offer some very brief observations on current actuarial practices in this field.

In recent years, UK actuaries have made significant and commendable developments in their analysis and modelling of this complex illiquid asset class. Insurance firms have built sophisticated stochastic modelling frameworks for the valuation and capital assessment of ERM portfolios and their securitisations. Contemporary ERM valuation methods may include several stochastic risk factors and make allowance for complex and important features of the asset class such as the dynamic prepayment behaviour of borrowers. Solvency II Internal Models have been extended to include the risk calibrations that are key to ERM capital modelling. These developments mean that many UK insurance firms have modelling capabilities that can provide robust estimates of ERM portfolio values and capital requirements.

Insurance firms have also developed quite complicated ways of structuring their exposure to the ERM asset class through the use of securitisation vehicles. This re-packaging of ERM asset portfolios has been motivated by Solvency II's Matching Adjustment. The *consistent* valuation and assessment of capital requirements for the underlying mortgage portfolio *and* the securitisations is a more demanding technical task than 'merely' working with the underlying mortgages. Current valuation models have tended not to have been built with a primary focus on consistent valuation of mortgages and securitisations. As a result, various *ad hoc* or *ex post* valuation adjustments may be required to make the valuation of securitisations 'work' - for example, to ensure that the values of the securitisation tranches sum to the value of the underlying mortgage portfolio (a property sometimes referred to as the 'equation of value'). These adjustments may not be necessary in the base valuation, but the need for such valuation adjustments can often arise in the stresses that drive capital assessments. This can undermine confidence in what the securitisation capital models are producing.

This paper argues that a consistent approach to the treatment of the mortgages and their securitisations removes the need for these *ad hoc* adjustments to the model output. To achieve this, the securitisation valuation model must be directly based on the fundamental economic relationship between the securitisation tranches and the underlying mortgage portfolio: that is, by recognising that the securitisation tranches are well-defined derivative instruments of the mortgage portfolio. This paper will show that the use of standard derivative valuation principles to value the securitisation tranches as derivatives of the mortgage portfolio can *significantly simplify* the valuation of the securitisation tranches. This simplification occurs by recognising that the derivative

[1] In particular, see Jeffery & Smith (2019), Dowd et al (2019) and Tunaru (2019) for a diverse set of treatments.

relationship between the mortgage and securitisation means that *the valuation model for the mortgages almost completely defines the valuation model for the securitisation tranches*.

The other side of this coin, however, is that we will be asking more of the mortgage valuation model than current practice does. The consequences of any shortcomings in the mortgage valuation model may be amplified when the model is used as the foundation for the valuation of securitisation tranches as well as mortgages. Current mortgage valuation methods do tend to have some shortcomings that become more material when we attempt to use them to value securitisations. So, a central theme of this paper is that *the key to improving securitisation valuation and capital methods is to make the underlying mortgage valuation model more economically sound*.

This paper therefore offers what might be thought of as a 'second generation' ERM valuation methodology. It will propose a mortgage valuation approach that is more economically logical than current mortgage valuation methods. This will make it fit for the secondary, but ultimately fundamental, task of being the foundation for securitisation valuation (and hence securitisation capital requirements under a 1-year Value-at-Risk basis).

It is the objective of this paper to show that the significant actuarial progress that has been made in recent years in the development of sophisticated ERM valuation methodologies can be further enhanced in a pragmatic way. The ERM valuation methodology proposed in this paper will admittedly be slightly more complicated to implement than current ERM valuation methodologies. But the pay-off from this investment will be that the valuation and capital assessment of the securitisation tranches will be made *much* more straightforward than the current approaches that seek to fit square securitisation valuation methods into the round holes created by mortgage valuation models. The paper will also show that the enhanced valuation methodology can provide useful insight into the treatment of ERM-based assets under the Solvency II Matching Adjustment.

The Structure of this Paper

Chapter 2 develops the economic logic that is the basis for the mortgage valuation methodology proposed in this paper. This is one of the shorter chapters in the paper, but it is perhaps the most important as everything else in the paper follows from what is developed here.

Chapter 3 provides a couple of case studies of valuation methods that implement the methodological framework offered in Chapter 2. It is important to note that the valuation logic developed in Chapter 2 can be implemented in any number of ways by means of different modelling approaches and assumptions. Chapter 3 highlights this by showing two quiet different implementation approaches – one using a Black-Scholes-based approach, and the other using a cashflow simulation model.

Chapter 4 shows how the cashflow simulation valuation method implemented in Chapter 3 can be used to derive a capital requirement for the example mortgage portfolio. The calculation process for the capital requirement will be familiar to any actuary with experience of Solvency II Internal Models. The robustness of the valuation method under stress can provide significant insight into how the mortgage valuations behave in stressed conditions.

Chapter 5 discusses how the valuation method for an underlying mortgage portfolio can be naturally and easily extended to determine valuations for all tranches of a securitisation. These securitisation valuations will be economically consistent with the underlying mortgage valuations and hence will 'automatically' satisfy the equation of value in all possible scenarios.

Chapter 6 extends the valuation analysis of Chapter 5 to derive the capital requirements of securitisation tranches in a way that parallels the capital approach taken in Chapter 4 for the mortgages. We show that the securitisation capital requirements that result from this process are consistent with the capital requirements of the underlying mortgage portfolio.

Finally, Chapter 7 discusses the application of the Matching Adjustment (MA) in the presence of MA-eligible ERM-based debt instruments. Here it is argued that the valuation method and capital method that has been implemented for the underlying mortgage portfolio provides a direct and natural basis for allowing for MA in both liability valuation and the assessment of capital requirements.

Throughout Chapters 3 to 7, it should be noted that the quantitative implementation examples are intended for the purposes of illustrating the proposed methodologies 'in action'. The specific parameter choices and modelling assumptions are therefore intended solely for the purposes of demonstrating the methodology. They are not intended to suggest 'best estimates' of parameters such as house deferment rates, asset illiquidity premia, mortality rates, dynamic prepayment behaviour, 99.5th percentile house price falls, etc. etc. This is not to suggest that the empirical estimation of these variables is by any means straightforward or unimportant. On the contrary, these tasks merit their own dedicated paper(s).

2 A Valuation Methodology for Equity Release Mortgages

The Introduction argued that the mortgage valuation methodology is the fundamental piece of economic logic from which everything else ERM-related directly follows – the capital assessment of the mortgage portfolio; the valuation of the securitisation tranches; the capital assessment of the securitisation tranches; and the Matching Adjustment treatment of the securitisations. Everything in this paper that follows this Chapter is therefore directly related to and a consequence of what is developed here.

Given the fundamental importance being attached to the mortgage valuation methodology, we had better be clear what exactly we *mean* by valuation. In the context of this report, valuation means something quite limited and specific: *valuation refers to the economic process of determining a value for an asset that has no observable price by interpolating and extrapolating from the prices of other relevant assets whose prices are observable*. Importantly, this means that the valuation process does not require as an input any judgments on the reasonableness (i.e. cheapness or dearness) of the observable prices of relevant assets. The objective here is simply to produce values that are consistent with those observable prices.

We may hold a subjective view that current observable prices are cheap or dear. That view may be based on considerable expertise and extensive empirical analysis. If our task is to determine whether we should invest in a corporate bond or an equity release mortgage or a commercial mortgage, this view is highly relevant. Similarly, if our task is to decide on a pricing basis for writing equity release mortgages, these views will again be crucial. If we would like to forecast how profitable we expect writing that business to be, these views would yet again be crucial. But these views are not relevant to the valuation process. Of course, the implementation of the valuation model may help *inform* these subjective views of the asset's cheapness and dearness (i.e. these views cannot be an *input* into the valuation prices, but they may form part of the *output*). The valuation model could, for example, suggest that ERM prices are consistent with a house price volatility that we believe is a very prudent estimate of likely future house price volatility. Or we might find that low LTV mortgage rates are consistent with a higher house price volatility assumption than high LTV mortgage rates and conclude that the current prices of low LTV mortgages are more expensive. Such observations may inform analysis of the attractiveness of the asset class or the relative attractiveness of different parts of the market. But, again, such views are outputs rather than inputs into the valuation process.

So, the fundamental objective of the valuation process is to consistently value different assets. This process must therefore involve determining which assets we wish to be consistent with. And what exactly does *consistency* mean here? In a nutshell, we will specify a model for the valuation of assets with unobservable prices. Some of the parameters of this model may be derived or 'fitted' to the observable prices of other assets. So, the *relevance* or otherwise of other assets is determined by the extent to which their observable prices help to determine any of the valuation model parameters for the valuation of the asset of interest. If the valuation model produces a value for other assets equal to their observable price, the valuation method may be considered *consistent* with those asset prices.

How do specify the valuation model that is to be calibrated to relevant observable asset prices? There are usually many choices here, and the choice usually involves some trade-offs between simplicity, performance, ease of calibration, and the range of assets that can be simultaneously valued by the model / calibration. But there are certain basic properties that all good valuation models have. Most fundamentally, we would expect two assets that produce identical cashflows in all possible states of the world to have the same value according to the model (if we assume a non-zero illiquidity premium, the liquidity state of these two assets must also be the same). It is difficult to see how a valuation model can be said to be consistent with the prices of other assets if the valuation model is not *internally* consistent – that is, if it can place different values on identical streams of cashflows that are generated by the assets we are trying to value. We might loosely refer to this property as a form of the law of one price.

This is a topical issue in current actuarial ERM valuation methods. As will be discussed further below, typical current valuation practice may imply, for example, that a certain cashflow generated by mortgage #1 in 2040 has a different value to a certain cashflow generated in 2040 by mortgage #2. This might be described as implying the two mortgages offer different illiquidity premia, even though the two assets have identical levels of liquidity.

A valuation methodology is developed below that avoids these types of internal inconsistencies. This ensures that a valuation method can be implemented that can work demonstrably well in its basic task of obtaining asset values that are consistent with the observable market prices of other relevant assets; but equally importantly, it provides a solid foundation on which the other methodologies – for valuation of securitisations, and capital assessment of both mortgages and their securitisations – can be logically constructed.

2.1 Current ERM Valuation Practices

Equity Release Mortgages (ERMs) are complex illiquid assets. Their valuation inevitably relies upon estimates of many unobservable inputs, as well as inputs in the form of the observable prices of relevant assets. Valuations can take place in a number of different contexts – fair valuation for the purposes of financial reporting, economic valuations as part of regulatory tests and so on. This paper will not consider the specific requirements of particular accounting standards; it will consider the more general requirements of market-based asset valuation.

The valuation techniques and assumptions used by UK insurance firms vary both according to the valuation context and from firm to firm. These techniques and assumptions are not fully publicly disclosed. So, there is no single current 'industry' valuation approach and it is only possible to discuss current practices in some generality. However, there are some areas of commonality that we can tentatively identify in current ERM valuation practices:

1. The valuation model explicitly recognises ERM's No-Negative Equity Guarantee (NNEG) and values it as a distinct component of the mortgage cashflow(s). Option pricing techniques are used to value the NNEG (which may be implemented using a Black-Scholes formula or through the use of a stochastic cashflow simulation model).
2. The NNEG valuation is done using best estimate assumptions for relevant unobservable variables such as the volatility of the price of the residential property and its deferment rate (as well as borrower assumptions such as decrement rates, prepayment behaviour etc.). The estimation of house price volatility is usually based on an analysis of historical house price data (property indices and in some cases individual house price transaction data). The deferment rate has a less well-defined interpretation, and a variety of approaches are used to set the deferment rate assumption.
3. The valuation process aims to produce a value for a loan at origination that is equal to the starting loan amount (possibly plus an allowance for acquisition costs). Observed retail market ERM origination prices are usually used as the key reference asset in the calibration of the valuation model for in-force ERMs.

4. The initial model valuation that is produced for the reference ERM loans in 3. by the best estimate parameter assumptions developed in 2. is then adjusted so that the origination loan amount is obtained.
5. This adjustment is not made by adjusting the parameter values of one (or more) of the unobservable house-specific assumptions such as its price volatility or deferment rate. Instead, the adjustment factor may simply be applied as a ratio of whatever size is required to obtain the loan amount, or by finding the illiquidity premium parameter input (i.e. the increase to the risk-free rate) that recovers the origination value for that specific mortgage (or some combination of both).
6. The adjustments produced in step 5 are the fundamental output of the valuation calibration process. They are then mapped to the valuation of in-force ERMs (i.e. the assets with unobservable prices that are to be valued). That is, each in-force ERM is mapped in some way to one of the reference prices, and the adjustment that was applied to that reference price is then applied to the valuation of the asset of interest.

Beyond these areas of commonality, many variations currently exist. This form of valuation process can do a quite reasonable job of valuing a portfolio of in-force ERMs. But its limitations become more telling as the use of the model is extended beyond base valuation.

There are a couple of key limitations to the above valuation logic: if the valuation adjustment in step 4 is applied as a fitted ratio between the values produced by the model and the observed origination loan amounts, we have no obvious way of determining how the adjustment behaves in stressed valuation conditions (necessary for 1-year VaR capital assessment), or in the valuation of derivatives of the mortgages (i.e. securitisations). And if the derived valuation adjustment is expressed as a mortgage-specific illiquidity premium adjustment, it may result in valuations that do not follow basic economic properties such as the law of one price (i.e. the idea that two assets with identical cashflows and the same liquidity state should have the same value). As noted above, it implies that if two mortgages were to produce certain cashflows, they would have a different value, even though the cashflows and liquidity of the mortgages are identical.

2.2 Overview of the proposed methodology

Chapter 2.1 raised a technical concern with current valuation practices: they either use uninterpreted adjustments that are very difficult to confidently extrapolate, or they assume inconsistent (contradictory) general assumptions in the valuation of each mortgage (e.g. assuming the market's reward for bearing illiquidity risk is different for each mortgage despite them having the same level of illiquidity). This paper argues that a next generation of ERM valuation methods can build on and enhance the valuation methods that have been implemented thus far so that these concerns are removed. In doing so, the mortgage valuation model will provide a stronger foundation for tasks beyond current valuation of the mortgages (in particular, securitisation valuation). It can do so by using a slightly different valuation logic that is more aligned to conventional derivative valuation methods.

This paper proposes a methodology where the valuation adjustments that generate the fits to observed prices in steps 4 and 5 of the above process are made by adjusting assumptions that are explicit and specific to the given mortgage (e.g. the price volatility or deferment rate of the house underlying the mortgage) instead of assuming inconsistent universal assumptions (e.g. different illiquidity premia for assets with the same illiquidity). We will see that this shift in the logic of the valuation process produces a valuation methodology that avoids the need for arbitrary or *ad hoc* further adjustments, and that will naturally extend to securitisation valuations and the assessment of capital requirements for both the mortgage portfolio and securitisations.

The concept of fitting a derivative value to an observed price by finding the valuation model parameter for the behaviour of the underlying asset of the derivative that produces a model value equal to the observed price is, of course, completely standard. There are real practical benefits to aligning ERM valuation methodology to this conventional derivative valuation logic. However, ERMs, and their NNEGs, are more complex than vanilla options on interest rates or equity indices. When valuing, say, an equity index option, we usually have only one variable that is unobservable (the index's volatility). The 'deferment rate' of the equity index can usually be derived from observable index forward prices. The question of an illiquidity premium tends not to arise as these securities are usually reasonably liquid (at least, more liquid than mortgages). In the case of an equity release mortgage, we have several variables that are not directly observable. Neither the volatility of the price of the underlying house nor the deferment rate of the house are directly observable. Moreover, the mortgage value depends on borrower characteristics and behaviours that also cannot be directly inferred from observable market prices. Finally, if we believe that illiquid assets earn an illiquidity premium, this also cannot be easily observed from market prices (the prices of illiquid assets naturally tend to be difficult to observe, at least frequently).

It is not, therefore, argued in this paper that the valuation methodology proposed below is a panacea that suddenly makes the valuation of this illiquid and complex asset easy. Rather, our objective is to ensure that the mortgage valuations are economically well-behaved and consistent with each other (for example, not using contradictory assumptions), because these properties make the model more transparent and more reliably extendable to other key purposes such as securitisation valuation and the assessment of capital requirements.

So, it has been argued above that the methodology should fit to observable ERM market prices by adjusting one of the mortgage-specific valuation model parameters that cannot be directly observed from other market prices. As there is more than one such parameter in the case of ERMs, we need to decide which one to mark to prices, and which to estimate using other methods. We align as closely as possible to derivative valuation conventions by choosing to fit the house price volatility. Therefore, instead of using best estimate volatility assumptions, the valuation methodology will derive implied house price volatilities from the observable prices of reference ERM assets. We will then assume the same form of house price volatility behaviour applies to the valuation of other ERM assets (that do not have observable ERM prices).

As we will discuss below, such an approach does not require us to ignore the possible effects of illiquidity on the ERM values. But it avoids the incoherent practice of explaining pricing differences by assuming different assets with the same level of liquidity / illiquidity have arbitrarily different illiquidity premia (or some other uninterpretable valuation adjustment factor).

This means that *the valuation methodology will not make use of historical data on house price returns at any stage of the mortgage valuation process*. This may strike some as very peculiar – how can we value a NNEG without estimating the price volatility of the house? But as noted above, our objective here is 'only' valuation. If we were tasked with deciding how to price a new ERM product, such analysis of price volatility data may be an important consideration in deciding what price to offer our product at; or in evaluating how cheap or dear current market pricing is. But the valuation task is merely to interpolate and extrapolate observable prices, and not to determine if they are 'right'.

The process of deriving ERM-implied house price volatilities may sound abstract and complicated, and a lot harder than deriving a scaling ratio to apply to ERM values produced by a model with a best estimate volatility. You may ask what great practical advantage is delivered by this approach. After all, both approaches will result in valuations that are in some sense consistent with observable prices. Let's use a simple example to illustrate this key point.

Suppose we have two observable ERM origination prices today, both for a borrower of age 55. Mortgage A has a Loan-to-Value (LTV) of 30% and a mortgage rate of 4%; mortgage B has an LTV of 20% and a mortgage rate of 3%. We assume a best estimate for future house price volatility of 15% and that the houses have the same deferment rate. We find that our model under-values Mortgage A and over-values Mortgage B. Specifically, we find that we need to apply a 110% adjustment factor to the Mortgage A value, and a 90% factor to the Mortgage B value in order to recover the starting loan amounts from the valuation model. We therefore decide to apply these valuation factors in the valuation of our portfolio of in-force ERMs.

We next decide to assess capital requirements for Mortgages A and B. What valuation adjustment factor should be applied to Mortgage A when house prices fall by 30%? What should it be when house prices rise by 30%? Well, we might argue that the house price rise would make Mortgage A look quite like the current Mortgage B, and so we could reduce the valuation adjustment factor from 110% to 90% or so in this stress. By this logic, the valuation adjustment factor should be higher when the Mortgage A house price is stressed down. Perhaps we should change the Mortgage A valuation adjustment factor from 110% to 120% or 130% in that stress? What if the house price halves and the LTV doubles to 60%? Suppose we project over many years and the LTV reaches 110%, what should the adjustment factor be then? Maybe 150% or perhaps 200%? We cannot easily say.

Now suppose we instead calibrate our valuation model by assuming Mortgage A and Mortgage B are valued using different house price volatilities. For example, suppose the value of Mortgage A is found to equal the starting loan amount when the house price volatility is assumed to be 13%. We will call this Mortgage A's implied volatility. Suppose we find Mortgage B has an implied volatility of 17%. Because this valuation adjustment has taken the form of a change to a parameter of the valuation model, the changing impact of the adjustment in different scenarios arises more naturally and transparently. We would find, for example, that as LTV increases to high levels, the adjustment tends not to 150% or 200% as speculated above, but to 100% (i.e. the ratio of the valuation with our best estimate volatility of 15% to the valuation with the 'implied' volatility of 17% or indeed any other level tends to 1 as the LTV tends to a very high value). This is because a very high LTV mortgage looks increasingly like deferred possession of the house, and its value is not sensitive to the assumed level of house price volatility. The value of the deferred possession is only a function of the deferment rate. If we believe the two mortgages have different deferment rates, we can of course assume they have different deferment rates in their valuations. But if we believe they have the same deferment rate (and the same illiquidity if a non-zero illiquidity premium is assumed), then assuming the two mortgages have different values when their LTVs are very high implies two assets with identical cashflows have different values. We avoid this valuation contradiction when we use an approach that starts with a more economically logical form of valuation adjustment.

The purpose of the above example is not to suggest that current valuation methods result in capital numbers being 'wrong' in the magnitudes described here. The numbers in this example are plucked out of thin air. The point is that the current methods can accidentally result in economically unreasonable behaviour, and great care is likely to be required to be sure of avoiding this. Under the proposed enhancement to current methodologies, the valuation process is more economically logical, and this makes it simpler and more transparent and therefore less prone to these potential accidents.

2.2.1 The basic cashflow of an equity release mortgage

As is now well-documented in a wide range of ERM research literature, the basic cashflow of the equity release mortgage (ERM) has the form min [X(T), S(T)], where X is the loan balance at time T (usually a fixed amount for any given future date as ERMs generally have a fixed mortgage rate), and S(T) is the value of the house the mortgage is written on at time T.

The timing of the cashflow, i.e. when T occurs, is uncertain. For now, we will consider the valuation conditional on a known time T, and we will return to this point shortly below.

Conceptually, the valuation of this cashflow, conditional on a known maturity time T, can be considered to consist of three essential steps:

1. The valuation of the (fixed) cashflow X due at time T
2. The valuation of the future possession of the house, S, at time T
3. The functional transformation of these two valuations that is required given the ERM cashflow is the *minimum* of the two cashflows.

None of these three steps are completely straightforward, but each of them can be logically tackled in an economically coherent way. We consider these three valuation components in turn below.

2.2.2 Valuing the fixed cashflow

The valuation of a fixed cashflow would seem to be the simplest valuation of all. The cashflow is certain (riskless), and so the present value of the cashflow is found by discounting it at the risk-free rate applicable to the term of the cashflow. This rate would be derived from the observable prices of liquid risk-free assets such as government bonds or interest rate swaps in the usual ways.

However, one form of complication may arise. If the mortgage is illiquid, and it is assumed that a non-zero illiquidity premium exists, then it is necessary to adjust the risk-free discount rate so that it is consistent with the interest rate implied by the price of a risk-free asset with the same level of liquidity as that of the mortgage asset. If the mortgage is as liquid as the observable risk-free asset (or no illiquidity premium is assumed), then no adjustment is required.

The observation of the illiquidity premium is generally a difficult problem. It requires the observation of the price of an illiquid asset and, almost by definition, those prices are not easily observed. Nonetheless, if the mortgage valuation method is going to make an allowance for mortgage illiquidity and the presence of an illiquidity premium, then the right way to do it is by adjusting the assumed risk-free rate that applies to all mortgages of same illiquidity (rather than assuming mortgages with the same illiquidity earn different illiquidity premia). This approach means that we will retain the fundamental property that two assets with the same liquidity state and identical cashflows in all states of the world will have the same value according to the valuation model.

The mortgage valuation impact of the assumed illiquidity premium will reduce as the LTV of the loan increases. This is economically logical. As the LTV of the loan increases, the loan increasingly resembles deferred possession of the house. The value of the deferred possession of the house does not depend on the liquid risk-free rate or the assumed size of the illiquidity premium. It depends on the current value of the house and the deferment rate. If the house is assumed to be illiquid, then no adjustment for illiquidity is required when valuing the illiquid mortgage as a function of the illiquid house price. The illiquidity valuation adjustment is already in the house price[2].

The valuation methodology will therefore take two forms of interest rate input:

- The (liquid) risk-free yield curve;
- The illiquidity premium term structure that applies to risk-free assets that have the same degree of illiquidity as the mortgage.

Again, the latter component is only required if the mortgage is illiquid and a non-zero illiquidity premium is assumed.

The concept of an illiquidity premium can be a divisive topic. No view is offered in this paper on the empirical 'reality' or otherwise of an illiquidity premium in asset pricing in general or equity release mortgage pricing in particular. The proposed valuation methodology does not *require* a valuation adjustment for illiquidity. Those who prefer not to assume an illiquidity premium can, of course, simply specify the illiquidity premium term structure is zero at all terms and carry on.

2.2.3 Valuing deferred possession of the house

The value at time t of the deferred possession of the house to some future time T can be written as:

$$\text{Deferred House Possession } (t, T) = S(t)\exp[-d(T-t)]$$

where S(t) is the house price at time t and *d* is defined as the deferment rate for the house.

Standard derivative theory would consider *d* to be the income yield of the underlying asset, in this case the house. However, there are many complications that arise when considering how to set *d* in the context of deferred possession of a residential property:

- Contracts for deferred possession of residential properties do not trade in active markets. The value for *d* therefore cannot be directly derived from observable market prices[3].
- Economic theory suggests *d* is the income yield of the underlying asset. In the context of a residential property, this quantity would be net of costs such as ongoing property maintenance and an expected frequency of rental voids. The net rental yield of an individual house is not a directly observable quantity.
- The above formula assumes the quantity *d* is fixed (deterministic). Even where the current level of the net income yield can be accurately observed, we must acknowledge that its future level is not certain. Economic theory does not have a straightforward account of how to allow for income uncertainty in the valuation of derivative contracts such as deferred possession, and any such allowance will be sensitive to the form and size of uncertainty that is assumed. ERMs are often very long-term contracts, and so are potentially subject to material amounts of income uncertainty.
- If the mortgage in which the right to deferred possession is embedded is more liquid than the house, and a positive illiquidity premium exists, then the value of deferred possession needs to be adjusted upwards to reflect that it would be worth more than if it had the same illiquidity as the house. This adjustment can be made by deducting the assumed illiquidity

[2] For further discussion of this topic, see my article of 29th March 2020: https://craigturnbullfia.com/notes-on-derivative-valuation-and-illiquid-assets/

[3] If we could observe the price of an ERM with an ultra-high LTV, the deferment rate could be derived from that price (as the value of such a mortgage would not be sensitive to the house price volatility). But such mortgages are (hopefully!) quite rare, and they tend not to have observable prices.

premium from the deferment rate[4]. In the following discussion, we will assume the mortgage is as illiquid as the house, and hence no such adjustment to the value of deferred possession will be necessary. But there is nothing in the methodology that precludes making such an adjustment if the mortgage is assumed to be more liquid than the house.

Finally, we should note a further complication in the valuation of the right to deferred possession: the above formula shows that the current value of the deferred possession of the house is a function of the current house value, S(t). There may not be a reliably accurate way of observing this current house value. It is an illiquid asset and it is quite possible that the mortgage is written on a house that has not traded in years or even decades. This valuation methodology, like the current ERM valuation methods used by UK insurance firms, will assume the current house price is known.

All of the above issues combine to make the valuation of long-term deferred possession of a residential property undoubtedly difficult. The interpretation and calibration of the deferment rate has perhaps been the single ERM valuation topic for which there has been the least amount of consensus in the research output of the last couple of years. This valuation methodology will interpret *d* as the net income yield of the given house, and we will assume it can be estimated with reasonable accuracy. This approach is broadly consistent with the approaches advocated by Jefferey & Smith (2019), Dowd et al (2019) and also the IFoA ERM Working Party's 2020 Discussion Paper on ERM economic valuation[5].

The difficulties in observing S(t) and *d* are significant limitations of the valuation methodology and should be recognised as such. It is not obvious how to mitigate these difficulties, and they also arise in the current ERM valuation methods used by UK actuaries. There is no obvious escape from the fact that the mortgage value is undoubtedly related to the current value of the house, and this house value may be difficult to observe. Illiquid asset valuation has some inevitable limitations, and it is important to bear these in mind when interpreting and using the valuation results.

2.2.4 Allowing for uncertainty in cashflow timing

Thus far, it has been assumed that the mortgage cashflow arises at a known time T. An equity release mortgage does not have a fixed maturity date. The maturity is determined by when the borrower dies or enters long-term care (ore prepays, a topic we will return to later). It is conventional to write the equity release mortgage value as a probability-weighted average of the mortgage cashflows that may arise at the various possible future maturity dates. Each of these mortgage cashflows may be referred to as 'ERMlets' and we will use this terminology here. So, ignoring prepayment for now, and including an allowance for long-term care entry in the decrement table, we can write the value of the mortgage at time t with a borrower age x at time t as:

$$\text{ERM}(t) = \sum_{T=1}^{\infty} q_{x+T-1,t+T-1}\, P_{x,t,t+T-1} ERMLet(t, S(t), X_{t+T}, t+T)$$

where:

- $q_{x,t}$ is the probability of someone age x at the start of the year t dying during year t;
- $P_{x,t,T}$ is the probability of someone age x at start of year t surviving to the start of year T;
- ERMLet(t, S(t), X_T, T) is the value at time t of a mortgage cashflow arising at fixed time T, min[X_T,S(T)].

[4] Again, more on this topic can be found in my article of 29th March 2020: https://craigturnbullfia.com/notes-on-derivative-valuation-and-illiquid-assets/

[5] IFoA Equity Release Mortgages Working Party (2020)

Strictly, the above formula assumes the borrower's mortality / long-term care entry rates are independent of changes in house prices (and anything else that is assumed to be stochastic in the valuation). We will make this assumption throughout this paper. The methodology will assume mortality and LTC entry rates are known. Given these assumptions, the ERM can be valued as a portfolio of ERMlets as per the above formula.

The presence of a deterministic prepayment rate that is some function of age and time alters the ERM cashflows as follows:

$$\text{ERM(t)} = \sum_{T=1}^{\infty} q_{x+T-1,t+T-1}\, P_{x,t,t+T-1} ERMLet(t, S(t), X_{t+T}, t+T) \;+ \\ \ldots \sum_{T=1}^{\infty} r_{x+T-1,t+T-1}\, P_{x,t,t+T-1} RFValue(t,\; X_{t+T}, t+T)$$

where:

- $r_{x,t}$ is the prepayment rate of a borrower age x in year t;
- RFValue(t, X_T, T) is the value at time t of the certain cashflow X_T payable at time T;

Here, P has a slightly different interpretation to the previous formula: instead of being the probability of being alive at the start of the year, it is the probability of the mortgage being in-force at the start of the year. Again, in this case of deterministic prepayment rates, every ERM can then be valued by valuing its constituent ERMlets.

2.2.5 Valuing the mortgage cashflow

The current value of the mortgage cashflow due at time T, min [X, S(T)], will be a function of the current value of the certain cashflow X payable at time T (as determined in Chapter 2.2.2) and the current value of deferred possession of the house to time T (as determined in Chapter 2.2.3). And the size of this cashflow will be determined by the probability weights attached to the mortgage maturing at that time, as determined by the formulae in Chapter 2.2.4.

It is well-established that the current value of the mortgage cashflow can never be greater than the current value of the certain cashflow X or the current value of deferred possession of the house. When one of these quantities is much larger than the other, the mortgage cashflow value will be closely approximated by the value of the lower quantity. But when the two quantities are of similar magnitude, a model is required to obtain a functional relationship between the value of the mortgage cashflow and the values of the fixed cashflow and deferred possession of the house.

This functional relationship is usually developed by writing min [X, S(T)] as X - max (0, X - S(T)). Written in this way, we can see that the value of the mortgage cashflow is the value of the certain cashflow X less the value of a put option on the house. This put option is commonly referred to as the No-Negative Equity Guarantee (NNEG). So, we can write:

Value [ERMLet(t, S(t), X_T, T)] = Value [X_T] + Value [0, X_T - S(T)]

The value of the certain cashflow, X_T, is determined by discounting it at the appropriate liquidity-adjusted risk-free rate:

Value [X_T] = X_T.exp[-{r(t,T) + λ(t,T)}(T-t)]

where r(t,T) is the liquid risk-free rate at time t of term (T-t) and λ(t,T) is the illiquidity premium at time t of term (T-t).

The second component of the ERMlet, the NNEG, is a put option on the house with underlying asset value S(t); strike price X_T; term (T- t); risk-free rate of r(t,T) + λ(t,T); and underlying asset income yield of *d*. The probability distribution of S(T) has not yet been considered, and so, thus far, we do not have a terminology for describing it or its impact on the NNEG valuation.

The key challenge in the valuation of the mortgage cashflow can therefore be considered as the valuation of the NNEG put option. This requires an option valuation model. An option valuation model is based on the recognition of uncertainty in the option's cashflow. This uncertainty is captured by assuming at least one of the factors that determine the option's cashflow are stochastic. The form of the option valuation model is determined by the choice of which valuation risk factors are assumed to be stochastic, and the form of probability distributions that are specified for them.

In the context of the cashflow of the No-Negative Equity Guarantee, several factors could potentially be modelled as stochastic processes in the valuation process:

- Most obviously, the house price could be assumed to have a stochastic process.
- Interest rates (and price inflation). Interest rate uncertainty may be a factor that influences long-term house price uncertainty. Furthermore, we might expect mortgage prepayment behaviour to be correlated with interest rate levels.
- The deferment rate. As noted in the discussion of Chapter 2.2.2, houses' net income yields are uncertain, especially over the long-term. Variation in deferment rates could be a factor in house price volatility.
- Prepayment rates.
- Mortality / Long-Term Care rates.

The methodological arguments presented so far do not compel us to assume that any of the variables above should or should not be considered stochastic within the valuation model. Our choices here will ultimately be guided by more general modelling considerations. Most importantly, unnecessary complexity is best avoided (Occam's razor). Therefore, if the stochastic process does not materially help us in our task of valuing mortgages consistently with observable mortgage market prices, it is preferable to assume the variable is deterministic. These choices are *implementation choices*, rather than a fundamental part of the methodological framework. We will therefore defer further discussion of these choices to Chapter 3, where some implementation case studies are developed.

The key point from a methodological perspective is that we must specify which underlying valuation factors are stochastic, and which forms of stochastic process they follow. This dictates the form of valuation function that will then be calibrated to available market prices. Once this valuation function has been determined, the valuation strategy will look similar to that used in standard derivatives valuation methodology:

1. We determine a mortgage valuation function by specifying which variables follow stochastic processes. If these stochastic processes satisfy some standard conditions, this will imply a theoretical arbitrage-free relationship between the value of the mortgage cashflow, the value of the fixed cashflow and the value of deferred possession.
2. We will identify a set of relevant ERM assets with observable prices. In the implementation case studies that follow in Chapter 3, these will be assumed to be retail origination prices, but the methodology does not require them to be.
3. We will also identify the other relevant inputs discussed above, i.e. the risk-free yield curve, the risk-free asset illiquidity premium term structure, the house prices and deferment rates associated with the ERMs identified in step 2.

4. We will make the various necessary assumptions about borrower characteristics and borrower behaviour such as mortality rates, rates of entry into long-term care and assumptions about prepayment behaviour for the borrowers who transact the reference ERMs.
5. We then calibrate the parameters of the valuation model determined in step 1 above to the observable prices identified in step 2, given the other input assumptions specified in steps 3 and 4. That is, we will find parameters for the specified stochastic processes for house prices (and any other variables assumed to be stochastic) that result in valuations of the reference ERM assets in step 2 that are equal to their observable price. (We know from standard option pricing theory that the valuation function determined in step 1 will not require an estimate of the expected return of the house).
6. We will then use this calibrated valuation model to value ERM assets with unobservable prices (e.g. a book of in-force mortgages). The borrower characteristics and behavioural assumptions used in the valuation may or may not be the same as the borrower assumptions used in step 4 for the reference ERMs.

Step 5 above is the part of the valuation methodology that is most distinct from current ERM valuation approaches. The details of the implementation of this step of the process depend on how the valuation model is specified in Step 1. Chapter 3 will provide a couple of detailed examples. However, before proceeding to those, there is an important general point to consider that is relevant to the calibration process.

When the valuation model is calibrated to produce values for the reference ERMs that are equal to their observable prices, we will do so by valuing each ERMs constituent ERMlets (according to the formulae developed above). Different reference ERMs may have exposures to ERMlets of identical or very similar forms. For example, if there are two 30% LTV mortgages with the same mortgage rate, one with a current borrower age of 55 and one with a current borrower age of 56, the form of the 1-year ERMlet cashflow will be the same, though the two mortgages will have slightly different portfolio weights to this option. It would be unsatisfactory for the valuation model to assume these identical ERMlets have different prices.

So, when the valuation model is calibrated, this will be done by 'looking-through' the reference ERMs to their ERMlet portfolios and finding the ERMlet valuation assumptions that provide the best fit to the collection of observed ERM prices. Put another way, the valuation methodology will assume that the value of ERMlet(t, $S(t)$, X_T, T) does not vary by mortgage.

The above discussion sets out the defining features of an ERM valuation methodology and its economic logic. To those with a background in derivative valuation, it may seem like a rather obvious approach to valuation and they may wonder why so many tedious pages have been used in discussing it. But the valuation methodology described here has some fundamental differences with the existing ERM valuation practices of UK insurers. As far as I am aware, no UK insurance firm values ERMs in this way in any context.

We now move on to some implementation case studies that aim to illustrate how the logic of this methodology can be applied in practice.

3 Valuation Implementation Case Studies

The valuation methodology that was set out in Chapter 2 leaves many implementation choices open. This chapter explores some of these choices and illustrates how they can be implemented in the valuation of an example portfolio of 1000 mortgages. Two alternative valuation implementations are developed for this same portfolio. These two implementations are discussed in detail in turn in Chapters 3.2 and 3.3. Chapter 3.1 sets out the common valuation assumptions and observable prices that will be used in both implementations.

3.1 Some Illustrative Valuation Assumptions and Observable Prices

3.1.1 Which things do we assume are stochastic?

Chapter 2 noted that there are many economic and demographic variables that may be assumed to be stochastic in the valuation of NNEGs. The methodological framework of Chapter 2 does not make any specific prescriptions in respect of which are assumed stochastic to be and which are not. It is an implementation choice.

It is, of course, quite difficult to envisage a NNEG valuation model performing well without stochastic variation in house prices, so let us take that as a given.

What about interest rates? There is, on the face of it, a strong case for assuming interest rates are stochastic in mortgage valuation, especially as it is natural to expect mortgage prepayment rates to vary with the level of interest rates. This sounds like it could be a fundamentally important factor in how mortgage values behave. However, recent research[6] suggests that stochastic variation in interest rates has a very second-order effect on equity release mortgage valuation, even in the presence of dynamic borrower prepayment behaviour that is fully rational. Given this result, the following two valuation implementations will assume interest rates are deterministic.

The house deferment rate is another natural candidate for stochastic modelling in valuation. We noted in Chapter 2.2.3 that there is long-term uncertainty in the level of houses' net rental yields. A stochastic process for the deferment rate would reflect this uncertainty. Stochastic variation in the deferment rate would also tend to create an upward-sloping term structure in house price volatility, which is a feature of empirical house price data. However, it is not clear that these benefits are worth the cost in modelling complexity. There are much simpler ways of fitting to implied volatility term structures. So, in the valuation implementations that follow, the deferment rate will be assumed to be deterministic rather than stochastic.

Uncertainty in mortality and LTC entry rates contributes to NNEG risk and would be expected to contribute to NNEG costs. As such, there is a case for considering their behaviour as stochastic processes in the model. My tenuous expectation is that the contribution of stochastic variation in mortality rates to NNEG risk will be second-order relative to the impact of the uncertainty in an individual house price, and so here again we will work with a deterministic assumption for mortality and LTC rates.

Stochastic behaviour in prepayment rates is a particularly interesting case and we will find that one of our implementation approaches cannot easily incorporate it whilst the other can. We will therefore defer discussion of prepayment modelling assumptions to within each of the implementation case studies below.

[6] See my 11th May 2020 article: https://craigturnbullfia.com/prepayment-risk-and-equity-release-mortgages/

Again, it is worth emphasising that there is nothing from a methodology perspective that precludes the assumption of stochastic variation in any of the above factors (or indeed others). It is simply an implementation choice that is based on a trade-off between valuation sophistication and Occam's razor. It is largely a matter of modelling judgment - alas, even 'mere' valuation still inevitably needs lots of it. I have erred on the side of simplicity in these examples because there is going to be enough for us to explore with only house prices generating stochastic variation. Different mortgage products and portfolios may warrant different choices in terms of the implementation set-up and the choice of stochastic valuation factors.

3.1.2 Risk-free rates

The assumed liquid risk-free yield curve at the valuation date is shown in Exhibit 3.1. The valuation assumes a risk-free illiquidity premium of 0.50% applies at all terms. The resultant illiquid risk-free term structure is also shown in Exhibit 3.1.

Exhibit 3.1: Risk-free yield curves

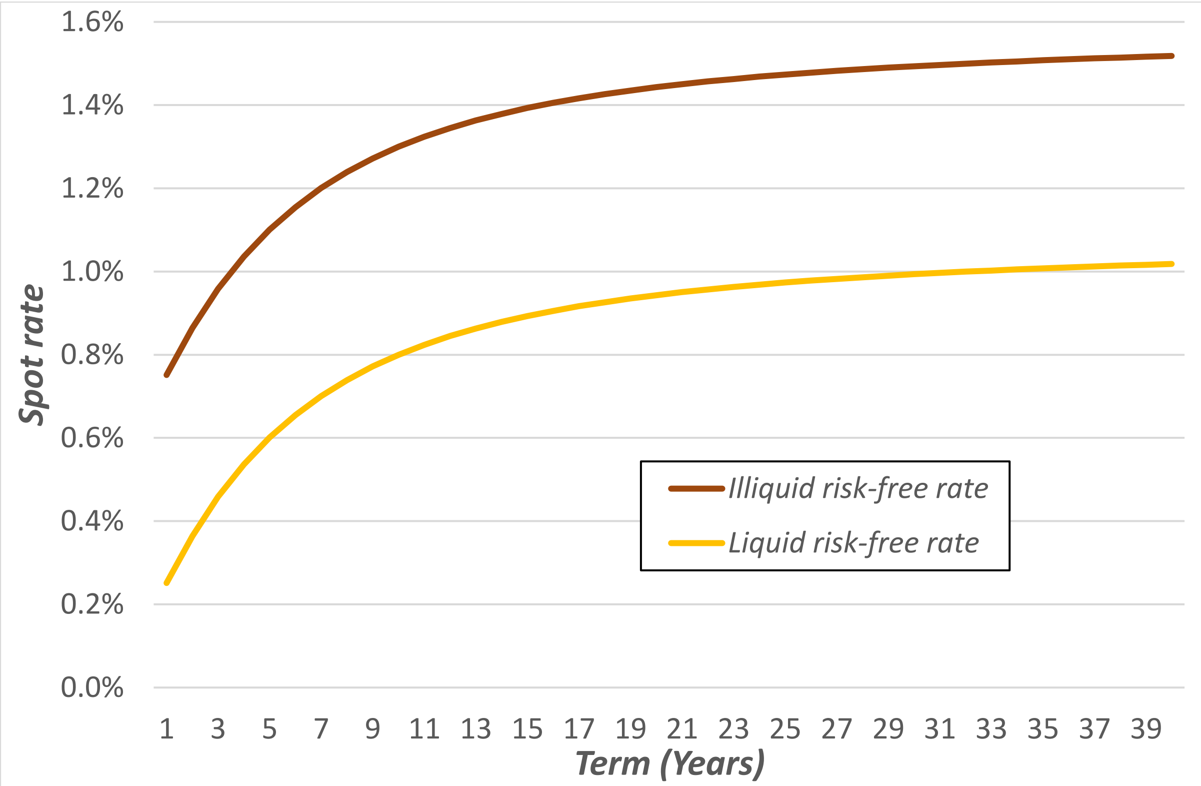

3.1.3 Mortality rates and rates of entry into Long-Term Care

The mortality table PMA08 with CMI2017 mortality improvements has been used as the mortality basis for all equity release mortgages (both those with observable prices that we will reference in the calibration of the valuation model; and those that we value with the valuation model). All mortality rates have been scaled up by a factor of 110% to make a simple allowance for the rate of long-term care entry.

3.1.4 House Deferment Rates

The valuation implementations will assume the current deferment rates of all houses for all terms is 3.5%pa. In a more sophisticated implementation, we could assume deferment rates vary by region, house size and so on, based on empirical study of rental income yields (as noted in Chapter 2.2.3, we cannot directly fit this valuation parameter to market prices as we cannot easily observe the value of contracts for deferred house possession).

3.1.5 Prepayment Rates

A deterministic prepayment rate will increase the value of the mortgage (all else equal), as it implies some fixed fractions of the NNEGs are simply being foregone by the borrower. The valuation effect will be more pronounced for mortgages with borrowers of younger ages, as the borrower will tend to live longer and will therefore have more time to prepay.

A deterministic prepayment rate assumption is arguably quite a sanguine description of borrower behaviour. It implies that, for example, a 90-year old borrower with an LTV of 150% is as likely to prepay as a 60-year old borrower who has just seen his house price double. But why would the 90-year old pay the full loan balance now when it is likely the NNEG will bite and reduce the repayment amount substantially when he or she dies, as is likely in the next few years? It is of course possible, but we might reasonably expect the likelihood to be reduced. And why wouldn't the 60-year old be more inclined to refinance their mortgage and obtain the lower LTV mortgage rate that the increased house value can now command?

At the other end of the behavioural spectrum from a deterministic prepayment assumption, we could assume borrowers behave rationally, and only prepay when it is in their financial interest – that is, when the prepayment amount is lower than the fair value of the mortgage if they do not prepay. This will tend to occur when their house price increases since origination, and / or when interest rates fall. Such an assumption would *reduce* the current mortgage value relative to the zero-prepayment case. The reality is, doubtless, somewhere in between these extremes – there may be some 'churn' in the portfolio that cannot be explained by direct financial incentives related to prepayment amounts relative to mortgage value (this will increase the mortgage value); and there may be a tendency to see higher prepayment frequencies when prepayment amounts are low relative to the no-prepayment mortgage fair value (this will decrease the mortgage value).

The implementation approach of Chapter 3.2 can easily accommodate deterministic prepayment rates but cannot easily accommodate dynamic prepayment behaviour. Allowing for deterministic prepayment frequencies without allowing for dynamic effects will, in my view, produce a distorted and biased view of valuation, so in the implementation case study considered in Chapter 3.2, we simply assume zero prepayment. The implementation approach of Chapter 3.3. can easily allow for both these types of prepayment behaviours, and the assumptions used there are described in full in that section of the paper.

It should be noted that different borrower demographic and prepayment assumption can be made for the reference ERMs and the ERM portfolio that we ultimately wish to value. In our implementations, for the sake of simplicity, we will use the same assumptions for both the reference portfolio and the mortgage portfolio that we value.

3.1.6 Observed ERM Origination Prices

Ten (hypothetical) observed ERM price quotations are used as the set of reference ERM asset prices that the valuation models of Chapters 3.2 and 3.3 will be calibrated to. The prices vary by Loan-to-Value and borrower age as set out in Exhibit 3.2 below.

Exhibit 3.2: Table of Ten Hypothetical ERM Prices

ERM #	Age	Loan-to-Value	Mortgage rate
1	55	25.5%	4.0%
2	60	30%	4.0%
3	65	35%	4.0%
4	70	41%	4.0%
5	75	47.5%	4.0%
6	55	19.1%	3.0%
7	60	22.5%	3.0%
8	65	26.3%	3.0%
9	70	30.8%	3.0%
10	75	35.6%	3.0%

The ERM prices apply to borrowers aged between 55 and 75. At each age there is a loan amount that attracts a 4.0% mortgage rate, and another LTV, that is (proportionally) 25% lower than the first, for which a 3.0% mortgage rate applies. These prices are plotted in the chart below.

Exhibit 3.3: Chart of the Ten Hypothetical ERM Prices

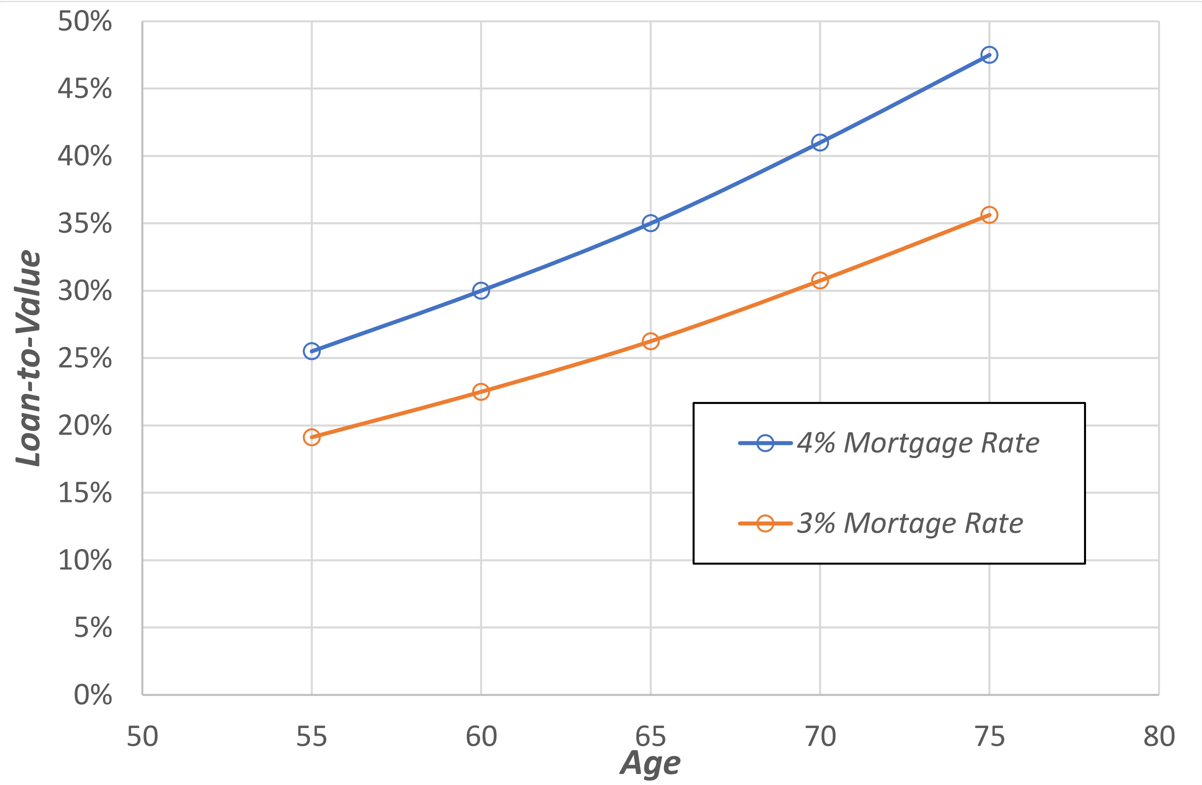

The price of these mortgages at the time of origination is the observed transaction price, i.e. the loan amount. It may be deemed reasonable to add a loading for origination acquisition costs to this transaction price. In the valuation case studies that follow we will assume an allowance for acquisition costs of 4% of the loan amount in the valuation calibrations.

These ERMs are assumed to have the same level of illiquidity as the house. The illiquidity premium term structure that is specified in Chapter 3.1.2 is assumed to correspond to this level of illiquidity.

3.1.7 The ERM portfolio to be valued

The two valuation implementations in Chapters 3.2 and Chapter 3.3 will use the above assumptions and observations to value a hypothetical in-force portfolio of 1000 mortgages. The in-force portfolio has a range of ages (between 55 and 90) and a range of current Loan-to-Values (between 16% and 120%). The average current borrower age is 72 and the average current LTV is 49%.

The distribution of current ages and current LTVs of this portfolio is shown in exhibits 3.4 and 3.5 below.

Exhibit 3.4: Distribution of Current Ages of Borrowers in the 1000 mortgage Portfolio

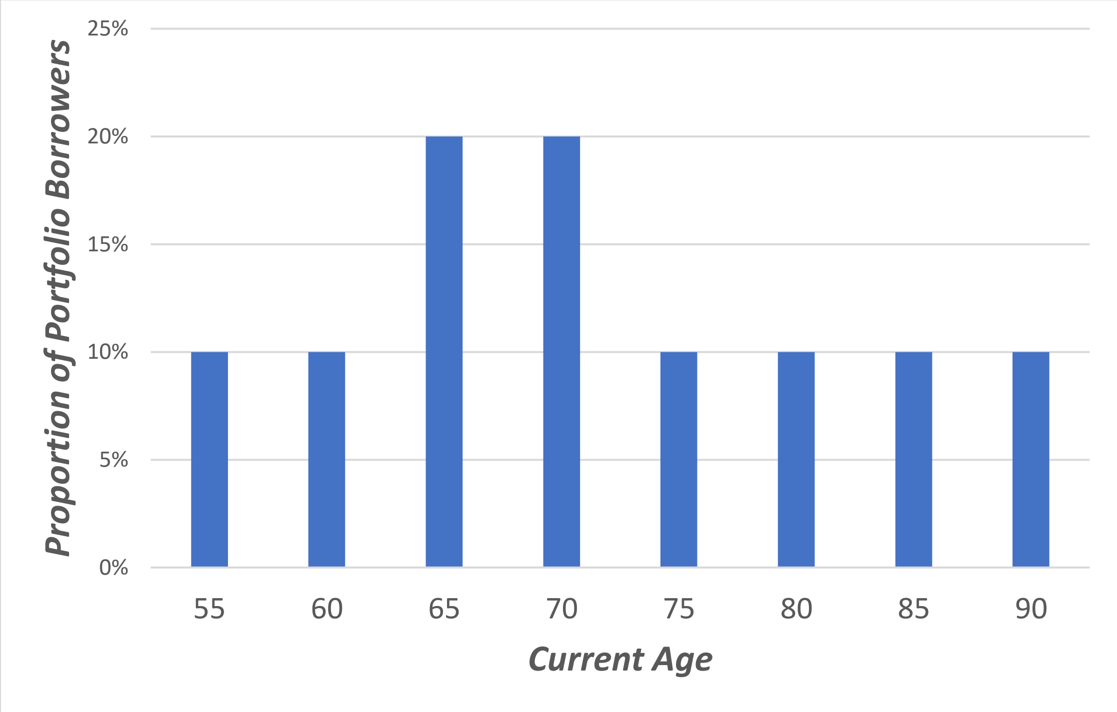

Exhibit 3.5: Distribution of Current LTVs in the 1000 mortgage Portfolio

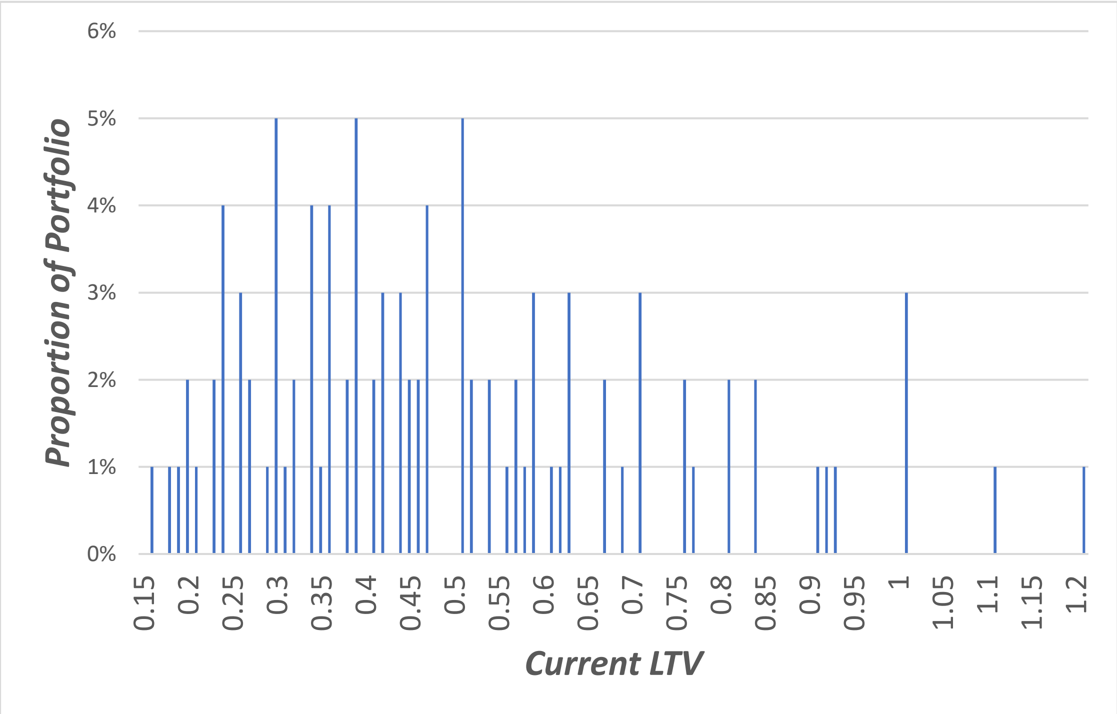

The current loan balances of the 1000 mortgages total £180m. The current house values underlying the mortgages total £410m.

These ERMs are assumed to have the same level of illiquidity as the reference ERMs.

In summary, in both of the following valuation implementations, the observable prices that we will use to value this in-force ERM portfolio of 1000 mortgages will be the 10 reference ERM prices in Chapter 3.1.6, together with the liquid risk-free yield curve and illiquidity premium term structure of Chapter 3.1.2. We also assume that the house prices associated with the 10 reference ERM prices are known – i.e. when a £300,000 loan is written with the 30% LTV pricing rate, the current value of the house is known to be £1m.

Given these prices, the assumption about the size of the house deferment rates (Chapter 3.1.4), and the demographic and prepayment assumptions for the mortgage borrowers (Chapters 3.1.3 and 3.1.5), we are now in a position to use a model to find values for in-force ERMs (that do not have observable market prices) that are consistent with the observed prices of the reference ERM assets.

As described in Chapter 2.2.5, the valuation methodology will seek to find a description of ERMlet house price volatility that correctly values the reference ERM assets. We will then apply these ERMlet volatility assumptions to the valuation of the in-force ERM portfolio.

3.2 A Black-Scholes-based Implementation Example

The discussion of Chapter 3.1.1 determined that this paper's valuation implementation models will assume that variation in house prices is the only source of stochastic variation in the valuation model. This still leaves many possibilities open to us in terms of the form of this variation, and we will explore a number of possibilities in this and the following sections.

We will begin with the simplest and most standard form of assumption for the behaviour of the underlying asset of a put option: we will assume the house price follows a geometric Brownian motion with a constant volatility. We will also assume the total return on the house (i.e. price change plus net rental income) includes a risk premium above the risk-free rate. We do not need to specify the size of this risk premium for the purposes of ERM valuation, as option prices do not depend on it. However, later, when we consider the expected return of the mortgage, we will specify an assumed size for the house risk premium.

The assumption that the house price follows a geometric Brownian motion with constant volatility means that the NNEG valuation is obtained by the Black-Scholes put option formula. That is, in the terminology of Chapter 2.2.5:

NNEG of ERMLet(t, S(t), X_T, T)= Value [0, X_T - S(T)] = Black-Scholes Put Option Formula

where the Black-Scholes formula assumes:

- Risk-free rate = r(t,T) + λ(t,T)
- Time to maturity = T - t
- Underlying asset value = S(t)
- Strike price = X_T
- Income yield = *d*, the house deferment rate.
- And a volatility parameter, σ, a free parameter to be calibrated to the reference ERM prices.

As in 'conventional' derivative valuation, the volatility parameter is a free parameter that we will calibrate so as to recover the observed ERM prices of the reference ERM assets (via ERMlet valuations).

3.2.1 Fitting the ERM-implied volatilities

We begin the calibration process by assuming all ERMlets have the same constant volatility. That is, we assume that the same volatility parameter enters all NNEG valuations, so we can write:

$$\text{NNEG of ERMlet } (t, S(t), X_T, T) = \text{Black Scholes } (t, S(t), X_T, T, \sigma)$$

As a first easy step, let us begin the calibration analysis by initially only considering the first reference ERM price – the mortgage on a 55-year old life with a mortgage rate of 4.0% and a starting LTV of 25.5%. We have one degree of freedom and one unknown (the house price volatility). We find that a volatility of 14.65% produces a value equal to the assumed observed price of 0.265 (i.e. 0.255 x 1.04, which allows for the 4% acquisition cost). So, the house price volatility implied by the price of ERM #1, together with the various assumptions set out in Chapter 3.1, is 14.65%.

Now suppose we wish to find a calibration of the model that fits to the first five reference ERMs, i.e. all the mortgages with a mortgage rate of 4.0%. We could simply repeat the above exercise and find the volatility assumption that is consistent with each ERM. But, as per the discussion of Chapter 2.2.5, this would mean different mortgages attach different values to the same constituent ERMlets. To avoid this, we will keep our volatility specification focused at the ERMlet level.

So, if we stay, for now, with the simple assumption that all ERMlets are valued using the same house price volatility, and we find the volatility that best fits to the first five reference ERM prices, we have one degree of freedom (the house price volatility) to fit to 5 observed prices. This simple approach may or may not result in 5 ERM values that are close to their 5 observed prices.

We find that a volatility of 14.25% produces the best valuation fits to the 5 prices. The valuation results are summarised in Exhibit 3.6.

Exhibit 3.6: Valuation Fits to the 5 ERM Prices (One Volatility Parameter)

ERM #	Age	Loan-to-Value	Mortgage rate	Observed Price[7]	Fitted Value	% Error
1	55	25.5%	4.0%	0.265	0.268	+0.9%
2	60	30%	4.0%	0.312	0.313	+0.4%
3	65	35%	4.0%	0.364	0.365	+0.2%
4	70	41%	4.0%	0.426	0.424	-0.6%
5	75	47.5%	4.0%	0.493	0.489	-0.9%

Exhibit 3.6 shows that the assumption of a volatility of 14.25% in the valuation of all ERMlets has produced some reasonable fits to the five ERM prices, with the valuation error never exceeding 1% of the mortgage value.

However, it can also be observed from Exhibit 3.6 that this volatility assumption tends to result in an overvaluation of the longer-duration ERMs and an undervaluation of the shorter-duration ERMs. This suggests we should be able to value the ERMs more accurately by assuming shorter-term ERMlets are valued using a lower volatility assumption than longer-term ERMlets. We can specify how the implied volatility varies with the term of the ERMlet so that it allows us to use the

[7] Note this includes the 4% allowance for acquisition costs.

specification to interpolate and extrapolate from these ERMlet valuations to the valuation of the ERMlets of ERMs with unobserved prices. To achieve this, we specify a functional form for how the volatility of each ERMlet is determined as a function of the term of the ERMlet:

$$\text{NNEG of ERMlet }(t, S(t), X_T, T) = \text{Black Scholes }(t, S(t), X_T, T, \sigma(t,T))$$

$$\sigma(t, T) = \sigma_0 + [1-\exp(-\theta(T-t))]\sigma_{inf}$$

So, we now have three volatility parameters instead of one. The parameter values $\sigma_0 = 0.1$; $\theta = 0.088$; $\sigma_{inf} = 0.15$ produce the following valuation fits to the five observed ERM prices:

Exhibit 3.7: Valuation Fits to the 5 ERM Prices (Three Volatility Parameters)

ERM #	Age	Loan-to-Value	Mortgage rate	Observed Price	Fitted Value	% Error
1	55	25.5%	4.0%	0.265	0.266	+0.1%
2	60	30%	4.0%	0.312	0.312	-0.1%
3	65	35%	4.0%	0.364	0.365	+0.2%
4	70	41%	4.0%	0.426	0.426	-0.1%
5	75	47.5%	4.0%	0.493	0.494	+0.0%

Exhibit 3.7 shows that allowing the house price volatility used in EMRlet valuations to vary as a function of the term of the ERMlet produces a very good fit to the five observed ERM prices. Exhibit 3.8 below shows the term structure of house price volatility that has been assumed in producing these ERM values.

Exhibit 3.8: Fitted ERMlet Implied Volatility Term Structure

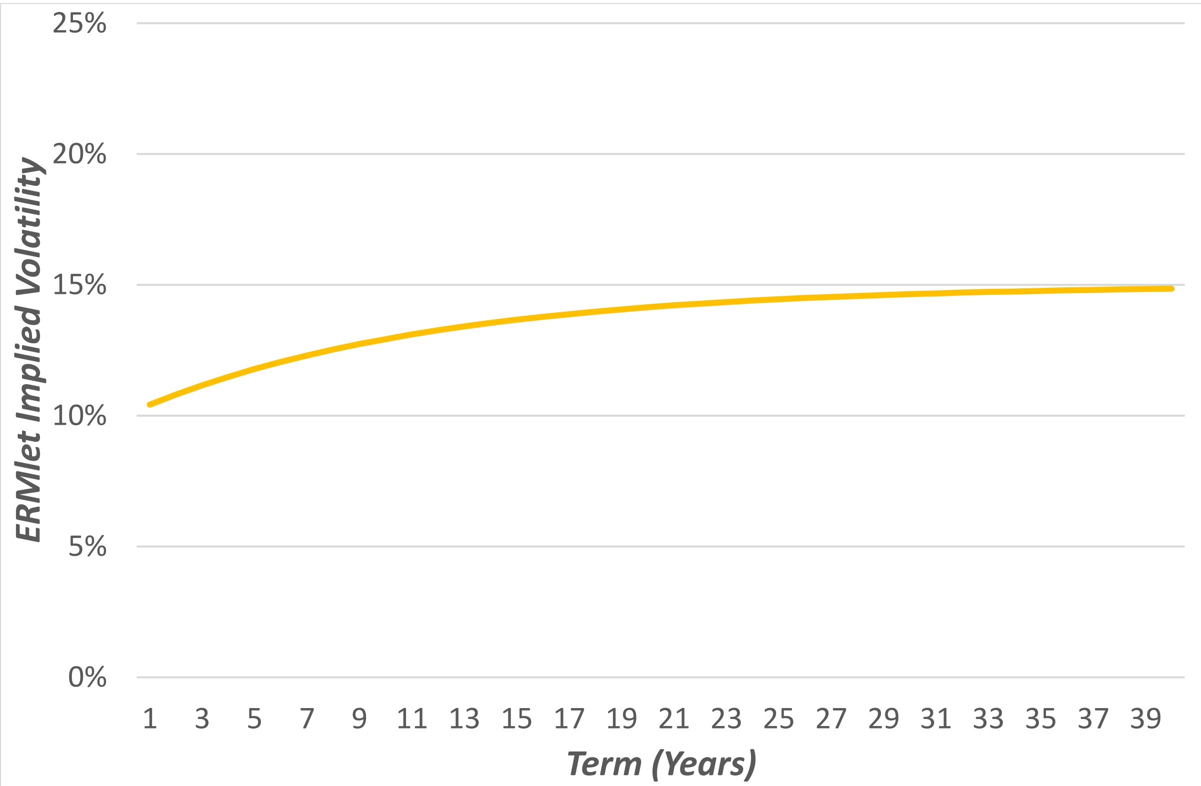

So far, we have only considered the fit to the observed prices of ERMs #1 to #5. If we take the 3-parameter volatility function fitted above to the prices of ERMs #1-#5, we obtain the following valuations for the ten reference ERM prices:

Exhibit 3.9: ERM Values for the 3-parameter volatility function (Volatilities fitted to ERMs #1-#5 only)

ERM #	Age	Loan-to-Value	Mortgage rate	Observed Price[8]	Fitted Value	% Error
1	55	25.5%	4.0%	0.265	0.266	+0.1%
2	60	30%	4.0%	0.312	0.312	-0.1%
3	65	35%	4.0%	0.364	0.365	+0.2%
4	70	41%	4.0%	0.426	0.426	-0.1%
5	75	47.5%	4.0%	0.493	0.494	+0.0%
6	55	19.1%	3.0%	0.199	0.205	+3.2%
7	60	22.5%	3.0%	0.234	0.242	+3.4%
8	65	26.25%	3.0%	0.273	0.284	+3.9%
9	70	30.75%	3.0%	0.320	0.332	+3.9%
10	75	35.63%	3.0%	0.371	0.385	+4.0%

You can see that when the volatility function that was fitted to the first 5 mortgage prices is used to also value the second 5, lower-LTV, mortgages, the model produces mortgage valuations that are 3%-4% too high in all five cases. The second 5 mortgages differ from the first in that they have lower LTVs and a lower mortgage rate. In option parlance, their NNEGs are lower strike put options and are further out-the-money than the first 5 mortgages. The volatility function that was fitted to the first 5 mortgages overvalues the second 5 mortgages. In other words, it undervalues the NNEGs of the second 5 mortgages. Put another way, to recover the mortgage values of the second 5 mortgages, higher house price volatility assumptions are required.

So, we have a choice. We could find a volatility function of the above form that best fits the ten prices. But this will inevitably undervalue the high LTV mortgages and overvalue the low LTV mortgages. The absolute scale of these errors may or not be judged to be material, but the inescapable point is that something is going on in the pricing behaviour that this valuation model is simply missing. This should make us wary of using the valuation model to extrapolate to values outside of the LTV range that we can observe in the reference prices. But that is exactly what we need to do when it comes to valuing the in-force portfolio. This leads to the conclusion that we ought to consider if there is some form of extension to the valuation model that can capture the pricing behaviour more satisfactorily.

The behaviour we are observing in these ERM prices is not an unusual feature of option pricing – it is often observed that lower strike options are priced with higher volatilities than higher strike options in equity option markets. This phenomenon is referred to as 'skew' in option-implied volatilities.

The above results suggest our volatility function needs one more dimension – as well as varying the ERMlet volatility by the term of the ERMlet, we could also vary the volatility parameter with the 'moneyness' of the ERMlet's NNEG.

We define a measure of NNEG moneyness as the ratio of the forward price of the house at the term of the ERMlet to the strike of the ERMlet:

[8] Including allowance for 4% acquisition costs.

$$f(t,T) = S(t).\exp\{ [r(t,T) + \lambda(t,T) - d](T-t)\}$$

$$m(t,T) = f(t,T) / X_T$$

r and λ are the liquid risk-free rate and illiquidity premium as defined above and *d* is the house deferment rate.

We can then generalise our volatility function so that the volatility assumed for an ERMlet is defined by the formulae:

$$\text{NNEG of ERMlet } (t, S(t), X_T, T) = \text{Black Scholes } (t, S(t), X_T, T, \sigma(t,T, m))$$

$$\sigma(t, T, m) = \{\sigma_0 + [1-\exp(-\theta t)]\sigma_{inf}\}m(t,T)^{\alpha.\exp(-\beta(T-t))}$$

We generally expect the skewness effect to be less impactful at longer maturities, and the exponential function that scales the α parameter produces this rate of 'decay'. So, the volatility function for the ERMlet implied volatility now has two dimensions and a total of 5 parameters. With parameter values $\alpha = 1$, $\beta = 0.05$, $\sigma_0 = 0.1087$; $\theta = 0.088$; $\sigma_{inf} = 0.1587$, we obtain the following fits:

Exhibit 3.10: Valuation Fits to the 10 ERM Prices (Five Volatility Parameters)

ERM #	Age	Loan-to-Value	Mortgage rate	Observed Price	Fitted Value	% Error
1	55	25.5%	4.0%	0.265	0.266	+0.3%
2	60	30%	4.0%	0.312	0.313	+0.2%
3	65	35%	4.0%	0.364	0.365	+0.3%
4	70	41%	4.0%	0.426	0.426	-0.1%
5	75	47.5%	4.0%	0.493	0.493	-0.3%
6	55	19.1%	3.0%	0.199	0.198	-0.5%
7	60	22.5%	3.0%	0.234	0.233	-0.3%
8	65	26.25%	3.0%	0.273	0.273	+0.0%
9	70	30.75%	3.0%	0.320	0.320	+0.1%
10	75	35.63%	3.0%	0.371	0.372	+0.3%

You can see that this volatility function produces a fit the 10 reference ERM prices that is within 0.5% of the observed price in all cases.

The volatility function here is an illustrative implementation example. The volatility function could be further generalised so that it fits exactly to all ERM prices if we wished. But the quality of fit produced by this volatility function suggests we have obtained a description of how to vary ERMlet implied volatilities by moneyness and term such that we can produce consistent valuations for other ERMs with unobservable prices.

The chart below shows how the model assumes ERMlet implied volatility varies with the moneyness and term of the ERMlet. Note a high *m* implies a low strike, i.e. an out-the-money put option. That is, an *m* of 2 means the house forward price is twice as large as the strike of the option.

Exhibit 3.11: ERMlet Implied Volatilities as a Function of Term and Moneyness

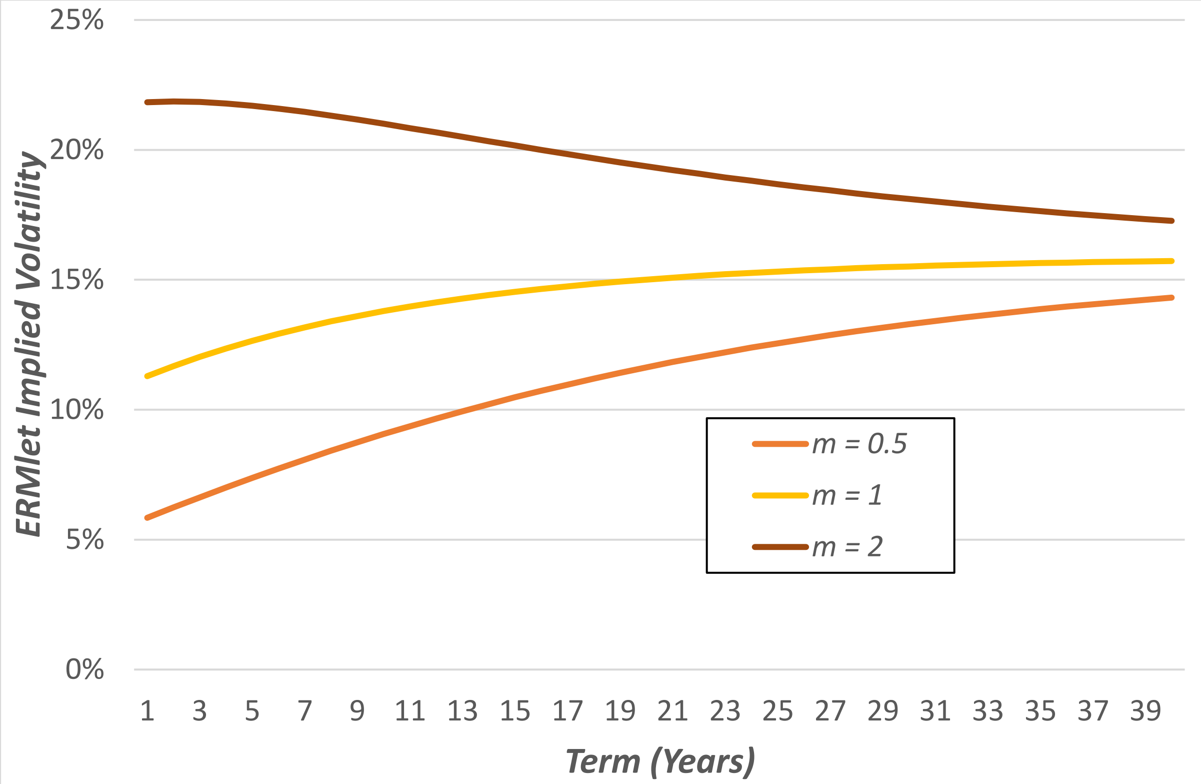

You can see that ERMlets with a lower strike (higher *m*) are valued with a higher volatility assumption, and the scale of this effect gradually reduces with the term of the ERMlet.

We have now determined an economic valuation basis for the ERMlets that can be used to value ERMs with unobservable prices.

3.2.2 Valuing the ERM portfolio

We now apply the valuation basis developed in Chapter 3.2.1 to the portfolio of 1000 in-force ERMs described in Chapter 3.1.7. To re-cap, the portfolio has total current loan balances of £180m and the total value of the houses underlying the mortgages in the portfolio is currently £410m.

The application of the valuation basis described by the assumptions in Chapter 3.1 and the two-dimensional implied volatility function basis fitted in Chapter 3.2.1 produces a current mortgage portfolio value of £170.3m.

Exhibit 3.12 plots the mortgage valuation results[9]. It shows how the mortgage valuations, expressed as a percentage of each mortgage's current loan balance, varies with the current LTV of the mortgage, for mortgages of various ages and mortgage rates.

[9] There are less than 1000 distinct plots on the chart because some of the mortgages have the same age, mortgage rate and LTV.

Exhibit 3.12: Valuation Results

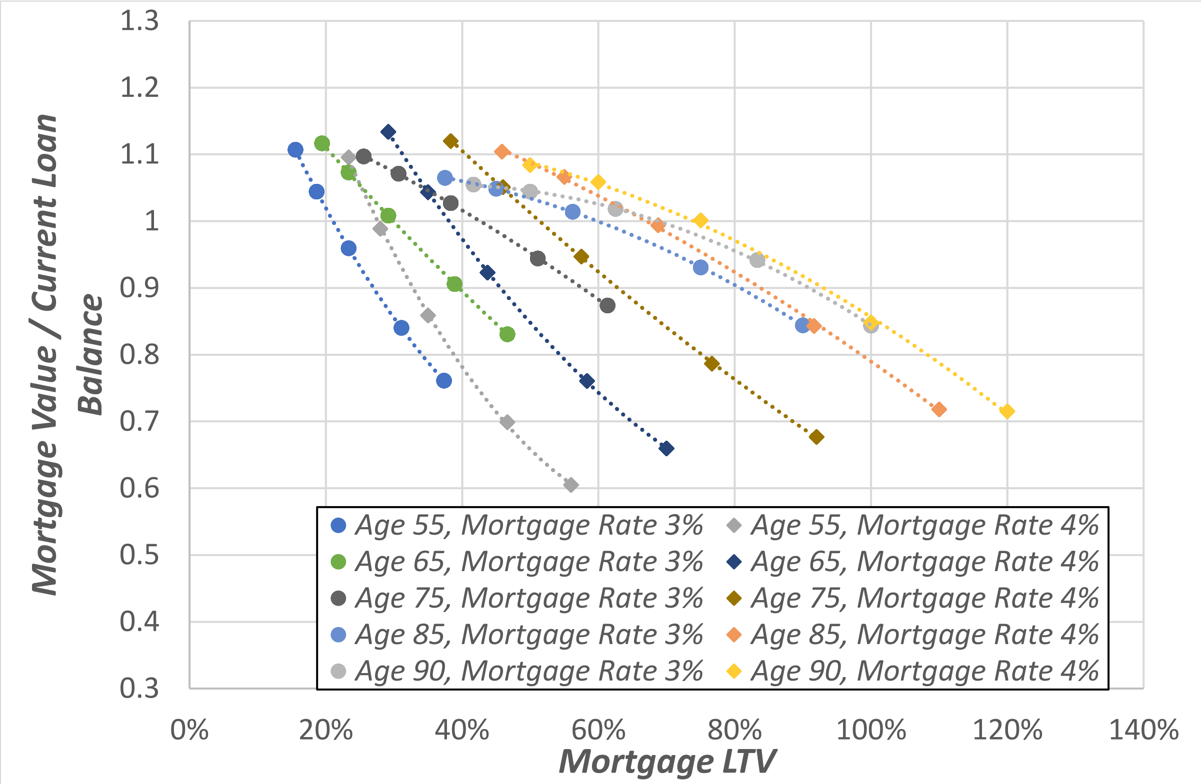

The reference ERM portfolio for which we could observe market prices only featured newly-originated mortgages. As a result, all the observed prices had a ratio of mortgage value-to-current loan balance of 1.04 (after allowing for the assumed 4% acquisition costs). The reference ERM prices directly determine, for example, that a 55-year old with a mortgage rate of 4% and an LTV of 25.5% will have a mortgage value of 1.04 relative to the current loan balance. But the reference portfolio does not directly tell us how this mortgage value will vary as the LTV moves to a higher or lower level than 25.5%. This is what the valuation model does.

The chart highlights how the relationship between mortgage value and LTV varies according to the valuation model. For example, for older lives, when the LTV is below 60% or so, further falls in LTV do not result in a significant increase in the mortgage value. This is because a 60% LTV mortgage with a 90-year old borrower is very low risk, and further reductions in LTV only make a very low risk mortgage even lower risk. For younger ages, however, a 60% LTV mortgage can be very risky, meaning that the mortgage value is very low and will increase significantly if LTV is reduced (i.e. if house prices increase).

The important point to note here is that we have used the patterns in the relative valuation of the ERMs – as identified by the implied volatility functions for their ERMlets – to infer these unobserved mortgage values. Naturally, increases in LTV result in a lower mortgage value and vice versa. But the chart highlights that the sensitivity of the mortgage value to a change in LTV varies significantly with age and mortgage rate and in non-linear ways.

This simple case study implementation of the valuation methodology described in Chapter 2 has shown how a valuation model can be calibrated to a set of reference ERM prices observed in the primary origination market. Once calibrated, the valuation model can then be applied to provide consistent and reasonable values for ERMs that do not have current observable prices.

In this case study, we considered a simple valuation model that was based on the assumption that house prices followed a geometric Brownian motion. This allowed us to use the Black-Scholes formula to value the NNEGs of the ERM's constituent ERMlets. To fit the various ERMs, we assumed that the ERMlets' Black-Scholes implied volatilities varied by the term and moneyness of the NNEG.

This analytical valuation approach makes the model very easy to use and interrogate. Valuations are calculated instantaneously, and the burdens of simulation run-times and sampling errors are avoided. This valuation method could be immediately applied to produce stressed valuations for the purposes of assessing the mortgage portfolio's capital requirements.

However, the Black-Scholes approach inevitably has its significant limitations. In particular, the model is unable to (easily) allow for dynamic prepayment behaviour. Whilst the degree of rational dynamic behaviour that may reasonably be expected of equity release mortgage borrowers may be open to much debate, the inability to allow for some dynamic prepayment behaviour is nonetheless a notable limitation.

The Black-Scholes modelling implementation is also limited in the extent to which it naturally extends to the valuation of securitisations. Our Black-Scholes implementation was easy because we effectively used a different valuation model for every ERMlet. But a securitisation depends on the *joint* behaviour of the valuation *paths* of a portfolio of ERMs. To model dynamic prepayment and the joint behaviour of a number of ERM assets through time, we may need a more flexible valuation model that uses stochastic simulation instead of analytical formulas. This does not affect the fundamental logic of the valuation methodology discussed in Chapter 2, but it requires a different implementation method to that developed here in Chapter 3.2. The implementation of such a modelling approach is developed next.

3.3 A Cashflow Simulation Model Implementation Example

This section develops a cashflow simulation-based implementation approach[10] of the valuation methodology presented in Chapter 2. Like the Black-Scholes implementation example above, it will use the valuation assumptions and observable ERM prices described in Chapter 3.1 to value the in-force portfolio presented in Chapter 3.1.7.

The simulation approach requires a different and more explicit approach to specifying the dynamics of house prices. In the Black-Scholes approach, the calibration process effectively specified the different probability distributions that could be assumed for house prices in order to value NNEGs with various terms and strikes. In the simulation approach, the model will explicitly specify a stochastic process for the path of the underlying house price. This path-modelling capability extends the potential use of the model. It allows us to capture the effect of potentially path-dependent mortgage features such as dynamic prepayment features. Moreover, securitisation valuation is also likely to be a path-dependent problem, even in the absence of dynamic prepayment behaviour (as the house price path may determine when cash is released from the securitisation, for example).

Of course, in the case where we assume house prices follow a geometric Brownian motion with constant volatility, the simulation model will produce identical valuations to the Black-Scholes implementation with uniform implied volatility (within the bounds of the sampling errors that are an inevitable feature of the simulation model valuation results). But to produce a good fit over the range of NNEG strikes, we will require a stochastic process for the paths of house prices that is explicitly different to geometric Brownian motion. The development and calibration of these house

[10] In all the following analysis produced by the cashflow simulation model, the model has been run with 10,000 simulations and annual time-steps.

price path dynamics will be our first task in the implementation of the simulation-based valuation approach.

3.3.1 A Stochastic Volatility Model

In our simulation-based approach, we begin by specifying a risk-neutral geometric Brownian motion for the house prices, implemented in discrete annual timesteps as follows:

$$\Delta S(t) = (f(t) - d)S(t) + \sigma S(t)z(t)$$

where S(t) and *d* are defined as above; z(t) is a random drawing from a standard normal variate and f(t) is the forward rate derived from the starting illiquid yield curve:

$$f(t) = [r(t) + \lambda(t)].t - [r(t-1) + \lambda(t-1)].(t-1)$$

So, for example, if the simulation model is run with the above house price stochastic process and σ = 14.25%, the same reference ERM valuations are obtained as in the Black-Scholes implementation with the uniform 14.25% volatility assumption (i.e. the results of Exhibit 3.6).

We found in the Black-Scholes analysis that the use of a uniform volatility assumption provided an inadequate fit to the ten reference ERM prices. The same conclusion will clearly apply to a simulation model that assumes a geometric Brownian motion with a constant volatility. As in the Black-Scholes implementation, a deterministic term structure can be specified for house price volatility so as to produce a better fit to ERMs of different borrower ages. In this case, the house price process can be specified as follows:

$$\Delta S(t) = (f(t) - d)S(t) + \sigma(t)S(t)z(t)$$

$$\sigma(t) = \sigma_0 + [1-\exp(-\theta t)]\sigma_{inf}$$

Note that the volatility function and its parameters now have a slightly different meaning to their Black-Scholes usage of Chapter 3.2. $\sigma(t)$ no longer refers to the volatility assumption that is used in the Black-Scholes option pricing formula for NNEGs of term *t*. Instead, it now refers to the volatility of the house price return at time *t*. But, just as in the Black-Scholes implementation, this volatility term structure can be calibrated to provide a very good fit to the reference ERM prices as a function of age.

However, once again as in the Black-Scholes model, this deterministic volatility term structure approach will not be able to provide a good fit to the reference ERM prices as they vary by LTV. To achieve this, we need a stochastic house price process that departs from geometric Brownian motion. In particular, we need something that will produce skew in the house probability distribution so that the left-hand tail is made fatter.

Stochastic volatility models provide a natural and well-trodden route to this objective. This valuation implementation will use a very simple form of stochastic volatility for the house price return. We will suppose that the house price volatility is a mean-reverting process with a lognormal distribution, and we will assume the (log) house price return, conditional on the level of volatility, is normal. We will also make the simplifying assumption that the variation in house price volatility is perfectly negatively correlated with the house price return[11].

[11] The assumption of perfect negative correlation between the house price return and the volatility means the volatility process is technically a deterministic function of the house price return, rather than a separate stochastic process.

Expressing this modelling specification in equations, the risk-neutral stochastic process for the house price can now be set out as follows:

$$\Delta S(t) = (f(t) - d)S(t) + V(t)S(t)z(t)$$
$$V(t) = \sigma(t)\exp(\text{Fat Tail Adjustment}(t))$$
$$\sigma(t) = \sigma_0 + [1-\exp(-\theta t)]\sigma_{inf}$$
$$\text{Fat Tail Adjustment } (t) = (1-\alpha)\text{Fat Tail Adjustment } (t-1) - \gamma z(t-1)$$

Fat Tail Adjustment (0) is assumed to be zero in this implementation.

There is nothing special or uniquely optimal about the above stochastic process for the house. It is one of many possible structures that may be able to deliver a good fit to observed ERM prices. Exhibit 3.13 shows the fits obtained by a calibration of the above model parameters to the prices of ERMs #1-#10. This calibration has assumed parameter values of σ_0 = 0.114; θ = 0.1; σ_{inf} = 0.1575; α = 0.25 and γ = 0.35.

Exhibit 3.13: Valuation Fits to the 10 ERM Prices (Stochastic Volatility Model)

ERM #	Age	Loan-to-Value	Mortgage rate	Observed Price[12]	Fitted Value	% Error
1	55	25.5%	4.0%	0.265	0.267	+0.7%
2	60	30%	4.0%	0.312	0.313	+0.4%
3	65	35%	4.0%	0.364	0.366	+0.5%
4	70	41%	4.0%	0.426	0.426	-0.1%
5	75	47.5%	4.0%	0.493	0.492	-0.4%
6	55	19.1%	3.0%	0.199	0.197	-1.2%
7	60	22.5%	3.0%	0.234	0.233	-0.6%
8	65	26.25%	3.0%	0.273	0.273	+0.0%
9	70	30.75%	3.0%	0.320	0.321	+0.3%
10	75	35.63%	3.0%	0.371	0.373	+0.6%

You can see that this calibration of the valuation model produces values within 1% or so of the ten ERM prices. The fitted prices are slightly less accurate than those that were produced by the Black-Scholes implementation with the 2-dimensional implied volatility 'surface' (see Exhibit 3.10, which showed the maximum valuation error for the 10 reference prices was 0.5%). This probably simply reflects the inadequacies of my elementary calibration fitting process rather than a more fundamental truth.

Now that we have obtained a satisfactory calibration of the simulation model to the reference ERM prices, we can use it to value the in-force portfolio of Chapter 3.1.7. A natural question arises: how do the values produced for the 1000 mortgages by this stochastic volatility simulation model compare with the mortgage valuations produced by the Black-Scholes implied volatility surface model (i.e. the results of Chapter 3.2.2)? Exhibit 3.14 shows the scatterplot of the valuations of each of the 1000 mortgages produced by the two different model implementations.

[12] Including the 4% allowance for acquisition costs.

Exhibit 3.14: Valuations of the 1000 mortgages in the example in-force portfolio

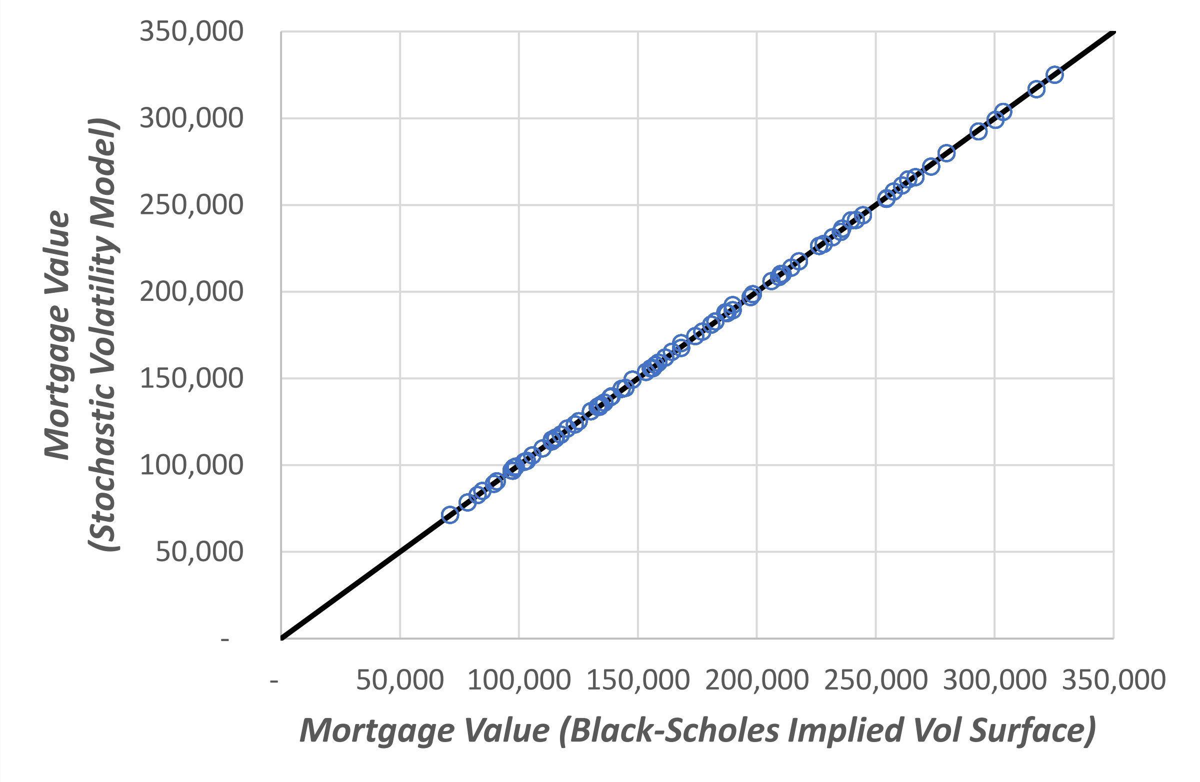

The chart highlights that the two models / calibrations produce very similar valuations for each of the mortgages in the in-force portfolio. We found in Chapter 3.2.2 that the Black-Scholes approach produced a portfolio valuation of £170.3m. The stochastic volatility model produces a valuation of £170.4m. This shows that the two model implementations, although quite different in approach, are producing very similar valuations when calibrated to the same set of reference prices.

It may also be interesting to consider the probability distribution of house prices that is implied by the stochastic volatility model calibration. This may sound like a very simple thing to do given we have explicitly stochastically modelled the value of houses when obtaining the ERM values plotted in Exhibit 3.14. However, a (very well-known) complication arises. To value the ERMs, we have only required risk-neutral house price probability distributions. We cannot derive a unique 'real-world' probability distribution from the ERM prices because the ERM values do not depend on the size of the house risk premium. So, any number of real-world distributions are consistent with those ERM prices and the assumed house price dynamics. Nonetheless, we can postulate a reasonable estimate of the house price risk premium – say, 3.5% - and examine the distribution that this assumption implies.

Exhibit 3.15 below shows the 10-year house price probability distributions produced by the simulation model. In the case of deterministic volatility, the α and γ parameters of the stochastic volatility process have been set to zero. In this case the distribution of the 10-year log house price return is normal with an annualised volatility of 13% (which is consistent with the 10-year ERMlet implied volatility with a lognormal house price, see Exhibit 3.8). In the stochastic volatility case produced by the calibration used in Exhibit 3.13, the left-hand tail has been substantially fattened. This feature has allowed us to value both the low LTV and high LTV ERMs with a single model calibration.

Exhibit 3.15: 10-year ERM-implied probability distribution for a house price

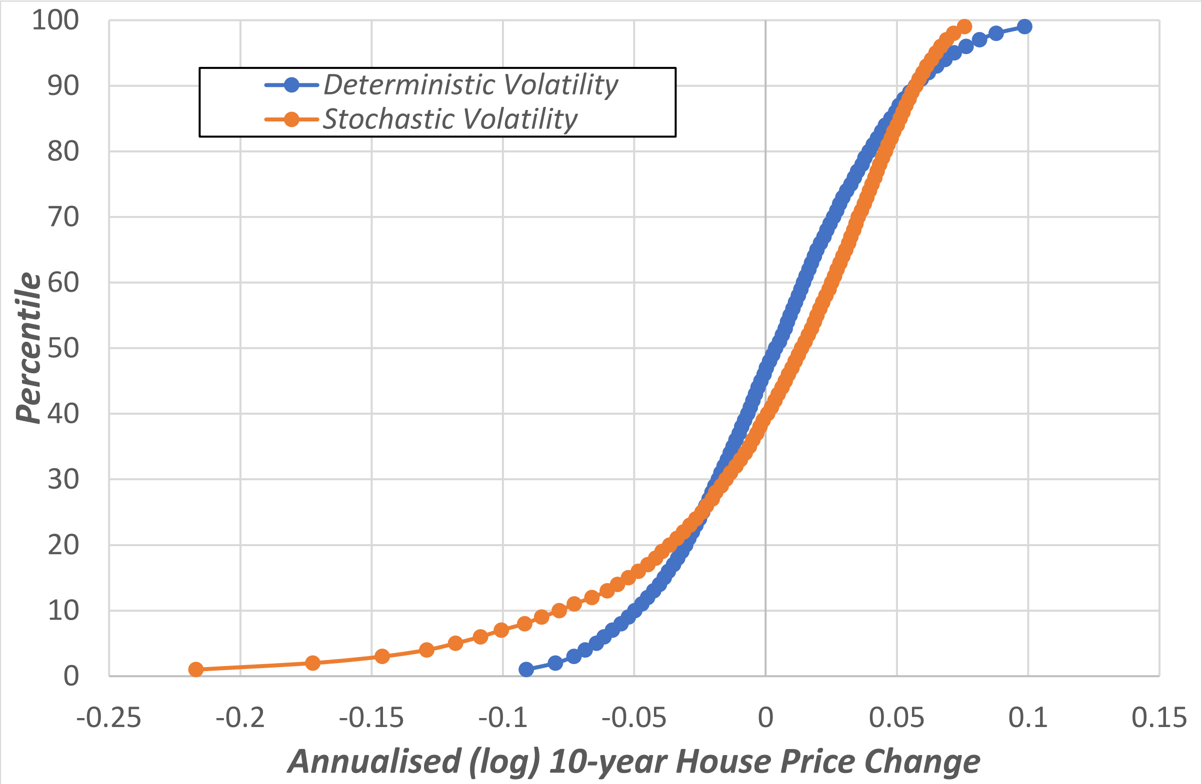

3.3.2 Dynamic prepayment modelling

Thus far, the cashflow simulation model has not taken us any further than the Black Scholes analysis of Chapter 3.2 – it has simply produced very similar valuation results when calibrated to the same prices with the same demographic and borrower behaviour assumptions.

But now we can consider an important aspect of ERM risk that we could not capture in the Black-Scholes framework: prepayment risk and, in particular, *dynamic* prepayment risk. By dynamic prepayment risk we mean a tendency for prepayment rates to increase (decrease) when the no-prepayment mortgage value is high (low) relative to the prepayment amount. Such a phenomenon will reduce the value of the mortgage (all else equal). This, in turn, will imply lower ERMlet volatilities for a given observable ERM price.

The calibration of dynamic prepayment behaviour is subject to a lot of judgement. 'Optimal' or 'rational' borrower prepayment behaviour can be approximated in a valuation model. But many would argue that the empirical prepayment behaviour of ERM borrowers suggests that their behaviour cannot reasonably be assumed to be fully optimal. So here we take a simpler approach. We specify a functional form for dynamic prepayment that tries to capture the essence of the dynamic selection effect: house price rises make the mortgage more valuable to the lender, and make the mortgage rate more expensive for the borrower (their new lower LTV could, in principle, allow the borrower to re-finance a new mortgage at a lower mortgage rate).

We will assume a 'base' prepayment rate of 1% for the first 10 years after origination of the mortgage, and 3% thereafter. We also assume that no prepayment occurs after the age of 80. We then specify that the actual prepayment rate for a given mortgage in year t is:

Prepayment rate (Mortgage, t) = Base prepayment rate (Mortgage, t) x

$$[(1 + \text{house price change since origination (Mortgage, t)})]^{\eta}$$

Again, this function is simply for illustration of the implementation. The valuation methodology described in this paper does not prescribe a particular form of prepayment assumption. In our implementation case study, we will assume that $\eta = 2$. So, if the house price has doubled between origination and year 10, the prepayment model says that the year-11 prepayment rate will increase by a factor of 4 from 3% to 12%. Whereas if the house price halved over the first 10 years of the mortgage, the year-11 prepayment rate would decrease by a factor of 4 from 3% to 0.75%.

The crucial point for our purposes of calibrating a valuation model is that a change in prepayment assumptions for the borrowers associated with the observed ERMs implies a change in the valuation model calibration (for a given set of prices). Exhibit 3.16 below shows how the mortgage values produced by the calibration of Exhibit 3.13 compare following the introduction of the dynamic prepayment model.

Exhibit 3.16: Valuation Fits to the 10 ERM Prices (Stochastic Volatility Model)

ERM #	Age	Loan-to-Value	Mortgage rate	Observed Price[13]	Fitted Value (no prepayments)	Value (with prepayments)
1	55	25.5%	4.0%	0.265	0.267	0.273
2	60	30%	4.0%	0.312	0.313	0.318
3	65	35%	4.0%	0.364	0.366	0.368
4	70	41%	4.0%	0.426	0.426	0.427
5	75	47.5%	4.0%	0.493	0.492	0.492
6	55	19.1%	3.0%	0.199	0.197	0.197
7	60	22.5%	3.0%	0.234	0.233	0.233
8	65	26.25%	3.0%	0.273	0.273	0.273
9	70	30.75%	3.0%	0.320	0.321	0.321
10	75	35.63%	3.0%	0.371	0.373	0.373

As noted in Chapter 3.1.5, the assumption of a *deterministic* prepayment rate will *increase* mortgage values. It has a particularly pronounced impact on the values of mortgages written with younger borrowers. For example, in the absence of the *dynamic* prepayment model (i.e. setting $\eta = 0$ above), ERM #1's value is increased by 17% to 0.309 by the presence of the base prepayment assumptions described above. However, the dynamic selection effect that we have modelled largely offsets this effect, so much so that the previously calibrated volatility model is still producing a very good fit to the reference prices. So, in this specific implementation, a re-calibration of the stochastic volatility model is not required by these changes in prepayment assumptions, but we would not expect this to generally be the case.

Although the valuations have not been changed materially by the introduction of prepayment, the sensitivities of the valuation may well have changed significantly (especially to interest rate risk).

Now that the model has been satisfactorily calibrated, we next consider the valuation of the 1000-mortgage in-force ERM portfolio in the presence of the assumed prepayment behaviour described above.

[13] Including 4% allowance for acquisition costs.

3.3.3 Valuing the ERM Portfolio

We noted in the above section that the stochastic volatility model produced a portfolio valuation in the absence of any prepayments of £170.4m. When we move from the no-prepayment to with-prepayment basis, the total portfolio value increases by around 1.7% from £170.4m to £173.2m. Exhibit 3.17 re-produces the valuation results analysis that was provided in Chapter 3.2.2 (Exhibit 3.12).

Exhibit 3.17: Valuation Results – with dynamic prepayment model

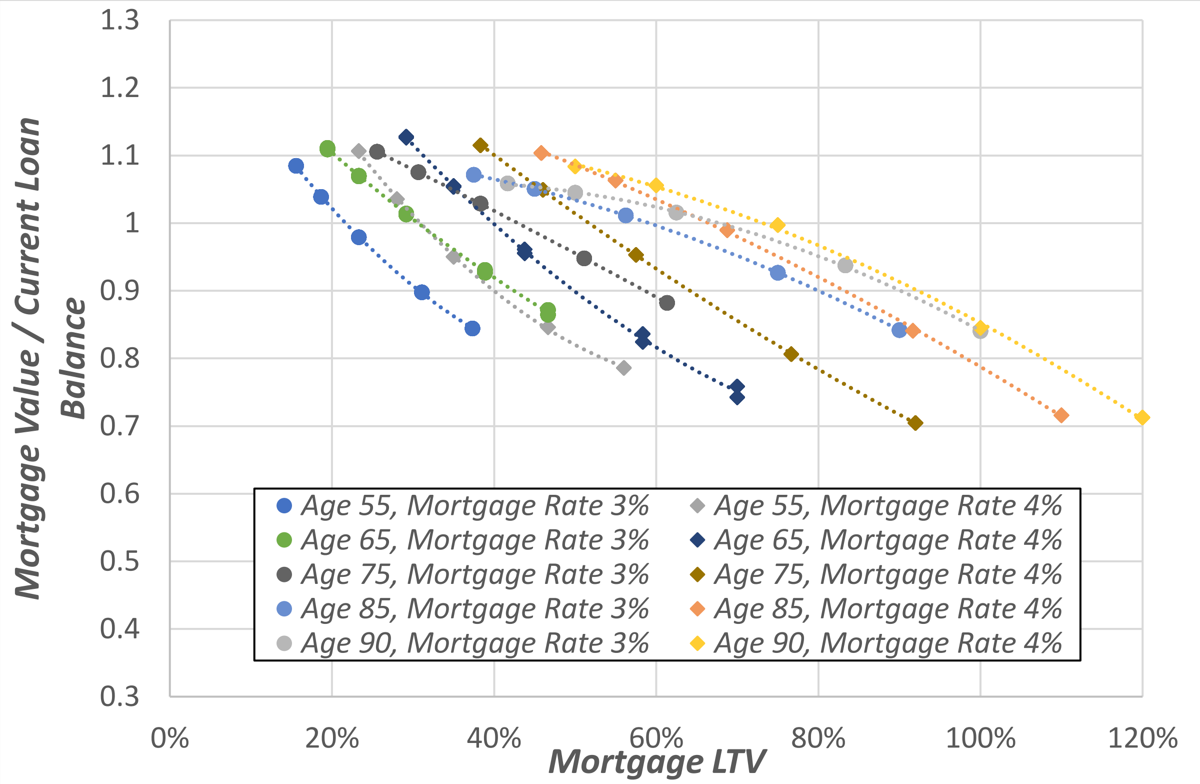

A comparison of Exhibits 3.12 and 3.17 shows that the presence of prepayment has tended to increase the value of the high LTV mortgages. This is because the underlying deterministic base prepayment rate has a very beneficial impact on those mortgage values.

This concludes our implementations of the ERM mortgage valuation methodology of Chapter 2. The case studies of Chapter 3 have shown that two quite different implementations of the methodology can produce similar qualities of fit to a set of reference ERM prices and can result in similar valuations of a portfolio of in-force mortgages. However, the simulation model has greater flexibility and can capture important ERM features such as dynamic prepayment behaviour. It is also more suited to the task of securitisation valuation. For these reasons, we will use the simulation model and calibration of Chapter 3.3 for the remainder of this paper.

4 Assessing Capital Requirements of Equity Release Mortgages

Chapter 4 considers the assessment of the 1000-mortgage ERM portfolio's capital requirements under a 1-year 99.5% Value-at-Risk (VaR) definition of capital. If using the Black-Scholes valuation method of Chapter 3.2, the 1-year VaR capital requirement can be assessed by re-calculating the analytical mortgage valuation functions under any number of 1-year scenarios. The analytical nature of the valuations makes the computational implementation of the capital assessment very straightforward. However, as discussed in Chapter 3, such an approach is unable to capture dynamic prepayment effects. We therefore consider the capital requirement that arise from the use of the cashflow simulation valuation method in the presence of dynamic prepayments, as described in Chapter 3.3.

The assessment of the capital requirement under the cashflow simulation method is a more computationally demanding implementation challenge. Every re-valuation of the portfolio requires another run of the simulation model, and simulation runtimes can make the full 'nested' implementation approach impractical. Of course, this is not a new type of problem, and European insurance firms have developed sophisticated quantitative techniques that can allow them to develop 'proxy functions' for the 1-year change in valuation that arises in 1-year VaR implementations. These proxy functions are usually applied to complex insurance liability types such as with-profit guarantees, but the same concept can be applied to the asset side of the balance sheet.

Whilst developing a proxy function for the ERM portfolio's 1-year valuation change would doubtless be an interesting exercise, it would be a digression from the central purpose of this paper. We therefore use a simpler implementation method for the illustrative purposes of our current exercise: we use the more basic 'stress-and-correlate' approach. Under this approach, the portfolio is re-valued under a 99.5th percentile stress for each of the risk factors under consideration. So, this approach typically only requires a handful or so of re-valuations rather than many thousands. We then use some correlation assumptions to aggregate these individual capital requirements into a total aggregate capital requirement for the portfolio.

Seven risk factors naturally arise from our implementation of the cashflow simulation valuation method in Chapter 3.3:

- 1-year change in house prices
- 1-year change in the liquid risk-free term structure
- 1-year change in the illiquidity premium term structure
- 1-year change in the house deferment rates
- 1-year change in assumed future mortality and / or long-term care entry rates
- 1-year change in assumed future behaviour of prepayment
- 1-year change in ERM-implied house price volatility (i.e. changes in the house price volatility implied by the observed prices of equity release mortgages).

Given the valuation method, it is straightforward to re-value the portfolio under stresses to each of these risk factors. Below we consider each of these risk factors in turn. In each case we will postulate an illustrative 99.5% 1 -year risk factor stress, and then re-value the 1000-mortgage portfolio under this stress (for simplicity, we assume the stress and re-valuation is instantaneous; that is, we do not roll the portfolio forward one year). In all cases, the specified size of the risk factor stress is simply illustrative. In a 'real' VaR capital implementation, the calibration and validation of these stresses would be a major element of the capital calculation. It is the part of the capital assessment process that inevitably involves the most judgment.

4.1 A Fall in House Prices

The ERM portfolio is naturally exposed to falls rather than rises in house prices. We will assume a 30% house price fall as the 1-year 99.5th percentile stress, i.e. we assume every house loses 30% of its value. In a well-diversified portfolio, we would not anticipate that all house prices would move in perfect correlation, so the uniform stress is a simplifying assumption.

The 30% house price stress reduces the 1000-mortgage portfolio value by 16% from £173.2m to £144.7m according to the Chapter 3.3 valuation model. The portfolio's houses have an aggregate starting value of £410m and a 30% property fall implies a fall in the portfolio's house values of £123m to £287m. So, the mortgage value's 'delta' over the 30% fall can be calculated as (£144.7m - £173.2m)/(£287m-£410m) = +0.232.

Naturally, not all mortgages exhibit the same sensitivity to house price falls – this sensitivity varies from 3% to 30% across the mortgages in the portfolio. Exhibit 4.1 below shows how the capital requirement of each individual mortgage varies as a function of LTV for mortgages of various mortgage rates and borrower ages.

Exhibit 4.1: Valuation Impact of the 30% House Price Stress

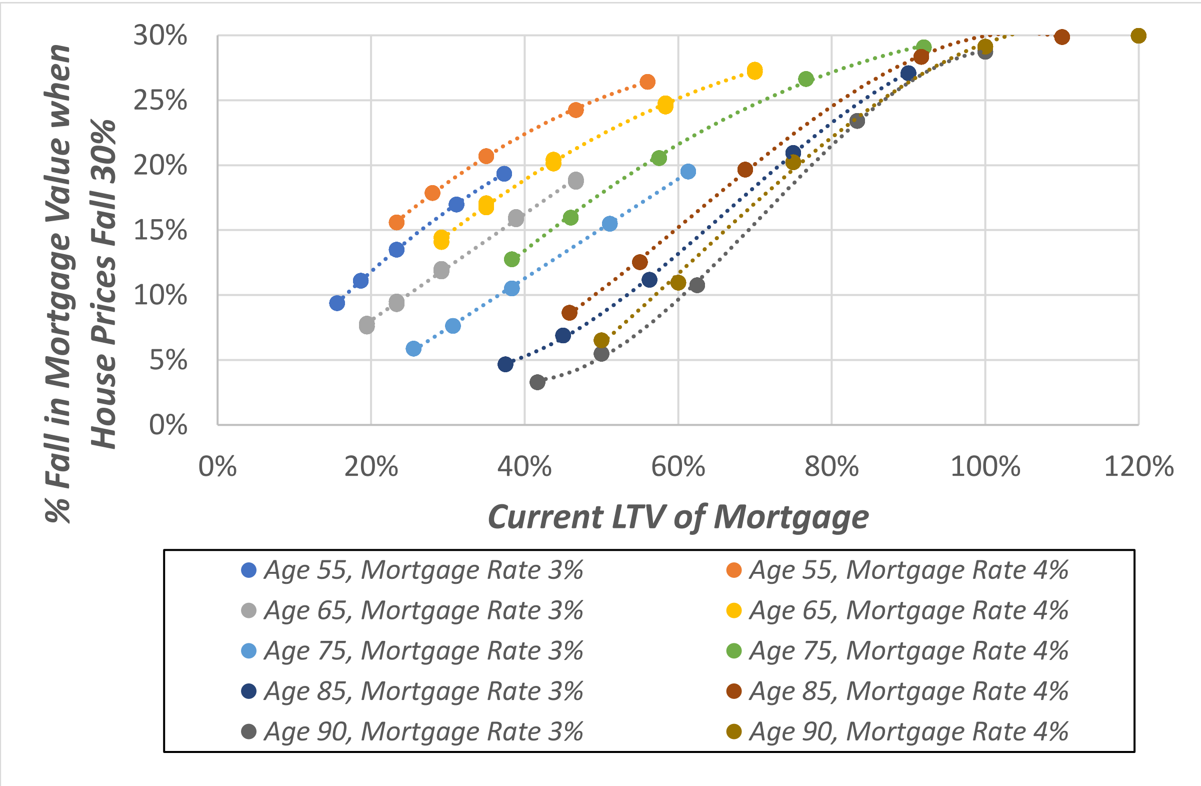

Unsurprisingly, the greatest variation in mortgage value sensitivity to LTV occurs at older ages – at age 90%, the short expected remaining lifetime makes a mortgage with an LTV of 40% particularly safe, whereas an LTV of 120% has a high likelihood of its NNEG finishing 'in-the-money'. These differences in 'moneyness' result in different NNEG deltas just as we would expect of any (short) put option.

4.2 An Increase in the Liquid Risk-free Rate

The ERM portfolio value is exposed to rises rather than falls in the risk-free term structure. For our illustrative purposes, we assume a 1-year 99.5th percentile upward stress of 200 basis points across the liquid risk-free term structure.

A 2.00% increase in the liquid risk-free yield curve reduces the portfolio value by 13% from £173.2m to £151.1m.

Exhibit 4.2 below shows how the individual mortgage values fall under this stress.

Exhibit 4.2: Valuation Impact of the 2% Risk-free Yield Curve Up Stress

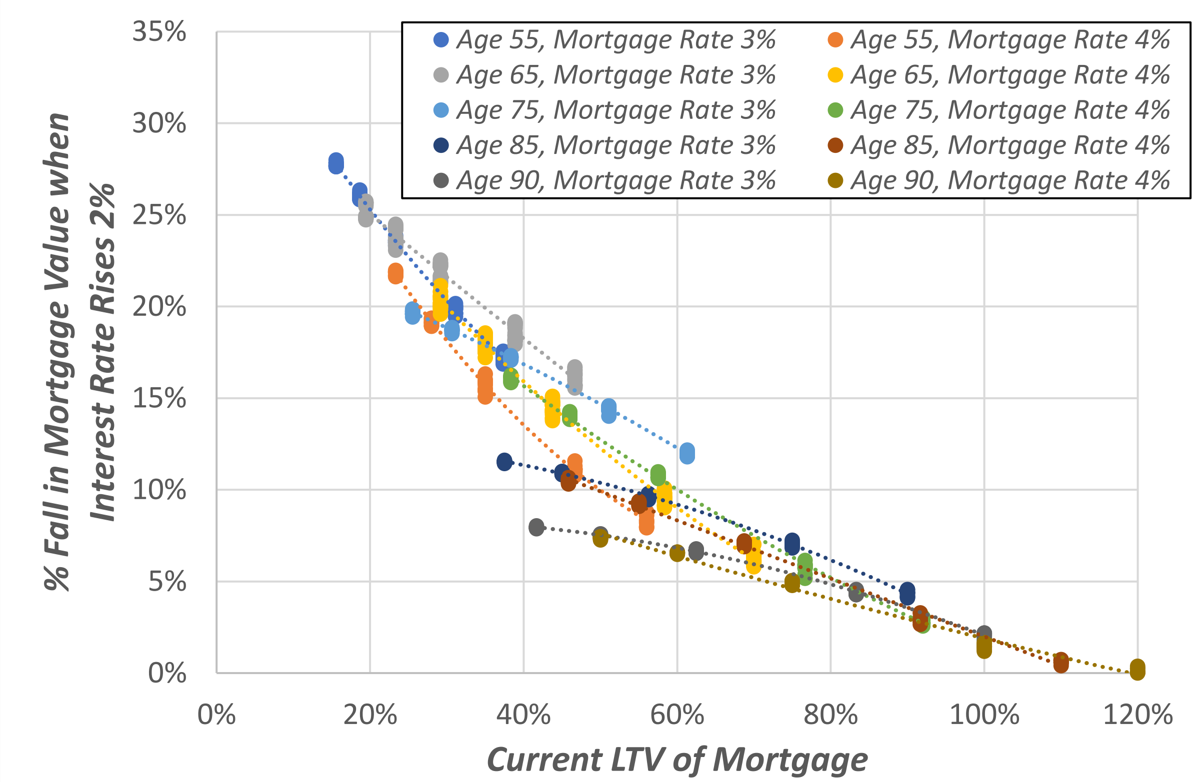

As we would expect, the mortgages with younger borrowers exhibit a greater interest rate sensitivity than the mortgages with older borrowers. The chart also shows that the valuation model produces greater interest rate sensitivity for lower LTV levels. This is also an intuitive result – as LTV falls, the loan more closely resembles a risk-free bond rather than a house, and it is therefore more sensitive to changes in risk-free bond prices.

Finally, you may find the overall levels of interest rate sensitivity surprisingly low given the long-term nature of these mortgages. The interest rate sensitivities above are significantly lower than would be implied by the Macaulay duration of the mortgages' expected cashflows. This effect arises because the mortgages' NNEG values are highly sensitive to interest rate changes, and, as the mortgage has a negative exposure to the NNEG, this 'dampens' the total interest rate sensitivity of the mortgage[14].

[14] For further discussion of this effect, see my article https://craigturnbullfia.com/prepayment-risk-and-equity-release-mortgages/

4.3 An Increase in the Illiquidity Premium

The case study has assumed the mortgages are illiquid. The valuation methodology of Chapter 2 does not require us to make this assumption, but if we do make this assumption, then a change in the illiquidity premium has exactly the same valuation impact as a change in the liquid risk-free rate. The relevant risk-free rate for the illiquid mortgage is the illiquid risk-free rate. The valuation is not affected by how that illiquid risk-free rate is comprised of the liquid risk-free rate and the illiquidity premium.

So, the ERM portfolio value is exposed to an increase in the illiquidity premium. For our illustrative purposes, we assume a 1-year 99.5th percentile illiquidity premium stress of a 100 basis points increase at all points in the term structure. The mortgage portfolio value falls by 6% from £173.2m to £162.2m in this stress.

4.4 An Increase in the Deferment Rate

The ERM portfolio value is reduced by an increase in the deferment rate. Our base valuation assumption is that all houses have a deferment rate of 3.5%. We assume the 99.5th percentile 1-year deferment rate stress is a 1.5% increase to 5.0%.

The 1.5% increase in the deferment rate of all mortgages reduces the portfolio value by 10% from £173.2m to £155.5m.

Exhibit 4.3 shows how the values of the individual mortgage holdings fall under this stress. The long duration mortgages exhibit a greater sensitivity than short duration mortgages to increases in the deferment rate; and, for a given age, deferment rate sensitivity is an increasing function of LTV (as the ERM increasingly resembles the deferred possession of the house rather than the (illiquid) risk-free bond, and its valuation is therefore more directly impacted by a change in the deferment rate).

Exhibit 4.3: Valuation Impact of the 1.5% Deferment Rate Up Stress

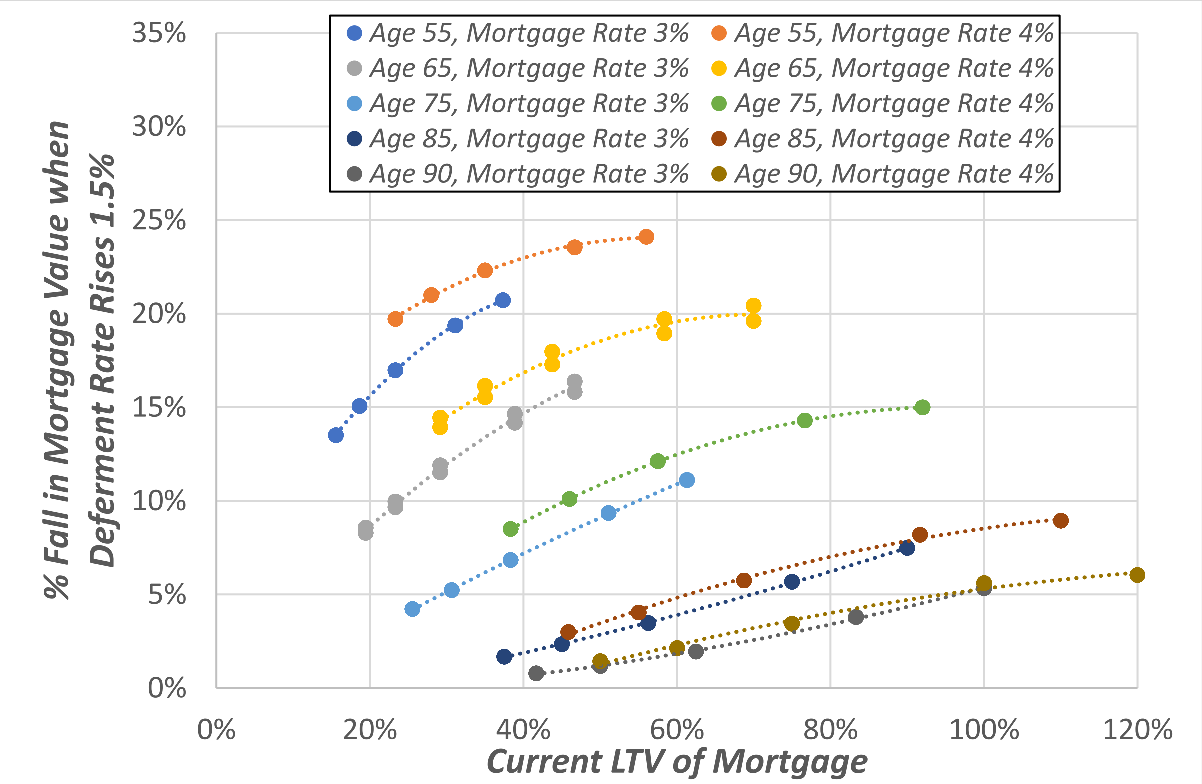

4.5 A Fall in Mortality and LTC Entry Rates

The ERM portfolio value is, in aggregate, exposed to a fall in mortality rates. The 99.5th percentile stress assumes mortality and LTC entry rates decrease by a proportion of 10% at all ages in all future years.

Unlike the other stresses considered in our capital assessment, however, not all mortgages have the same directional sensitivity to this risk factor – the valuation sensitivity to the fall in mortality rates varies between 1.7% and -0.1%. In particular, mortgages with borrower age 90 and LTVs of less than 55% generate a (very small) *increase* in value when mortality rates fall. These loans are very safe, even if their term is lengthened; the mortgage rate is more than enough compensation for their risk, and so an extension of the term of the loan is value-enhancing, even though it increases the NNEG value.

The 10% reduction in mortality / LTC entry rates reduces the portfolio value by 1% from £173.2m to £171.7m.

4.6 A Fall in Base Prepayment rates

Borrower behaviour can deviate from that assumed by our prepayment model in a couple of distinct ways: the deterministic element of the base rates could be lower than assumed in the valuation; the degree of dynamism in the prepayment could be greater than assumed in the valuation. The latter, in particular, is a very difficult stress to calibrate. Here we will simply consider a downward stress to the deterministic element of the prepayment behaviour. We assume the 99.5th percentile 1-year stress to the deterministic prepayment rate is a 50% proportional reduction at all terms of mortgage (so the base prepayment rate over the first 10 years of the mortgage is reduced from 1.0% to 0.5%; and from 3.0% to 1.5% thereafter, all subject to the assumption that no prepayments take place beyond age 80).

The valuation impact of this stress is small - it reduces the portfolio value by 1% from £173.2m to £171.9m. It is worth noting the important interaction between the impact of base prepayment rates and the assumed dynamic prepayment behaviour: the valuation impact of the base prepayment rate stress would be materially greater in the absence of dynamic prepayment behaviour.

4.7 An Increase in ERM-Implied House Price Volatilities

It is quite possible that the observable prices of the reference ERMs will change for reasons that cannot be explained by changes in the above risk factors. In the valuation methodology of Chapter 2, this would manifest itself as a change in the house price volatility calibration of the valuation model. The ERM portfolio value is exposed to increases in ERM-implied volatility.

This is a difficult stress to calibrate. We could consider a times series of empirical ERM prices and derive the ERM-implied hose price volatility calibrations for each historical date. This would give some indication of the historical variation in the house price implied volatility. But the limited length of available historical time series for ERM prices and the sensitivity of the results to how other valuation assumptions have changed over time would likely limit the usefulness of such an exercise. A forward-looking perspective could consider how much an estimate of future long-term house price volatility could change by over one year. Again, this is no easy task. This is a general difficulty that arises from applying a 1-year VaR capital framework to illiquid assets, rather than a specific issue related to the methodology proposed of this paper.

For our illustrative purposes, we assume the ERM-implied volatilities increase by 5% at all terms and strikes[15]. This implied volatility stress reduces the portfolio value by 6% from £173.2m to £162.3m.

Exhibit 4.4 shows how the stress impacts on the individual mortgage values. The long duration mortgages are impacted more significantly by the volatility stress – the mortgages with 55-year old borrowers fall by 6% to 12% in value whereas the mortgages with 90-year old borrowers fall by 1% to 4%.

Exhibit 4.4: Valuation Impact of the 5% Volatility Up Stress

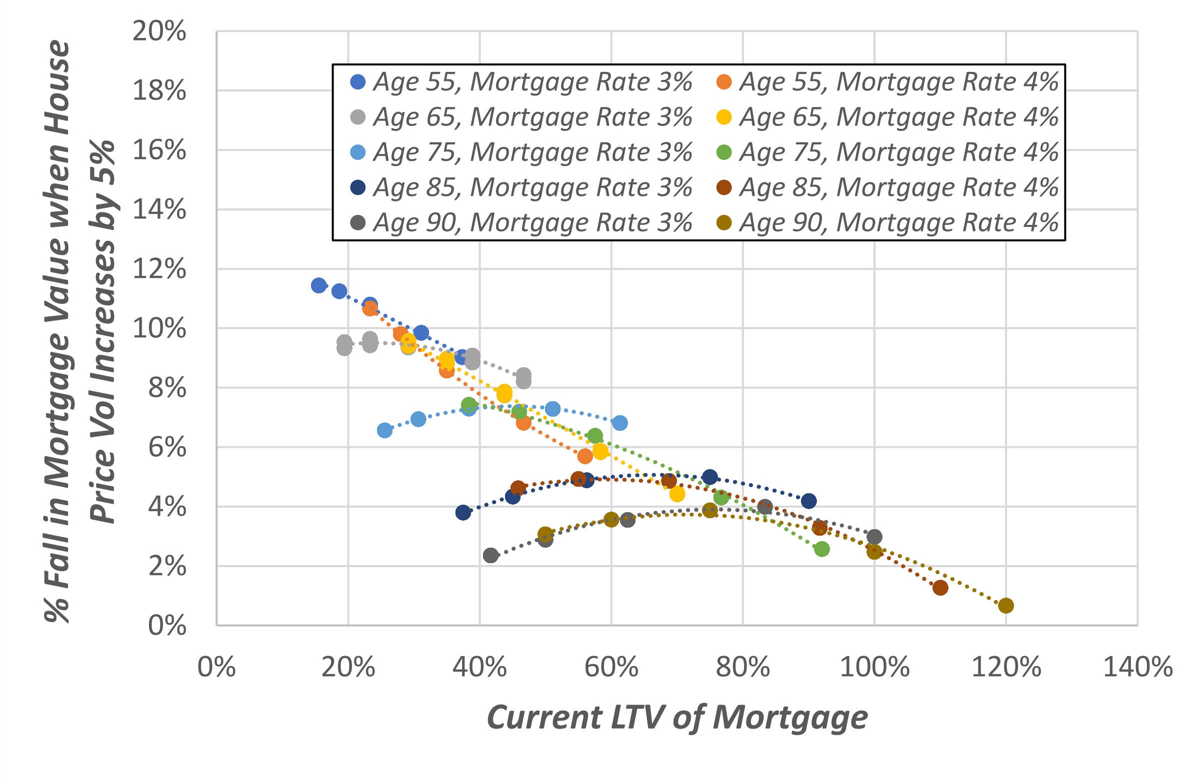

4.8 Risk aggregation

The individual capital stresses presented above are summarised in the table below.

Exhibit 4.5: Capital Requirements by Risk Factor

Risk factor stress	Capital Requirement
30% House Price Fall	16%
2% Liquid Risk-free Rate Rise	13%
1% Illiquidity Premium Rise	6%
1.5% Deferment Rate Rise	10%
5% House Price Vol Rise	6%
10% Mortality / LTC Rate Fall	1%
50% Fall in Base Prepayment Rate	1%

[15] More specifically, we assume the σ_0 and σ_{inf} parameters of the stochastic volatility model each increase by 0.05. There is a good argument for short-term implied volatility increasing by more than long-term volatility and this could be implemented by only stressing the σ_0 parameter. But for simplicity here we simply assume a parallel shift in the whole term structure of implied house price volatility.

If no allowance were made for diversification amongst the risk factors, the aggregate capital requirement would be 54%. We may expect some of these risk factors to have high correlation. For example, asset prices and asset volatility are usually strongly negatively correlated, which implies the house price fall stress and house price volatility rise stresses are strongly positively correlated. But some of these stresses may be expected to have low or even negative correlations. This could result in a material diversification benefit arising in the capital aggregation.

The joint behaviour of the deferment rate with house prices and interest rates is perhaps the most interesting and important relationship for the capital aggregation. Different economic arguments may be offered that imply quite different joint behaviours amongst these variables. The joint behaviour of the deferment rate with interest rates and house prices depends on how the deferment rate is interpreted and estimated. We have noted earlier in this paper that long-term house deferment rates tend not to be readily implied from observable market prices. In the valuation methodology set out above, it has been argued that, in the absence of such prices, the deferment rate is most naturally defined as the net rental yield on the house (with a deduction of the house illiquidity premium if the mortgage is a liquid asset; in all analysis in this report we assume the mortgage is illiquid).

Given this interpretation of the deferment rate as a net rental yield, the joint behaviour of the deferment rate and house prices has clear implications for the behaviour of rental income and vice versa. If rental income is assumed to be stable, a downward move in house prices would tend to result in net rental yields moving upwards, and vice versa. The correlation would not be perfectly negative because rental income will also vary. But the assumption that house rental income levels are stable would imply the correlation between prices and yield is strongly negative, perhaps -0.9. This implies a strongly positive correlation between the house price fall stress and the deferment rate rise stress.

On the other hand, if we assume that rental income is highly variable and moves in lockstep with house prices, this implies the net rental yield is constant. In this case, the deferment rate stress disappears by assumption – any change in house price is associated with a corresponding change in the level of house rental income, and the yield remains constant.

But not everyone works with a net rental yield interpretation of the deferment rate. For example, the PRA's SS3/17 and Effective Value Test clearly does not use the net rental yield as the basis for the parameterisation of the deferment rate. Under this alternative interpretation, the deferment rate is assumed to be strongly correlated with real interest rates. This alternative correlation has important implications not only for capital, but also for the implied duration of the mortgages. If we assume the deferment rate moves in perfect correlation with interest rates, the implied interest rate sensitivity (duration) of the mortgages is made much longer.

So, the above discussion identifies three different sets of correlations between the deferment rate and house prices and interest rates that could be argued to be reasonable:

1. Deferment rate as net rental yield, assuming rental income stable.
2. Deferment rate as net rental yield, assuming rental income volatile.
3. Deferment rate moving in lockstep with interest rates[16].

[16] This illustrative implementation model has not explicitly modelled changes in real interest rates, and so we assume here that deferment rate moves in lockstep with the nominal interest rate. The analysis could be further enriched by modelling both nominal and real interest rates as correlated risk factors, without fundamentally altering the logic of any of the analysis.

Let's analyse what these alternative perspectives imply for the mortgage portfolio's aggregate capital requirement. For simplicity, here we will assume the house price fall is perfectly positively correlated with the house price volatility rise; and that the liquid risk-free rate, illiquidity premium, mortality / LTC entry rate fall stress and prepayment rate fall stress are each independent of all other risk factors.

1. Deferment rate as net rental yield, assuming rental income stable

In this case, the house price fall and deferment rate rise are perfectly correlated (as a limiting case). This implies[17] an aggregate capital requirement of **36%.**

2. Deferment rate as net rental yield, assuming rental income volatile

Here the deferment rate stress disappears, as the deferment rate is assumed to be constant (again, as a limiting case). All other relationships are as in case 1. This implies[18] an aggregate capital requirement of **27%**.

3. Deferment rate moving in lockstep with interest rate

In this case, the deferment rate rise stress and the interest rate rise stress are perfectly correlated. This implies[19] an aggregate capital requirement of **33%**.

Finally, if the mortgages are held to back long duration liabilities, it may be more meaningful to consider the capital requirements with the interest rate capital requirement's contribution excluded. In this case, the capital requirement in the three cases are **34%**, **24%** and **24%** respectively[20]. (Note that cases 2 and 3 are identical when interest rate risk capital is excluded, though the duration of the implied matching liability will be materially different in the two cases).

The aggregated capital requirements are summarised in the table below.

Exhibit 4.6: Aggregated Capital Requirements

Capital aggregation method	With Interest Rate Risk	Without Interest Rate Risk
Case 1 (stable rental income)	36%	34%
Case 2 (stable rental yield)	27%	24%
Case 3 (deferment rate with fixed relationship with interest rate)	33%	24%

Clearly, the difference between a capital requirement of 24% and 34% is highly material. The interpretation of the deferment rate and its consequent implied behaviour under stress is a crucial element of the assessment of the mortgage portfolio's aggregate capital requirement.

[17] i.e. sqrt $[(16\%+10\%+6\%)^2+13\%^2+6\%^2+1\%^2+1\%^2]$
[18] i.e. sqrt $[(16\%+6\%)^2+13\%^2+6\%^2+1\%^2+1\%^2]$
[19] i.e. sqrt $[(16\%+6\%)^2+(13\%+10\%)^2+6\%^2+1\%^2+1\%^2]$
[20] Case 1: sqrt $[(16\%+10\%+6\%)^2+6\%^2+1\%^2+1\%^2]$; Case 2: sqrt $[(16\%+6\%)^2+6\%^2+1\%^2+1\%^2]$; Case 3: sqrt $[(16\%+6\%)^2+6\%^2+1\%^2+1\%^2]$

5 A Valuation Methodology for ERM Securitisations

Chapter 3 presented two distinct implementations of the valuation methodology of Chapter 2, one using Black-Scholes and the other a cashflow simulation model. Whilst it was shown that the two approaches could produce very similar portfolio valuations under the same demographic and borrower behaviour assumptions, one of the main advantages of the simulation approach is that it can be more easily extended to the valuation of mortgage portfolio securitisations. This chapter will show how the cashflow simulation valuation model can be applied to the valuation of securitisations of the mortgage portfolio.

The valuation method for securitisation tranches is perhaps the topic where there is greatest distance between the approach proposed here and current industry approaches. As noted in the Introduction, there currently tends to be a quite limited connection between the industry valuation methods for mortgages and their securitisations. The securitisations are a well-defined form of derivative security of the mortgage portfolio. The valuation behaviour of the underlying portfolio therefore has some very direct and natural implications for the behaviour of each tranche of a portfolio securitisation. But this linkage tends to currently only be recognised in terms of *ex post* valuation adjustments that are necessary to recover the equation of value. As this chapter will demonstrate, an explicit appreciation of the economic linkage between the value of the mortgage portfolio and the value of the securitisation tranches removes the necessity for any such adjustments.

5.1 Mortgage Portfolio Cashflows

A preliminary analysis of the cashflows produced by the mortgage portfolio is useful for contextualising the behaviour of a securitisation.

The stochastic cashflow simulation model of Chapter 3.3 can provide a probability distribution for the cashflows produced by the mortgage portfolio in every future year. However, as per our discussion of the house price probability distribution of Chapter 3.3.1, we should note that our model has not been designed or calibrated with the objective of producing a realistic distribution of the mortgage cashflows – it has been used to produce current values for the mortgages (and the portfolio). To obtain a probability distribution from the model, we need to make assumptions about the risk premia associated with the model's risk factors. The model has only one fundamental risk factor: the house price return[21].

For the purposes of producing mortgage cashflow probability distributions, we assume all house assets have a risk premium of 3.5% (as we did when producing the house price probability distribution in Exhibit 3.15). It should, however, be noted that this does not necessarily produce a fully realistic distribution of house prices or mortgage cashflows. We have not concerned ourselves with many of the empirical features of house prices that may be important in cashflow projection (but less important for current valuation). We also have not considered uncertainty in risks such as mortality and long-term care rates that can impact on the distribution of cashflows.

When using the model's valuation calibration to produce probability distributions, we also make the implicit assumption that the option-implied volatility behaviour that has been derived from ERM prices provides a realistic description of future house price volatility. If we believed that various costs and risks impact on ERM prices in a way that is not explicitly captured by the valuation model, then we might expect the derived option-implied volatility to include a 'loading' that reflects those

[21] Whilst we have a stochastic volatility process, it is perfectly correlated with the house price return, and so does not require any additional market price of risk assumptions.

costs and risks. In this case our best estimate of house price volatility may be lower than the derived ERM-implied volatility behaviour.

With these caveats borne in mind, the mortgage cashflow probability distributions produced by the simulation model with a 3.5% house risk premium can nonetheless provide a useful indication of the potential future behaviour of the mortgage cashflows. There is one final complication that we must consider, and which will be an important topic in the treatment of securitisations: the modelling of the probability distribution of the mortgage portfolio cashflows requires an assumption about the *correlations between house prices*.

This correlation question did not arise in Chapter 3's mortgage portfolio valuation implementations because the portfolio value is simply the sum of the individual mortgage values. These individual mortgage values do not depend on the correlation between the different mortgages. The question could have arisen in Chapter 4 when the portfolio house price stress was specified, but we effectively removed the issue by simply assuming the same house price stress applied to every house in the portfolio.

In the securitisation analysis, we will make the assumption that all houses in the portfolio have a correlation with each other of +0.55. This is consistent with a simple one-factor index model of house returns, where the index return and house-specific return have similar levels of volatility[22]. For the purposes of this illustration, we are using a very simple correlation structure .There is no methodological barrier to applying a more sophisticated, multi-factor house price correlation structure that uses factors such as property type, location, property size and so on to describe the correlation relationships between the houses of the mortgages in the portfolio.

Finally, we should note that, as securitisation tranche values will depend on the assumed level of house price correlation, observed prices of ERM securitisation tranches would be a natural source of information on market-implied house price correlations. But in the absence of observable securitisation market prices, we must use an economic approach to infer the correlations that are used in the securitisation valuation model.

The following chart shows the cashflow probability distributions for the mortgage portfolio generated by the cashflow simulation model of Chapter 3.3 with a house risk premium of 3.5% and a house-to-house correlation of +0.55 for all houses in the portfolio.

[22] For example, if index volatility is 12% and house-specific volatility is 10.75%, the house-to-house correlation is $12\%^2/(12\%^2+10.75\%^2) = +0.55$.

Exhibit 5.1: Projected Probability Distributions of the Annual Mortgage Cashflows

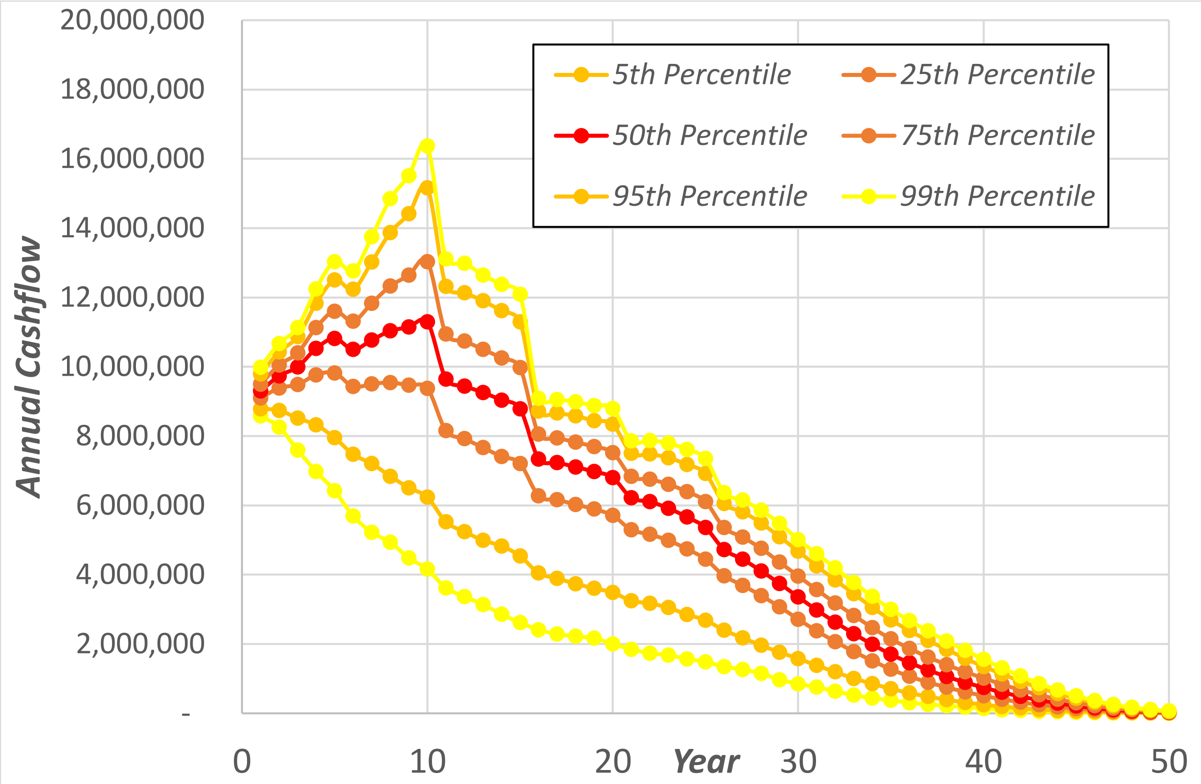

This chart provides an illustration of the overall profile of the cashflow distribution and the uncertainty associated with it. Note the chart is plotting the percentiles of the probability distribution for each individual year – it is not plotting the *paths* of outcomes. There may be some complex interactions between the cashflow outcomes from year to year. For example, strong house market returns in the early years of the projection will reduce NNEG losses and this will have a positive impact on cashflows; it will also increase dynamic prepayment rates, which will increase short-term cashflows but reduce those arising at a later date (this effect explains the sharp drops in the right-hand tails of the cashflows between years 10 and 20).

Recall from above that the mortgage portfolio has been valued at £173.2m. We can calculate various yield and return statistics from this valuation and the above cashflow distributions. For example, we can take the expected values of the mortgage portfolio cashflows projected for every year and find the discount rate that equates the present value of those cashflows with the portfolio value of £173.2m. If we apply a discount rate to the cashflow in year *t* of the current liquid risk-free *t*-year spot rate plus a spread *x*, we find a value for *x* of **1.85%** equates the present value of the cashflows with the portfolio valuation of £173.2m. **This provides a measure of the mortgage portfolio's expected return in excess of the liquid risk-free rate.**

Let's consider the size of this estimate of the mortgage expected return. The cashflow projection has assumed that houses offer a total return in excess of the liquid risk-free rate of 4.0% (3.5% risk premium + 0.5% illiquidity premium). The mortgage portfolio is assumed to have the same degree of illiquidity as houses, and it bears some positive exposure to house prices, but less than the exposure associated with outright (deferred) possession of the houses. So, the expected return of the mortgage portfolio over the liquid risk-free rate ought to be between 0.5% and 4.0%. The riskiness of the portfolio will determine where in this range the portfolio expected return sits. As the LTV of the mortgages tend to zero (infinity), their expected return in excess of the liquid risk-free rate will tend to 0.5% (4.0%), as they will increasingly resemble risk-free illiquid loans (deferred possession of the house). The current LTV profile of the mortgages is producing exposure to the house risk premium of the size of around 1.35%, plus the 0.5% illiquidity premium.

1.35% is roughly two fifths of the assumed house risk premium of 3.5%. Is this an intuitive level of participation in the house risk premium? Recall from the capital analysis that we found that a 30% house price fall would result in a 16% fall in the mortgage portfolio value. This may suggest that the portfolio has a bit more than half of the exposure to house price risk as directly owning the house. But the mortgage's exposure to property risk is highly non-linear, and we would expect the measure produced by a 30% stress to overstate the 'typical' exposure the mortgage portfolio has to house price risk across all possible scenarios. A 40% 'average' participation in the house risk premium for the mortgage portfolio is therefore intuitively reasonable.

The expected cashflows of the mortgage portfolio have a Macaulay duration of 12 years. The assumed 12-year liquid risk-free rate is 0.84% and we have assumed a 0.5% illiquidity premium for all terms. The total return de-composition for the mortgage portfolio based on these calculations is shown below in Exhibit 5.2.

Exhibit 5.2: De-composition of the Mortgage Portfolio Total Expected Return

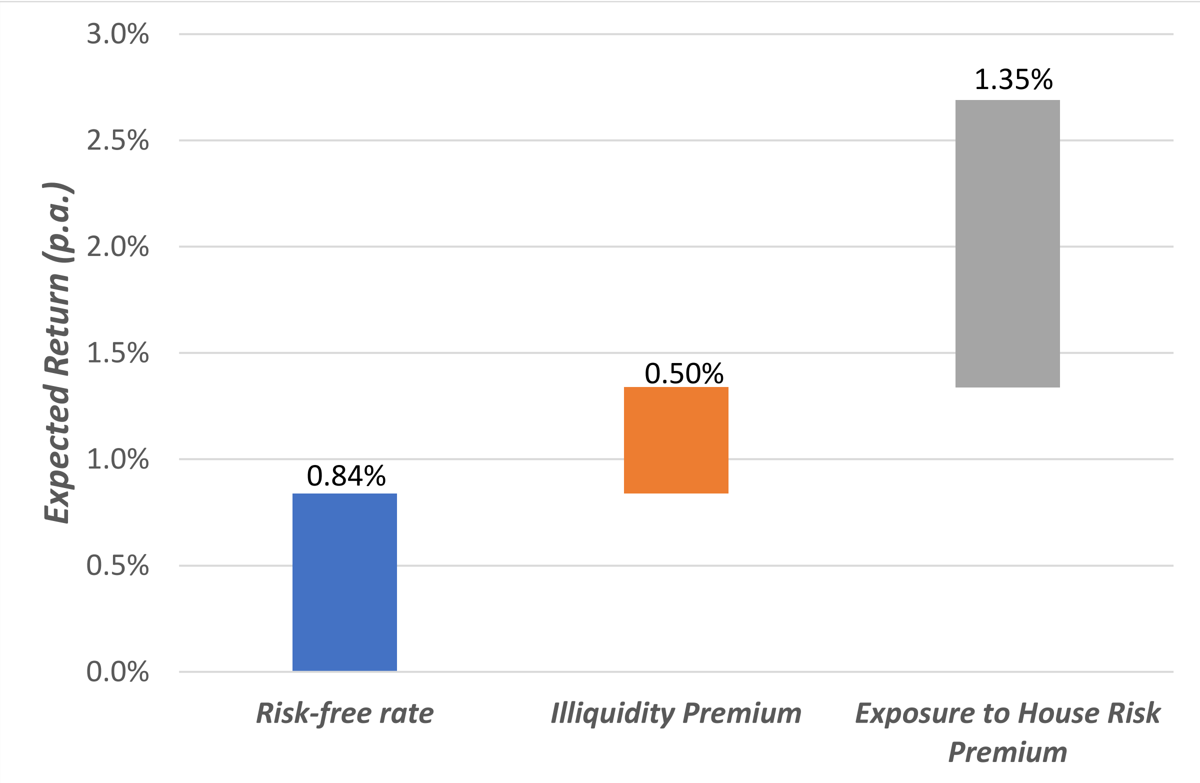

Having familiarised ourselves with the mortgage portfolio's cashflow probability distributions and what they imply about the expected return on the portfolio, we now turn our attention to the securitisation of the portfolio and the valuation of the securitisation tranches.

5.2 An example securitisation

We consider a simple securitisation structure that consists of only two elements: a debt security with a schedule of fixed promised cashflows due over the years in which the mortgage portfolio generates cashflows; and an equity tranche that owns the residual value of the portfolio after the debtholder is paid. We need some securitisation rules about how mortgage cashflows are distributed to the debt tranche and equity tranche over time in all possible circumstances. For our purposes of illustration, we make the following assumptions:

- In each year, a net cashflow is calculated as the mortgage portfolio cashflow less the promised debt tranche cashflow;
- If the net cashflow is positive, the debt tranche receives its full promised cashflow for that year;
- If the net cashflow is negative, the debt tranche receives all of the mortgage cashflow produced that year; the shortfall in the cashflow received by the debt tranche relative to the promised cashflow is added to next year's promised cashflow, plus interest as determined by the prevailing 1-year illiquid risk-free rate.
- In each year, any mortgage cashflow left after the payment to the debt tranche is paid to the equity tranche.

This simple example does not make allowances for the costs and expenses of setting up and administering the securitisation, but they could be incorporated into the cashflow model in a straightforward way.

In our example, we set the debt tranche's promised cashflows at 85% of the expected mortgage portfolio cashflows as produced in the mortgage portfolio cashflow projection of Exhibit 5.1. Exhibit 5.3 below updates Exhibit 5.1 to show the debt tranche's schedule of promised cashflows.

Exhibit 5.3: Mortgage Cashflow Distributions and the Debt Tranche Promised Cashflow Schedule

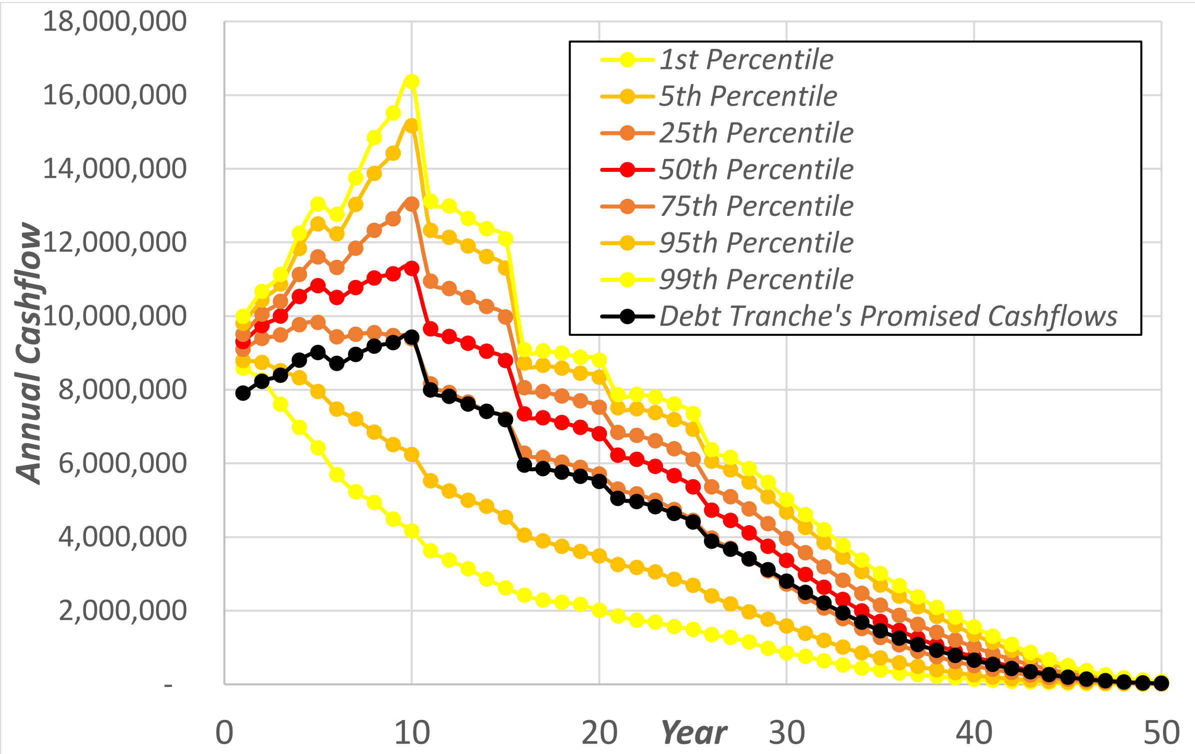

The above description of the securitisation unambiguously defines how the mortgage cashflows are distributed between the equity and debt tranches of the securitisation in every year of all possible scenarios. We can therefore re-visit the mortgage cashflow distributions shown above and calculate the debt and equity tranche cashflow distributions that they generate, given the debt tranche's promised cashflow schedule.

These cashflow distributions are shown below in Exhibit 5.4 and 5.5.

Exhibit 5.4: Projected Debt Tranche Cashflow Distributions

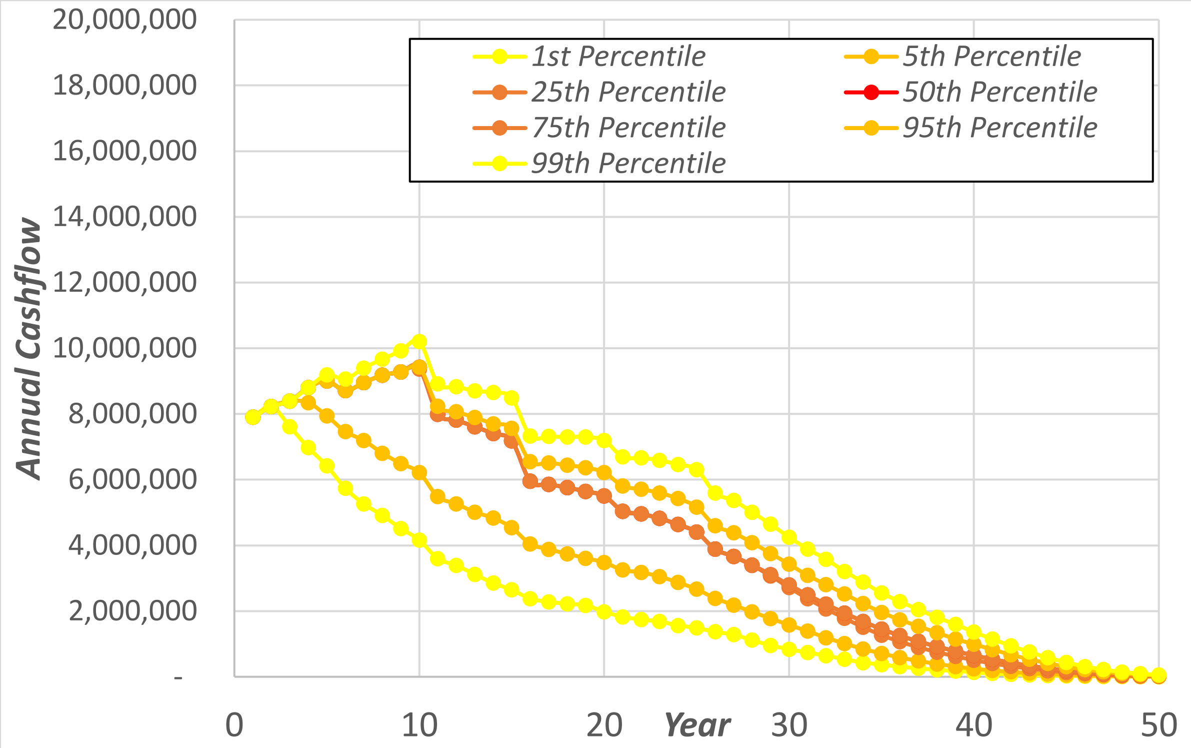

Exhibit 5.5: Projected Equity Tranche Cashflow Probability Distributions

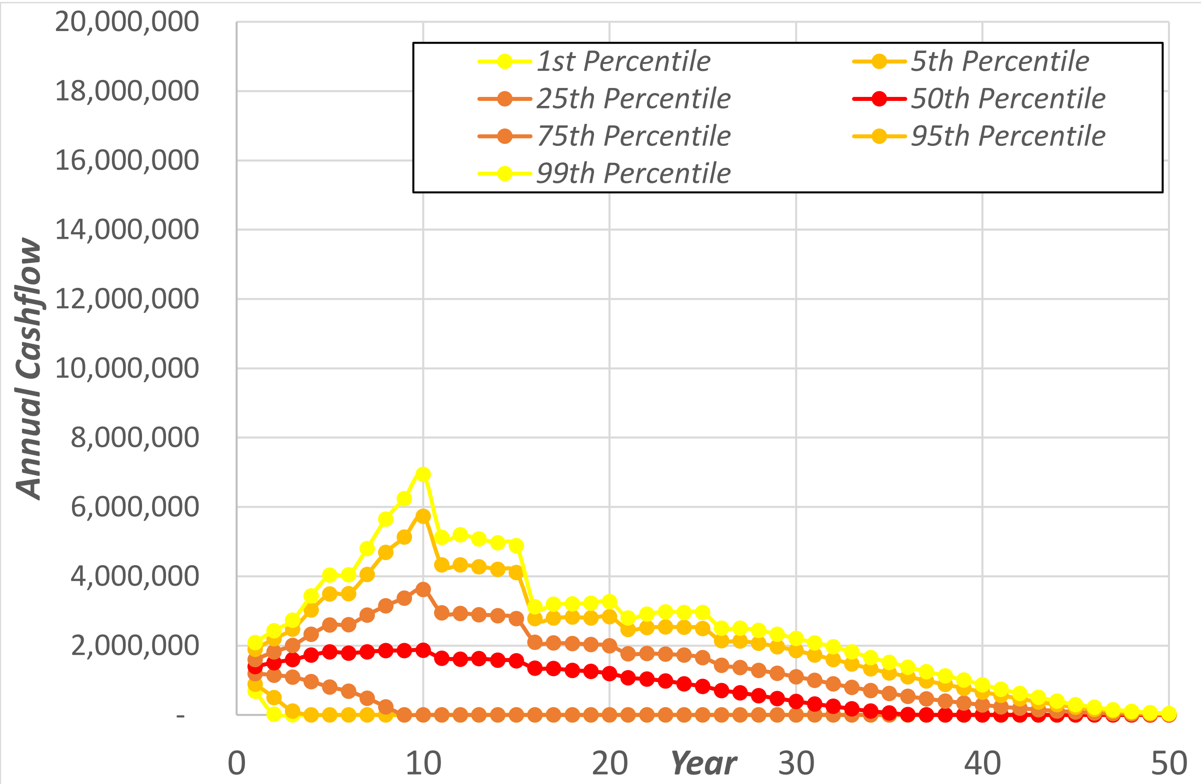

Exhibit 5.4 shows that most of our simulations result in the debt tranche receiving its promised cashflow schedule. In some years, more than the original promised cashflow is obtained as a result of the 'rollover' of previous years' shortfalls. But the debt tranche cashflow distributions have a fat left-hand that is mainly associated with NNEG losses in the underlying mortgage portfolio.

Exhibit 5.5 highlights that the equity tranche cashflows are naturally more uncertain, and the 25th percentile cashflow from year 10 onwards is zero. But as we shall see below, this uncertainty is rewarded with a substantially higher expected return.

So, the questions we now intend to answer are: If the mortgage portfolio described in the earlier implementation case studies is securitised, and the debt tranche has a cashflow schedule such as that above, and a set of cashflow distribution rules such as those above, how would we value the tranches of the securitisation? And how would we assess each tranches' capital requirements?

5.3 Valuing the example securitisation's tranches

There are no modelling choices left to make. The valuation calibration of Chapter 3.3 together with the house-to-house correlation of +0.55 and the securitisation cashflow rules fully define the risk-neutral probability distributions of the securitisation cashflows and how they are discounted. There is nothing more to the securitisation valuations than running the valuation simulation model, generating the mortgage cashflows, determining how those cashflows get distributed between the securitisation's equity and debt tranches according to the securitisation rules, and discounting all the cashflows in the usual way to obtain the tranche valuations.

One implicit assumption is worth making explicit. We are assuming all tranches of the securitisation have the same degree of liquidity as the underlying mortgages (which in turn have been assumed to have the same liquidity as the underlying house). With this assumption, the equation of value is always met: the sum of the value of the tranches of the securitisation will always equal the sum of the value of the underlying mortgages.

The table below summarises the valuation results that are obtained from the simulation model implementation.

Exhibit 5.6: Valuation of the Tranches of the Securitisation

Asset	Portfolio Value
Underlying Mortgage Portfolio	£173.2m
Securitisation's Debt Tranche	£157.2m
Securitisation's Equity Tranche	£16.0m

The debt tranche, which has a promised cashflow schedule equal to 85% of the expected mortgage cashflows, has a market value of 91% of the total mortgage portfolio. Why 91 and not 85? In a nutshell, because the debtholder gets paid first. The debt cashflows are less risky and hence are worth more.

Given the debt tranche's promised cashflows, its expected cashflows and its valuation, we can obtain a gross yield to maturity on the debt tranche and de-compose it into various components. Taking the liquid risk-free curve that is applicable to the projected cashflow in each year, we can find the spread that obtains the above debt tranche value of £157.2m when applied to its promised cashflow schedule. This spread over the liquid risk-free rate is 1.30%. We know, by assumption, that

0.50% of that spread is illiquidity premium. The remaining is a 'credit' (house price risk) spread of 0.80%.

The credit spread includes both a risk premium and the NNEG losses that the debt tranche is expected to bear (naturally, some of the NNEG losses will be borne by the equity tranche, but, as Exhibit 5.4 shows, the debt tranche has significant tail risk as a result of the mortgage portfolio's NNEG risk). We can obtain an estimate of the expected NNEG loss by comparing the promised cashflow schedule with the expected cashflows produced by the distributions in Exhibit 5.4 above. The spread over the illiquid risk-free rate that equates the expected debt tranche cashflows with the valuation is 0.56%. This implies that the participation in the house risk premium is 0.56% and so the expected NNEG loss must be 0.24%. The Macaulay duration of the debt cashflows is slightly longer than the mortgage cashflows (13 years versus 12 years). The 13-year liquid risk-free rate is 0.86%. These numbers are charted in Exhibit 5.7.

Exhibit 5.7: De-composition of the Debt Tranche's Gross Yield to Maturity

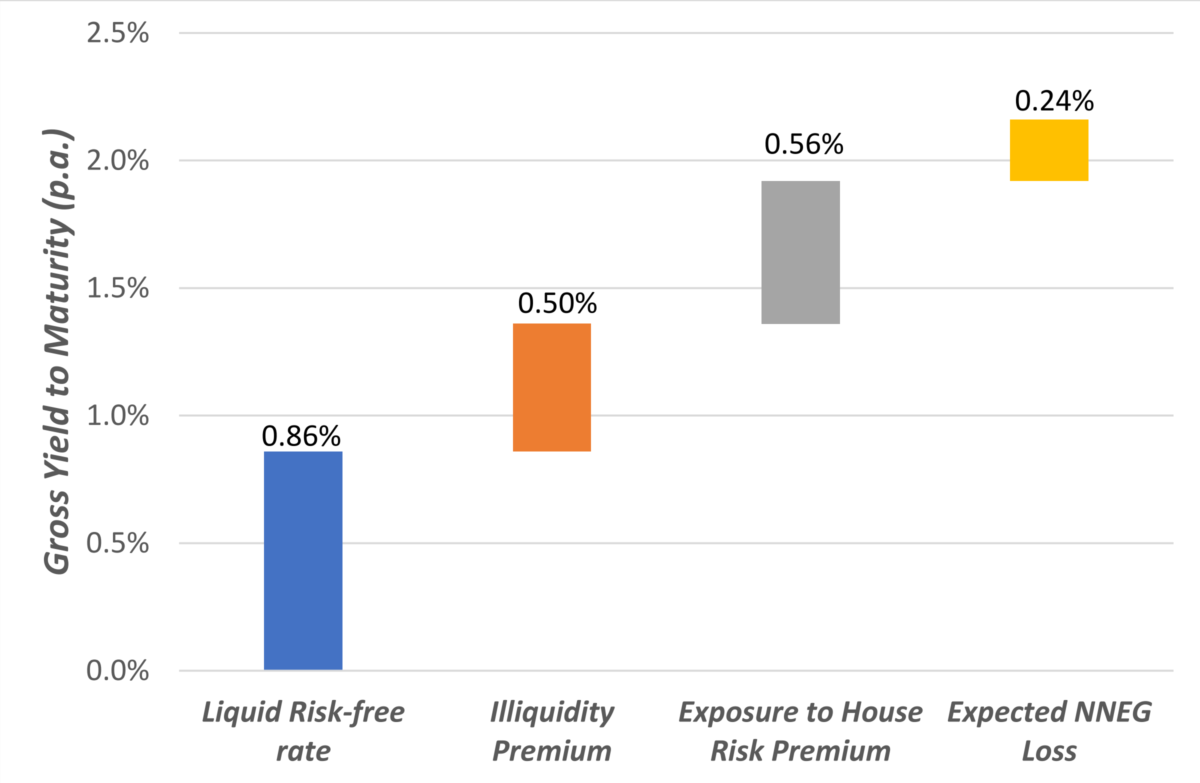

Finally, we can compare the expected returns of the debt and equity tranche (i.e. the return implied by discounting their expected cashflows to obtain a value equal to the valuations above). These are shown in Exhibit 5.8 below.

Exhibit 5.8: Expected returns in excess of the liquid risk-free rate

Asset	Expected Return in excess of the liquid risk-free rate
Illiquid Risk-free Asset	0.50%
House	4.00%
Mortgage Portfolio	1.85%
Securitisation's Debt Tranche	1.06%
Securitisation's Equity Tranche	7.83%

The expected returns on the illiquid risk-free asset and houses are inputs into this modelling process. The mortgage portfolio expected return of 1.85% results from these two expected return assumptions and the amount of house price risk in the mortgage portfolio (as was discussed in Chapter 5.1). The securitisation, in total, must generate the same expected return as the mortgage portfolio[23]. The debt (equity) tranche naturally produces a lower (higher) expected return than the mortgage portfolio because it is less (more) risky than the mortgage portfolio.

In summary, to extend the mortgage valuation methodology to value the tranches of the securitisation, we required one further form of assumption beyond the mortgage valuation model of Chapter 3.3: the correlations between house prices. Once the correlation relationship is specified, there is no remaining ambiguity about the securitisation valuations: the securitisation values are determined by the mortgage valuation model. As a result, the securitisation valuation method is quite simple, and properties such as the equation of value arise 'automatically' and hold in all circumstances.

[23] If we take a market value-weighted average of the expected returns of the debt tranche and equity tranche under these measures of expected return, we do not quite obtain the mortgage portfolio expected return. That is because these expected return statistics are not measures of the instantaneous expected return, but the expected return over the lifetimes of the securities (and the durations of the tranches are different).

6 Assessing the Capital Requirements of ERM Securitisations

This chapter considers the assessment of the capital requirements of the securitisation tranches that were valued in Chapter 5. We will use Chapter 3's stress-and-correlate approach and its set of stresses to assess the 99.5% 1-year VaR capital requirements of the securitisation tranches. The securitisation tranches are subject to a further source of risk that is not relevant to the mortgage portfolio value: a change in the house-to-house correlation. A change in this correlation will not impact on the total mortgage portfolio value, and therefore will not impact on the total value of the securitisation tranches. But it will impact on the value of each tranche.

In each of the tables below, we show the stressed mortgage portfolio values that have already been reported in Chapter 3, as well as the stressed values of the debt and equity tranches of the securitisation. In all cases, the sum of the changes to the tranche values is equal to the change in the mortgage portfolio. This is a demonstration of the equation of value holding in all circumstances – *it is not an assumption, but an outcome of the economic coherence of the mortgage and securitisation valuation methods*.

6.1 A Fall in House Prices

Exhibit 6.1: Stressed Valuations and Capital Requirements: 30% House Price Fall

Asset	Base valuation	Valuation under stress	Change in Value	Capital Requirement
Mortgage Portfolio	£173.2m	£144.7m	-£28.5m	16%
Debt Tranche	£157.2m	£141.5m	-£15.7m	10%
Equity Tranche	£16.0m	£3.2m	-£12.8m	80%

As we might expect, the equity tranche of the securitisation bears a disproportionate part of the mortgage portfolio's house price capital requirement, whilst the debt tranche generates a materially lower capital requirement than that produced by directly holding the underlying mortgage portfolio.

The equity tranche's capital requirement is eight times greater than the debt tranche. Can we sanity check the magnitude of this ratio? Yes, it is consistent with the relative sizes of the expected excess returns of the two tranches (see Exhibit 5.8).

6.2 An Increase in the Liquid Risk-free Rate

Exhibit 6.2: Stressed Valuations and Capital Requirements: Liquid Risk-Free rate Up 2.0%

Asset	Base valuation	Valuation under stress	Change in Value	Capital Requirement
Mortgage Portfolio	£173.2m	£151.1m	-£22.1m	13%
Debt Tranche	£157.2m	£130.0m	-£27.3m	17%
Equity Tranche	£16.0m	£22.1m	+£5.2m	-33%

Exhibit 6.2 delivers an interesting result – the equity tranche value *increases* when risk-free rates rise. The equity tranche essentially has a negative duration. This may seem quite a counter-intuitive result. How could this long-term asset have a negative effective duration?

The equity tranche is a call option on the mortgage portfolio value. As any actuarial student will tell us, call options have a positive *rho* - that is, their value increases with the level of the risk-free rate (all else equal). When we change the risk-free interest rate, this impacts not only on the rate at which the equity tranche's cashflows are discounted, but also on the rate at which the underlying house asset is projected. An increase in the projected rate of return of the house increases the equity tranche's expected cashflows. The gearing of the equity tranche results in this effect having a greater valuation impact than the increase in the discount rate. The equity tranche value therefore increases with an increase in the interest rate.

The other side of this coin is that the debt tranche must therefore have greater interest rate sensitivity than the underlying mortgage portfolio. This is illustrated by the results of Exhibit 6.2. It shows the debt tranche generates an interest rate capital requirement of 17%, which is some 4 percentage points more than the mortgage portfolio.

6.3 An Increase in the Illiquidity Premium

As in the case of the underlying mortgage portfolio, the 1.0% increase in the illiquidity premium term structure has an identical impact on each of the tranche valuations as a 1.0% increase in the liquid risk-free rate.

Exhibit 6.3: Stressed Valuations and Capital Requirements: Illiquidity Premium Up 1.0%

Asset	Base valuation	Valuation under stress	Change in Value	Capital Requirement
Mortgage Portfolio	£173.2m	£162.2m	-£11.0m	6%
Debt Tranche	£157.2m	£143.2m	-£14.0m	9%
Equity Tranche	£16.0m	£19.0m	+£3.0m	-19%

6.4 An Increase in the Deferment Rate

Exhibit 6.4: Stressed Valuations and Capital Requirements: Deferment Rate Up 1.5%

Asset	Base valuation	Valuation under stress	Change in Value	Capital Requirement
Mortgage Portfolio	£173.2m	£155.5m	-£17.7m	10%
Debt Tranche	£157.2m	£146.6m	-£10.6m	7%
Equity Tranche	£16.0m	£8.9m	-£7.0m	44%

The deferment rate stress generates a similar apportionment of the mortgage portfolio's deferment rate capital requirement as was produced for the house price stress capital in Chapter 6.1.

6.5 A Fall in Base Prepayment Rates

Exhibit 6.5: Stressed Valuations and Capital Requirements: Base Prepayment Rates Down 50%

Asset	Base valuation	Valuation under stress	Change in Value	Capital Requirement
Mortgage Portfolio	£173.2m	£171.9m	-£1.3m	1%
Debt Tranche	£157.2m	£158.0m	+£0.8m	-1%
Equity Tranche	£16.0m	£13.8m	-£2.1m	13%

The prepayment rate stress produces interesting results. The debt tranche value is increased by a fall in base prepayment rates, even though such a fall reduces the mortgage portfolio value. Meanwhile, the equity tranche is favourably affected by an *increase* in prepayment rates, as prepayment results in early surplus net cashflows that are distributed to the equity tranche. This, of course, leaves less mortgages behind to generate the cashflows to service the debt tranche's promised cashflows (and once the equity dividends have been paid, the debtholder cannot get them back).

This effect could be at least partially mitigated by a form of debt 'covenant' that requires any surplus net cashflows generated from unexpected levels of prepayment to be held in a cash reserve within the securitisation vehicle instead of being made available for immediate distribution to the equity tranche.

To assess the impact of uncertainty in base prepayment rates on the debt tranche capital requirement, we must consider an *increase* in base prepayment rates. Exhibit 6.6 shows the result of a 50% increase in prepayment rates.

Exhibit 6.6: Stressed Valuations and Capital Requirements: Base Prepayment Rates Up 50%

Asset	Base valuation	Valuation under stress	Change in Value	Capital Requirement
Mortgage Portfolio	£173.2m	£174.3m	+£1.2m	-1%
Debt Tranche	£157.2m	£155.7m	-£1.6m	1%
Equity Tranche	£16.0m	£18.7m	+£2.7m	-17%

6.6 A Fall in Mortality and LTC Entry Rates

Exhibit 6.7: Stressed Valuations and Capital Requirements: Mortality / LTC Rates Down 10%

Asset	Base valuation	Valuation under stress	Change in Value	Capital Requirement
Mortgage Portfolio	£173.2m	£171.7m	-£1.5m	1%
Debt Tranche	£157.2m	£157.7m	+£0.5m	0%
Equity Tranche	£16.0m	£14.0m	-£1.9m	12%

Here again we find the debt and equity tranches have opposite exposures to falls in mortality / LTC rates. The scale of the debt tranche's exposure is, however, very small.

6.7 An Increase in ERM-Implied House Price Volatility

Exhibit 6.8: Stressed Valuations and Capital Requirements: Volatility Up 5%

Asset	Base valuation	Valuation under stress	Change in Value	Capital Requirement
Mortgage Portfolio	£173.2m	£162.3m	-£10.9m	6%
Debt Tranche	£157.2m	£147.7m	-£9.5m	6%
Equity Tranche	£16.0m	£14.6m	-£1.4m	9%

An increase in ERM-implied house price volatility reduces both the debt and equity tranche values. We would always expect an increase in house price volatility to reduce the value of the debt tranche. It is not necessarily obvious that the equity tranche value will always be reduced by an increase in house price volatility, however. The 'long call option' nature of the equity tranche may result in it being positively impacted by increases in house price volatility. But the negative impact of the volatility increase on the underlying mortgage portfolio value clearly dominates in this case.

6.8 An Increase in House Price Correlation

The stresses in Chapters 6.1 to 6.7 were all applied to the mortgage portfolio in Chapter 3. Here we consider the sensitivity of the securitisation tranches to a change in the house price correlation. This change will have no impact on the mortgage portfolio value, and hence is irrelevant to the mortgage portfolio's capital assessment. But it does impact on the tranche valuations. We assume a 1-year 99.5th percentile house price correlation stress is an increase from +0.55 to +0.75.

Exhibit 6.9: Stressed Valuations and Capital Requirements: Correlation Up From +0.55 to +0.75

Asset	Base valuation	Valuation under stress	Change in Value	Capital Requirement
Mortgage Portfolio	£173.2m	£173.2m	£0.0m	0%
Debt Tranche	£157.2m	£154.3m	-£2.9m	2%
Equity Tranche	£16.0m	£18.9m	+£2.9m	-18%

As the mortgage portfolio value is unaffected by the change in correlation, we know the correlation stress must have an equal and opposite effect on the values of the equity and debt tranches. The increase in correlation increases risk without reducing the value of the underlying asset (the mortgage portfolio). As a long call option on the mortgage portfolio, this has a positive impact on the value of the equity tranche. It makes the debt tranche riskier without providing any additional upside, and hence reduces the value of the debt tranche.

6.9 Risk aggregation

We now consider the aggregation of the individual capital stresses into aggregate capital requirements for each of the equity and debt tranches of the securitisations. We use exactly the same aggregation approach as that developed in Chapter 4.8 for the mortgage portfolio. One further aggregation assumption is required beyond those used Chapter 4.8: how to aggregate the house price correlation capital requirement. Here, we make the simple assumption that this risk is perfectly correlated with house price volatility risk.

The aggregated results for the debt tranche are set out in the tables below for the three aggregation assumptions developed in Chapter 4.8. The equity tranche capital requirements are 100% for all three aggregation approaches, with and without interest rate risk.

Exhibit 6.10: The Debt Tranche's Aggregated Capital Requirements

Capital aggregation method	With Interest Rate Risk	Without Interest Rate Risk
Case 1 (stable rental income)	31%	26%
Case 2 (stable rental yield)	26%	20%
Case 3 (deferment rate with fixed relationship with interest rate)	31%	20%

You may recall from Chapter 4.8 that the 'without interest rate' results for the underlying mortgage varied between 24% and 34% across the three sets of aggregation assumptions. The debt tranche's capital requirement varies between 20% and 26%.

So, investing in the debt tranche of the securitisation instead of the underlying mortgage portfolio results in a material reduction in capital requirements. This reflects the reduction in risk that has been achieved by having the senior claim on the mortgage cashflows. Of course, this reduction in risk and capital also comes with a reduction in expected return. Recall Exhibit 5.8 showed that the annualised expected return in excess of the liquid risk-free rate of the mortgage portfolio was 1.85%, whereas for the securitised debt tranche it was 1.06%.

The market value-weighted sum of the equity and debt tranche capital requirements is slightly greater than the mortgage portfolio capital requirements in all three aggregation cases. This is because the equity and debt tranches are not perfectly correlated, as demonstrated by a number of the above stress results. As a result, a small diversification benefit would be obtained when aggregating the capital requirements of holding both the equity and debt tranches.

7 Equity Release Mortgages and the Matching Adjustment

The primary rationale for the industry practice of securitising ERM portfolios is to construct Matching Adjustment (MA)-eligible assets. The impact that the MA has on the assessment of the capital requirements of MA-eligible assets and the values of the MA-eligible liabilities that the assets are held to back are of first order importance to the financial management of these books of business. This chapter considers how the ERM valuation and capital methodologies developed in the preceding chapters can be applied when assessing the MA implications of holding MA-eligible ERM-based assets.

Suppose the securitised debt tranche considered in Chapters 5 and 6 was deemed to be MA-eligible. How would we assess the MA impact that it generates? There are two elements to this answer: how much can the liability discount rate be increased by as a result of backing the liability with this asset instead of, say, a liquid risk-free bond (the resulting reduction in liability value is referred to as the 'MA benefit'); and how is the assessment of the securitised debt capital requirement of Chapter 6 impacted by the MA-eligibility of the asset.

The standard approach to answering these questions would involve developing an internal credit rating process for the securitised debt tranche. In a nutshell, this internal credit rating could then be mapped to the MA's Credit Quality Step (CQS) scale. The CQS, along with the duration of the asset, is then used to determine the Fundamental Spread associated with the asset according to EIOPA's latest published tables. The Fundamental Spread is then deducted from the debt asset's yield to determine the net yield that can be used for the purposes of discounting the liability cashflows.

To take a numerical example, we saw in Chapter 5 that the securitised debt tranche valuation and its promised cashflow schedule imply a gross yield to maturity of 2.16%, which in turn implies a spread of 1.30% over the duration-equivalent liquid risk-free bond yield of 0.86% (see Exhibit 5.7). The duration of the debt tranche's promised cashflows is 13 years. This paper has not offered an internal credit rating process for the securitised debt tranche, but we might anticipate that the CQS mapping for this asset would result in a CQS of 1 or 2. EIOPA's end-May 2020 tables stipulate that the 13-year non-financial fundamental spread for a CQS of 1 is 31bps, and for a CQS of 2 it is 42bps. So, depending on the conclusion of the CQS mapping process, the liability discount rate would be increased by 0.88% (CQS of 2) or 0.99% (CQS of 1) as a result of switching an MA-eligible liability's backing assets from a liquid risk-free bond to the securitised debt tranche.

To assess the capital requirement in the presence of the MA, we need to calculate stressed values for the liability as well as the asset. The stressed liability values will be a function of how the CQS changes under each stress, and the impact this has on the applicable Fundamental Spread that is applied to the debt spread that arises in the stress.

7.1 The Effective Value Test

Nothing in the above discussion is specific to ERM-based assets – these forms of liability valuation and capital requirements calculations will apply to any MA-eligible unrated debt security. There are, however, also some ERM-specific MA considerations. In particular, the PRA's Effective Value Test (EVT) and its associated text in PRA SS3/17.

The EVT is a partially-prescribed calculation that places a quantitative ceiling on the amount of MA benefit that is derived from MA-eligible ERM assets. The associated text also sets out some more fundamental principles. Most notably, in a discussion of the treatment of ERM-based securitised notes and the impact that NNEG risk has on their valuation and the MA benefit that is consequently derived, Section 3.4 states: "Compensation for these NNEG risks should not lead to an MA benefit".

The intention and implications of this statement are no doubt open to interpretation. But if interpreted in the most straightforward way, it arguably simplifies the MA treatment of ERMs considerably. And such an interpretation would be very simple to implement within the valuation and capital framework developed in this paper. To see how, it is useful to again refer to Exhibit 5.7, which provides the de-composition of the gross redemption yield of the securitised debt tranche. In our example, the securitised debt tranche's gross redemption yield of 2.16% was de-composed into a 0.86% liquid risk-free rate, a 0.50% illiquidity premium, 0.24% expected NNEG loss and a 0.56% participation in the house risk premium. The latter two elements of this de-composition are only present because of NNEG risk – they represent the portion of the cost of the NNEG that is borne by the debt tranche. If we ran the valuation process for a securitised debt tranche that had zero exposure to NNEG risk, the expected NNEG loss would be zero and the participation in the house risk premium would be zero. If the MA benefit should not include any compensation for NNEG risks, this would imply the MA liability discount rate should not exceed the liquid risk-free rate by more than 0.50% (i.e. the assumed illiquidity premium). That is, the MA liability discount rate should not exceed the illiquid risk-free rate.

So, this direct interpretation of the 'no NNEG compensation' principle would imply reducing the net MA spread calculated above of 0.88% or 0.99% to 0.50% in our example[24]. We will refer to this as the 'no NNEG compensation' approach. If we were only guided by the statement of SS3/17's Section 3.4, we might therefore conclude there is a straightforward relationship between the mortgage valuation methodology of this paper and the maximum permitted MA benefit. But there is more to SS3/17 than this text. There is, of course, also the EVT calculation.

The EVT starts from the premise that directly-held equity release mortgages are not eligible MA assets, and that all ERM-based assets in MA portfolios will therefore be debt tranches of ERM securitisations. The EVT compares the valuation of all tranches of the re-structured ERM portfolio with a valuation of the ERM portfolio before it is re-structured. The underlying mortgage portfolio is valued using an economic basis that specifically makes no explicit allowance for illiquidity. The difference between this 'economic value'[25] and the market value of the total ERM securitisation (including MA and non-MA assets) is called 'Day 1 Gain' in the EVT. To pass the Test, the MA benefit derived from the ERM assets must not exceed the EVT Day 1 Gain measure.

Given the purpose of the test is ostensibly to measure the difference between the market value of the ERM assets and their value in the absence of any allowance for illiquidity, there is no obvious rationale for using different models, calibrations or assumptions in the valuation of both 'sides' of the test (other than the use of only the liquid risk-free rate in the 'left hand side' of the test). In that

[24] The mechanics of implementing this change to the MA spread within the prescribed MA yield calculation process is not completely straightforward. The MA rules stipulate that the EIOPA Fundamental Spread that corresponds to the CQS must be used in the MA calculation. So, the only way to adjust the MA spread is to adjust the asset CQS. Moreover, the EIOPA tables do not specify a minimum level for the FS, but the specific value for the FS that must be used. In the end-May tables, a CQS of 3 and duration of 13 years produces a FS of 57bps and hence would produce a net spread of 0.73% in our example. And a CQS of 4 had a FS of 173bps, which exceeds the gross yield on the asset. Perhaps approval could be sought for the prudent action of assuming a FS greater than the prescribed EIOPA value for the assessed CQS of the asset.

[25] The EVT economic valuation process de-composes the economic value of the mortgage portfolio into elements such as pre-NNEG value, expenses and NNEG. But this de-composition of the mortgage value does not affect the outcome of the Test, and so we do not consider it here.

case, in the language of our methodology, the EVT's Day 1 Gain is simply the difference between the mortgage portfolio valued without an illiquidity premium and the actual value that has been placed on the total securitisation (which may or may not have assumed a positive illiquidity premium).

These numbers are easily calculated under the methodology of this paper. Taking the example developed throughout this paper, the total value of the securitisation (the 'right-hand side' of the Test) is £173.2m (of which, £157.2m is the debt tranche value, and £16.0m is the equity tranche value). If we re-run the cashflow simulation valuation model with zero illiquidity premium, the value of the mortgage portfolio is found to be £178.4m. And so, in this example, using the same volatility and deferment rate assumptions on both sides of the Test, the EVT Day 1 Gain would be £178.4m - £173.2m = **£5.2m**.

Now suppose we have an MA-eligible set of fixed liability cashflows with exactly the same run-off pattern as the securitised debt tranche, and those liability cashflows are currently backed by (liquid) gilts. Further suppose the securitised debt tranche is deemed MA-eligible and is assessed to have a CQS of 2. So, as discussed above, before we apply the EVT, we would find that switching from gilts to the securitised debt tranche would permit an increase of 0.88% in the liability discount rate from the 13-year liquid risk-free rate of 0.86% to 1.74% (= 2.16% (Debt Gross redemption yield) - 0.42% (CQS 2 Fundamental Spread).

Under the gilt basis, these liability cashflows would have a present value of **£187.4m**. The MA discount rate of 1.74% results in a liability present value of £166.5m, and so the MA benefit generated by the debt tranche is **£20.9m** (= £187.4m - £166.5m). In this case, our Day 1 Gain EVT result of £5.2m calculated above would place a ceiling on the permitted MA benefit, and would consequently compel the firm to use a FS that produces an MA benefit of no more than £5.2m.

How do these results compare with the 'no NNEG compensation' interpretation of SS3/17 Section 3.4? Above, we argued the 'No NNEG compensation' principle of SS3/17 Section 3.4 would imply the liability discount rate should be increased by no more than the illiquidity premium obtained by switching from gilts to the debt tranche. In our example, the illiquidity premium has been assumed to be 0.50%. If we increase the gilt liability discount rate by 0.50%, we obtain a liability present value of **£175.0m**, which implies an MA benefit of **£12.4m** (= £187.4m - £175.0m). This is also materially less than the MA benefit produced by a CQS of 2 for the securitised debt tranche. But it is more than twice as big as the EVT Day 1 Gain of £5.2m. This seems odd given that the Day 1 Gain number has been driven entirely by the impact of the illiquidity premium on the mortgage portfolio valuation. What is going on here? Why is the assessed EVT Day 1 Gain only around 2.8% of the liability value, when we have assumed a 0.50% illiquidity premium and the liability duration is 13 and a bit years?

The relatively small size of the EVT Day 1 Gain calculated above arises because the 'left-hand side' of the EVT applies to the ERM portfolio value *as a whole*, rather than only to the MA-eligible element of the portfolio (the securitised debt tranche in our example). We noted above that the mortgage portfolio value increases from £173.2m to £178.4m when the illiquidity premium (ILP) is reduced from 0.50% to zero. We know our securitisation valuation method of Chapter 5 will always produce tranche valuations that are consistent with the equation of value, so we know the two tranche values will also add up to £178.4m when the ILP is assumed to be zero. But the change in mortgage portfolio value that arises when the ILP is set to zero is not proportionally distributed between the values of the securitised debt tranche and securitised equity tranche. On the contrary, reducing the assumed ILP of 0.50% to zero results in the securitised debt tranche increasing by £7.0m to £164.2m, and the equity tranche *falling* in value by £1.8m to £14.2m.

This fall in the equity tranche value may seem a highly counter-intuitive result. If we have assumed an asset is illiquid, and that investors require some compensation for bearing this illiquidity (i.e. an asset value discount), and then we suppose that this required illiquidity compensation no longer applies, then surely this must result in the value of the illiquid asset increasing? The answer is, all other things being equal, yes it would. But when the illiquidity premium is reduced to zero, two relevant things happen. The discount rate that applies to the securitised tranche cashflows is reduced (positive impact on valuation); and the cashflows of the securitised tranches are also altered. Why are the cashflows altered? Because the tranche cashflows depend on the behaviour of (illiquid) house prices. If we assume there is no illiquidity premium, this means the expected[26] rate of increase of the price of each underlying (illiquid) house is also reduced by 0.50%. This has a negative impact on mortgage valuations, and both tranche values. But for the mortgage portfolio and the debt tranche, the impact of the reduction in discount rate outweighs this impact. However, the highly geared nature of the equity tranche makes this cashflow effect bigger than the impact of the reduction in discount rate, resulting in the value increasing. It is the same effect which results in the equity tranche having a positive 'rho', as per the discussion of Chapter 6.2.

As a result of this effect, by setting the maximum MA benefit as the illiquidity valuation impact on the total mortgage portfolio rather than only on the MA-eligible element, the EVT arguably produces a number that is lower than the logical intention of SS3/17 (if we interpret the aim of the EVT as ensuring that the MA benefit does not exceed the value of the illiquidity valuation effect on the assets backing the MA liability).

The above argument may seem quite technical and abstract. It might be argued that the spirit of the EVT is much more straightforward – it is merely intended as a relatively simple prescribed calculation that checks that the MA benefit generated by ERM assets is not unreasonable. From this perspective, we could abandon our above approach of assuming the same volatility and deferment rate assumptions on both sides of the Test, and instead use the prescribed minimum parameters in the EVT economic valuation on the left-hand side of the EVT.

If we value our example mortgage portfolio using the PRA's most recent prescribed EVT minimum deferment rate of 0.5% (instead of 3.5%) and house price volatility of 13% for all terms and LTVs (instead of the volatility model developed in Chapter 3.3), and the liquid risk-free yield curve of Chapter 3.1.2 (instead of the illiquid risk-free curve), then we obtain a mortgage portfolio value of £220.9m. This implies an EVT Day 1 Gain, and hence maximum permitted MA benefit, of **£47.7m** (= £220.9m - £173.2m)...almost ten times bigger than the EVT Day 1 Gain calculated above! So, in this case, the MA benefit of £20.9m that was calculated when the debt tranche was given a CQS rating of 2 does not trouble the EVT's maximum permitted amount of MA benefit, and indeed leaves plenty of headroom.

The above discussion has highlighted that the MA can be quite complicated, and that different interpretations of the EVT and SS3/17 can produce materially different MA results for ERM-backed assets. In our example, an implementation of the EVT with the prescribed minimum parameters results in a large EVT Day 1 Gain and, consequently, an accommodative maximum MA benefit. If, however, we were to attempt to implement the EVT in a spirit inspired by with SS3/ 17 Section 3.4, with consistent volatility and deferment rate parameters used on both sides of the Test, the

[26] This change in expectation applies in both risk-neutral and real-world settings, as the illiquidity premium impacts on the risk-neutral drift rate of illiquid assets.

maximum permitted MA benefit produced in our example is materially lower, and less than the MA benefit that would be associated with the debt tranche if it were assessed to have a CQS of 2.

7.2 The Impact of the Matching Adjustment on the Capital Requirement

The implications of the MA for the ERM capital requirement depends on how the MA liability discount rate is assumed to behave under stress. When modelling the behaviour of the MA liability discount rate under stress, it is generally understood that the MA benefit produced under stress must always pass the EVT result in that stress.

To see how this can work, let us again use our example securitised debt tranche. Let's take the simplest case and assume we implement the EVT using the prescribed minimum parameters. And suppose we implement the MA using the 'No NNEG Compensation' principle, i.e. by assuming that only the illiquidity premium of the securitised debt tranche can be added to the MA liability discount rate. As noted above, this implies a sort of reverse fitting of the FS, which may either be implemented by finding the CQS that produces a similar FS, or assuming an FS which is greater than the EIOPA prescribed value for the CQS (if such an action were permitted). Such an approach may or may not be deemed acceptable to all relevant parties, but it is arguably economically logical and prudent in the sense that it always avoids capitalising unearned NNEG compensation in the liability valuation.

This interpretation of the MA would have quite a simple implication for capital. In the capital assessments of Chapters 4 and 6, variation in the illiquidity premium was explicitly treated as a risk factor. Under the 'no NNEG compensation' interpretation, the MA liability discount rate under stress will always be the illiquid risk-free rate (i.e. the liquid risk-free rate plus the illiquidity premium). The variation in the illiquidity premium on both sides of the balance sheet therefore results in it 'cancelling out' in the capital requirement. That is, the MA capital requirement of the securitised debt tranche is the Chapter 6 capital requirement with the illiquidity premium stress removed from the aggregation. This, again, is a very simple calculation to implement.

For example, we found in Chapter 6 that the Case 1 (stable rental income) aggregation approach generated a capital requirement for the securitised debt tranche when it is backing matching liabilities of 26% (see Exhibit 6.10). If the illiquidity premium stress is removed from that capital aggregation calculation, this capital requirement is reduced from 26.2% to 24.7%. The MA treatment would therefore reduce the capital assessment associated with the securitised debt tranche from 26.2% of £187.4m = £49.1m to 24.7% of £175.0m = £43.2m.

These numbers are summarised below in Exhibit 7.1.

Exhibit 7.1: Impact of MA on Liability and Capital ('No NNEG Compensation' approach)

	Without Matching Adjustment	With Matching Adjustment	% Reduction
Liability Value	£187.4m	£175.0m	6.6%
Capital Requirement	£49.1m	£43.2m	11.9%
Total Asset Requirement	£235.6m	£218.3m	7.6%

This (simple and prudent) interpretation of the PRA's requirements for the MA treatment of ERMs results in a reduction in the total asset requirement of 7.6% in our example. It should be emphasised here that the numerical result is entirely a function of illustrative parameter assumptions for the level of the illiquidity premium and its 1-year 99.5th stressed value. The purpose of the example is to illustrate the logic that is implied by the valuation and capital methodologies developed in this paper.

This illustrative MA implementation has one final task remaining. We must check that this approach to assessing the MA benefit does not result in failing the EVT in any of the capital stresses. The capital requirement for the securitised debt tranche has been assessed using a total of 8 stresses (see Chapters 6.1 – 6.8). So, we need to calculate the EVT results in each of these 8 stresses.

Exhibit 7.2 shows the EVT Economic Values calculated in each case, using the base calibration of the minimum prescribed parameters of a house price volatility of 13% and deferment rate of 0.5%. The EVT Day 1 Gain is calculated as the difference between these EVT Economic Values and the asset portfolio values that have already been calculated in Chapters 4 and 6.

Exhibit 7.2: EVT Day 1 Gain Under Stress

	EVT Economic Value	Asset portfolio value[27]	EVT Day 1 Gain
Base	£220.9m	£173.2m	£47.7m
House Price Fall 30%	£194.7m	£144.7m	£50.0m
Liquid Interest Rate Up 2%	£182.2m	£151.1m	£31.1m
Illiquidity Premium Up 1%	£220.9m	£162.2m	£58.7m
Deferment Rate Up 1.5%	£202.1m	£155.5m	£46.6m
Prepayment Rate Up 50%	£218.5m	£174.3m	£44.2m
Mortality Rate Up 10%	£220.1m	£174.4m	£45.7m
House Price Vol Up 5%	£205.4m	£162.3m	£43.1m
House Price Correlation Up	£220.9m	£173.2m	£47.7m

Exhibit 7.3 shows the calculation of the MA Benefit in each stress under the 'No NNEG Compensation' approach described above. The liability value without matching adjustment is calculated as the present value of the fixed liability cashflows discounted using the prevailing gilt yield (liquid risk-free yield curve). So, this value only changes in the liquid risk-free rate stress. Under the 'No NNEG Compensation' approach to MA, the liability discount rate with matching adjustment is always equal to the illiquid risk-free yield curve, and hence changes when either the liquid risk-free rate or the illiquidity premium change.

Exhibit 7.3: MA Benefit Under Stress

	Liability Value Without Matching Adjustment	Liability Value With Matching Adjustment	MA Benefit
Base	£187.4m	£175.0m	£12.4m
House Price Fall 30%	£187.4m	£175.0m	£12.4m
Liquid Risk-free Rate Up 2%	£144.5m	£136.0m	£8.4m
Illiquidity Premium Up 1%	£187.4m	£163.8m	£23.5m
Deferment Rate Up 1.5%	£187.4m	£175.0m	£12.4m
Prepayment Rate Up 50%	£187.4m	£175.0m	£12.4m
Mortality Rate Up 10%	£187.4m	£175.0m	£12.4m
House Price Vol Up 5%	£187.4m	£175.0m	£12.4m
House Price Correlation Up	£187.4m	£175.0m	£12.4m

[27] These values correspond to the asset values under stress reported in Chapters 4 and 6.

The MA Benefit arising under the ‘No NNEG Compensation’ approach is fairly stable under stress. As we would expect under this approach, the MA Benefit increases with increases in the illiquidity premium. The MA Benefit falls in the liquid risk-free rate up stress because the liability duration is longer without the matching adjustment than it is with the matching adjustment. But this stress does not contribute to the capital calculations of Exhibit 7.1 as the liability cashflows are matched by the securitised debt tranche cashflows.

Exhibit 7.4 compares the EVT Day 1 Gain generated under each stress (as reported in Exhibit 7.2) with the MA benefit under each stress (as reported in Exhibit 7.3).

Exhibit 7.4: EVT Test Results Under Stress

	EVT Day 1 Gain	MA Benefit	EVT Test: Pass or Fail
Base	£47.7m	£12.4m	Pass
House Price Fall 30%	£50.0m	£12.4m	Pass
Liquid Interest Rate Up 2%	£31.1m	£8.4m	Pass
Illiquidity Premium Up 1%	£58.7m	£23.5m	Pass
Deferment Rate Up 1.5%	£46.6m	£12.4m	Pass
Prepayment Rate Up 50%	£44.2m	£12.4m	Pass
Mortality Rate Up 10%	£45.7m	£12.4m	Pass
House Price Vol Up 5%	£43.1m	£12.4m	Pass
House Price Correlation Up	£47.7m	£12.4m	Pass

Exhibit 7.4 shows that, for our example, the MA Benefit that is generated by the ‘No NNEG Compensation’ approach comfortably passes the EVT test in all capital stresses.

The MA implementation rules and processes are quite complicated, especially for ERM assets. However, the above discussion may serve to highlight that there is a natural correspondence between the objectives of the MA and EVT and the valuation and capital methodologies developed in this paper.

8 Conclusions

This paper proposes what might be considered as a ‘second-generation’ ERM valuation methodology that builds on the significant ERM valuation modelling work done by UK actuaries in recent years. The proposed methodology aims to address some of the shortcomings that can naturally arise in the development of valuation approaches for this complex asset class and its securitisations. In particular, by focusing more closely on the economic logic of the valuation process, the paper has proposed a pragmatic and rigorous ERM valuation method that readily provides a basis for the consistent valuation of ERM securitisation tranches as well as mortgage portfolios. And, as is usual for a capital requirements definition based on Value at Risk, the valuation methodology also underpins the assessment of capital requirements for both mortgage portfolios and tranches of securitisations of mortgage portfolios. Finally, the paper discussed how the economic logic of the valuation methodology may have a natural correspondence with the Matching Adjustment and ERM-specific requirements of PRA SS3/17.

One of the fundamental differences between the proposed valuation methodology and the valuation approaches that have tended to be used in recent ERM valuation practice is that here mortgage values are made consistent with observed ERM prices by finding the house price volatilities implied by those prices, rather than by using other forms of valuation adjustment that are harder to interpret. Current valuation adjustments are sometimes labelled as illiquidity premia, despite their assumed size varying from mortgage to mortgage when the mortgages have identical levels of liquidity. The key point is that a valuation methodology should deliver an economically logical way of interpolating and extrapolating from the values of mortgages with observed prices (typically origination prices in the primary retail market) to the values of mortgages with unobserved prices (in-force ERMs). The proposed methodology can help firms meet this objective more effectively, thereby avoiding the subsequent need for *ad hoc* adjustments to ensure mortgage portfolio valuations and securitisation valuations meet basic economic criteria such as the equation of value.

The valuation methodology proposed in this paper aligns more closely with conventional derivative valuation practice. There is no doubt, however, that the house price volatilities implied by ERM prices are sensitive to a number of difficult assumptions about factors such as the prepayment behaviour of ERM borrowers and the market price of deferred possession of houses. Equity release mortgage valuations are inevitably sensitive to difficult assumptions about these complex features, and there is no easy escape from that reality. The proposed methodology is not assumption-free, but nor is any other.

The proposed methodology does not prescribe specific implementation assumptions. A wide array of approaches can be used in its implementation. As demonstrated in Chapter 3, these can vary from a simple Black-Scholes approach through to a quite complex simulation model with various assumptions for fat-tailed house price dynamics and dynamic borrower prepayment behaviour. Implementations could be made much more sophisticated than the examples developed in the paper. ERM-investing insurance firms will have developed their own implementation assumptions for the modelling of various important valuation factors (borrower prepayment behaviour, house deferment rates, etc.) in their current valuation methods, and this modelling can be naturally incorporated into the valuation methodology developed in this paper.

References

Dowd, Blake, Buckner and Fry (2019), 'The Valuation of no-negative equity guarantees and equity release mortgages', *Economics Letters*, Vol. 184.

IFoA Equity Release Mortgages Working Party (2020), 'A Discussion Note on the Economic Valuation of Equity Release Mortgages as part of the PRA's Effective Value Test'.

Jeffery and Smith (2019), 'Equity Release Mortgages: Irish and UK Experience', *Society of Actuaries of Ireland*.

Prudential Regulation Authority (2019), *Supervisory Statement 3/17*.

Tunaru (2019), 'UK Equity Release Mortgages: a review of the No Negative Equity Guarantee', *Actuarial Research Centre, Institute and Faculty of Actuaries*.

Turnbull (2020), '*Notes on Derivative Valuation and Illiquid Assets*', https://craigturnbullfia.com/notes-on-derivative-valuation-and-illiquid-assets/

Turnbull (2020), '*Prepayment risk and Equity Release Mortgages'*, https://craigturnbullfia.com/prepayment-risk-and-equity-release-mortgages/

www.craigturnbullfia.com