Equity Release Mortgages: Irish & UK Experience

Tony Jeffery and Andrew D Smith

Society of Actuaries in Ireland, 28 March 2019

Contents

Abstract: Equity release mortgages (ERMs), also called lifetime mortgages, have played an increasing role in generating income for retired home-owners, especially in the UK. As new liquidity rules have reduced the supply of bank lending, so insurers have stepped in, encouraged by generous regulatory treatment for annuity writers. Some methods for valuing ERMs have proved controversial. As the volume of these assets grows on insurance balance sheets, there are concerns that insurers' reliance on continued house price growth could make the industry less resilient to the next house market downturn.

This paper describes the basic products and illustrates alternative valuation methods with reference to Ireland and the UK. We summarise recent research and provide example calculations to illustrate the competing methods, highlighting areas of actuarial debate. We conclude with implications for Ireland where so far ERM volumes have remained modest. We consider technical approaches for Irish ERMs and discuss the value of these products – both positive and negative – to our society as a whole.

1 Introduction

1.1 Why have we written this paper?

Issue 1487 of Private Eye (11th to 24th of January 2019) included an article titled "Just about Managing". It was primarily about the Just Group, a large insurance group in the UK specialising in enhanced annuities and with substantial holdings of Equity Release Mortgages ("ERM"). The article claims that negative equity on ERMs pose risks that are not being adequately recognised. In the article was contained the following quotes.

"The PRA's recognition of the problem was itself delayed by years of lobbying from firms……and by the Institute of Actuaries. While accountants have faced some political heat recently, the even more easily ignored actuaries - who measure things like likely future losses – have avoided such scrutiny"

"Back in the 2005, in the wake of the collapse of Equitable Life the government's Morris Review recommended major improvements in the actuarial profession. Don't discount the possibility of another one being needed soon."

That Private Eye (the UK equivalent of the Phoenix journal) should be alleging that the IFOA has been lobbying to delay the recognition of a problem in an area where the Actuarial Profession would claim expertise and independence, is serious indeed. If the profession had the reputation we might desire, such allegations might be unlikely. In this paper we will touch briefly on some reasons why this claim might be thought to have some substance.

In Ireland, by contrast, the small market that existed before 2012 has dried up. However that does not mean that it could not (or indeed should not) reopen. We believe that if this is imminent then it would be sensible for the SoAI to consider carefully what its position on ERMs should be. Our purpose is to set out what we think should be considered and provide our views to stimulate debate for Ireland.

But why are ERMs controversial in the UK? There have been two consultations and policy statements from the PRA. The Institute and Faculty of Actuaries (IFoA) responded to both and then (jointly with the ABI) commissioned special research, which has only recently reported. It has been discussed at the Treasury Select Committee. There was even a program on Radio 4 devoted to the issue.

The sums involved are large and the risks substantial. Several large insurance companies are active in selling ERMs and backing their annuity liabilities with them.

But why is this controversial? Simply because of the no negative equity guarantee (or "NNEG"). Companies selling ERM guarantee that no matter how much the ERM grows in amount (and they usually do), it will not exceed the value of the property it is secured upon, when that property is sold. It is thus an option on individual residential properties, a market that is not active, deep, liquid or transparent. In this paper we discuss the valuation of such options.

More important to actuaries than valuation is the question of risk and capital. What is the prudent level of capital needed to make insurance companies safe? We discuss that also.

We also consider the question of whether these products are desirable? In the UK, it seems widely held that the existence of ERMs is desirable for society as a whole. This should not go unchallenged.

In a housing market where there is a shortage of family homes, do we want to encourage empty nesters to live in places too big for them?

The obverse of this issue is whether they are appropriate for individuals to take out. There is no doubt that being able to get hands on some value locked up in the house that one is living in, may be convenient but payday loans are also convenient. Are they a good idea? This is also explored in this paper.

It is also widely claimed (including by the PRA) that ERMs are a suitable product to back annuities. In a letter to Industry of 2nd July 2018, David Rule stated "We continue to believe that restructured ERMs are an appropriate asset to back annuities as part of a diversified portfolio." That last phrase could be interpreted as not ruling out that the part of the portfolio should be quite small. We would be comfortable only with quite a small proportion indeed.

Our aim in presenting this paper is simple. The market in the UK has grown large with active support from the IFOA while many actuarial issues unresolved. By raising these issues for open discussion now, the SoAI may debate whether it wants to take a similar path. The Central Bank of Ireland might well want to consider what regime to apply to ERM.

We set out below our views on the issues that should be considered to start that debate

1.2 Acknowledgements

The authors would like to acknowledge assistance from Dean Buckner and Kevin Dowd . Their web site "Eumaeus Project" is devoted to the issue of ERMs and is well worth visiting for those with an interest in this subject. Caroline Twomey made useful suggestions concerning longevity basis and Emily O'Gara was helpful with long term care information. Alan Reed made many insightful comments about the modelling and made numerous helpful suggestions that have improved the clarity of earlier drafts of this text. Seamus Creedon picked up several important points we had missed. We are also grateful for useful discussions with Andrew Cairns, Dan Georgescu, Guy Thomas and James Thorpe. Any errors are our own responsibility. As always, the views expressed in this paper are the authors' own. **In particular they in no way represent the views of our current or previous employers.**

2 Background and history

2.1 What is an ERM?

An equity release mortgage is a loan taken out by a property owner (or joint property owners; acting together for convenience we will use singular term throughout) which is secured against the value of that property. The loan is only normally only repayable when the property owner dies. There are other circumstances when the loan may become payable. If the borrower moves into long term care, then the loan is normally also repayable. Early redemption usually carries a penalty.

It cannot be repeated too often that ERM is a misnomer. The equity in a property is not released it is borrowed against. The value of the property to the owner becomes geared.

There are options for the borrower. The interest on the loan can be paid on a regular basis out of income or it may be rolled up against the value of the property. There is also offered a form where the borrower may take a certain amount now but have the option to take out a further tranche without further loan underwriting. Even loans that are fixed in value are in practice often topped up.

As the loan will roll up, the amount that will be offered is a lot lower than the value of the property. The percentage will depend on the age of owner (the older the more), the terms of the mortgage and location of property. A rule of thumb is that 30% of the value can be borrowed (Tony was offered 31%, he is aged 62), but we have seen reports that 50% is attainable.

When ERMs were first offered they were in the form of having no limit on the amount of the outstanding loan. This meant that if the amount of the loan was greater than the value of property (net of disposal costs), the ERM provider could claim against the residual estate. This was perceived as being contrary to the interests of borrowers. As a result the No Negative Equity Guarantee ("NNEG") was introduced and this is now pretty much universal in UK and Ireland. This reduces risk to borrowers but increases risk to the lenders. Most of the controversy about ERM's is related to the NNEG.

2.2 The ERM Market in the UK

The ERM market in the UK is dominated by members of the Equity Release Council ("ERC"). All statistics in this section are drawn from their web-site. They claim they represent "over 90% of the sector". To belong requires ERM providers to adhere to certain standards in market conduct, primarily in the sales process and in providing a NNEG.

In a press release of 24th January 2019, the ERC announced that 46,397 new ERMs had been taken out in 2018, with 32,759 existing borrowers taking drawdown and 3,644 taking loan extensions, this meant there were nearly 83,000 customers taking loans and they borrowed £3.94Bn.

That is a lot of money.

Roughly two-thirds of borrowers took products allowing them to make further drawdown later and their average initial loan was a little over £60,000. The other third were borrowing more their average was just over £95,000.

The Jan 2019 release does not give the size of the loan book outstanding nor does the previous release for Q3 2018, however both releases make it clear that while the market is now growing very fast, there have been borrowings of over £1bn per annum for many years now.

So the total loan book is an awful lot of money and nearly all of it is on the balance sheet of insurance companies backing annuity liabilities.

2.3 ERM Market in Ireland

There have been players in this market in the past, however our searches revealed nobody currently active.

The Bank of Ireland sold ERMs under its Life Loan brand, between 2001 and 2010. The final straw for lenders appears to have been the CBI's revised Consumer Protection Code, which came into force from 1 January 2012, and among other things required the following warnings:

> Warning: While no interest is payable during the period of the mortgage, the interest is compounded on an annual basis and is payable in full in circumstances such as death, permanent vacation of or sale of the property.
>
> Warning: Purchasing this product may negatively impact on your ability to fund future needs.

As far as we are aware there has been no new ERM lending in the Republic since then. According to the CBI in January 2016, at least 3,100 people owe a total of €300m on equity release loans in Ireland.

A company called Seniors Money which has sold ERMs in the past was reported to be looking to re-enter the market, but is not currently selling.

2.4 The Regulatory story in UK

Under the old Solvency I rules for both pillar 1 and pillar 2, the discount rate used on liabilities was derived from the assets held. Provided the terms of assets and liabilities were the same you took the gross redemption yield of the assets and knocked off whatever you thought was appropriate for the risk you were taking with those assets.

This was most evident in the approach to valuing annuities which were (and still are) largely backed by corporate bonds. It wasn't very scientific in its application. A percentage of the spread was taken, typically 50% to 60% (though a lot of work had been put into the theory to justify the magnitude). The argument was that a large part of the spread was there to compensate investors for the lack of liquidity of the asset, but that lack did not matter to insurance companies because annuity liabilities were not liquid either. The extra boost to the discount rate was called the liquidity premium (or illiquidity premium).

When Solvency II was first mapped out it had another approach to liabilities. The discount rate of liabilities was to be totally independent of the assets and to be based on the risk-free rate based on government bonds or interest-rate swaps.

This left the UK industry with a major headache. With its very large annuity portfolios the lack of liquidity premium left the apparent solvency of the industry in doubt and (the industry argued) would put up annuity prices hitting retiring voters. This issue was so important it even made it onto national TV.

Insurers elsewhere in Europe had various problems with Solvency II and, together, these concerns threatened the implementation of the whole project. Eventually compromises were found, continental insurers with long-term profit-sharing business got an ultimate forward rate for Euro liabilities far above market yields, while the UK got the matching adjustment ("MA").

The MA essentially allowed insurance companies to take credit for part of the spread that corporate bonds had above risk-free rate but subject to some fairly tight conditions. The rules did not say corporate bonds only, and for annuities only, but came pretty close.

Nevertheless, the principle of claiming higher discount rates than (liquid) risk-free because of the assets that you held was re-established, at least when those assets were fixed income investments. Naturally, the industry sought to extend that exemption to other asset classes with similar features. In particular ERMs came under consideration. Some companies were going to be hit very hard if they could not get their old illiquidity premium into the MA. However, the rules clearly said that the returns for the assets had to be fixed to qualify for MA.

An idea then emerged, if the portfolio of ERM could be divided up into senior and non-senior tranches maybe the senior bit could get the MA. This was generally accepted and happened. This could happen by internal restructuring of the ERM portfolio. An insurance company can shuffle its assets around without making any changes and, under Solvency II, the liability discount rate changes back close to what it would have been under Solvency I. Solvency II, however, remains a stronger basis than Solvency I for annuities because of the risk margin requirement (which did not exist under Solvency I).

Before Solvency II was live however the PRA issued its first discussion paper on the subject of ERM, DP 1/16. In it the idea of effective value was mooted.

This is the total value of all tranches of the restructured ERMs on the asset side of the balance sheet, plus the MA benefit arising from the restructured ERMs on the liability side of the balance sheet.

The Effective Value test sets an upper bound on how much benefit the company can take from the MA adjustment. It does not stop a lender making a "Day 1" profit.

The Discussion paper was then followed in December 2016 by a consultative paper CP 48/16. This included some basic principles laid out.

3.8 The PRA will assess the allowance made for the NNEG risk against its view of the underlying risks retained by the firm. This assessment will include the following four principles, which are explained in more detail below:

(I) securitisations where firms hold all tranches do not result in a reduction of risk to the firm; 1 The focus on the NNEG should not be taken to imply that other risks (eg prepayment risk) are not considered material by the PRA; and indeed Chapter 2 is clear that these other risks should all be considered in the internal credit assessment and FS mapping. 2 The PRA's rules on valuation are set out in Valuation 2.1 of the PRA Rulebook. 16 Solvency II: Matching adjustment - illiquid unrated assets and equity release mortgages December 2016

(II) The economic value of ERM cash flows cannot be greater than either the value of an equivalent loan without an NNEG or the value of future possession of the property providing collateral;

(III) The value of future possession of property should be less than the value of immediate possession; and

(IV) The compensation for the risks retained by a firm as a result of the NNEG must comprise more than the best estimate cost of the NNEG.

The CP also reinforced the PRA's endorsement of the EVT.

All four principles appeared again in the Supervisory Statement SS 3/17 which was released on 5 July 2017. The SS 3/17 was reissued with minor changes on 4 July 2018.

The next step came on 2nd July 2018 when the PRA published CP 13/18. This contained four key points

1) The rate of deferment for house prices should be at least 1% for valuing ERMs.
2) In assessing the value of the NNEG a central estimate of house price volatility should be 13%
3) That the old ICAS regime should be modified likewise
4) That companies should have 3 years to smooth in any impact of this change

At first sight it might seem strange that the old solvency regime that was supposed to have been superseded by Solvency II still needs updating. This is because of transitional measures on technical provisions ("TMTP"). These were intended to soften the impact of Solvency II by allowing said impact to be smoothed over 16 years. In practice in the UK this has been interpreted as allowing all pre-Solvency II business to be valued as it was in the good old days.

The IFoA provided a negative response to CP 13/18. They wrote

"Our main comments in relation to the proposals within CP13/18 are as follows:

- *it is in the public interest that the valuation of any NNEG on ERM assets is both robust and adequately reflects the corresponding risk;*
- *the proposals could have an adverse impact on the supply of equity release mortgages to consumers, and knock-on impacts on individual and bulk purchase annuity rates;*
- *the PRA should therefore have regard to both policyholder security and policyholder value for money when considering the impact of the proposals;*
- *we believe more research needs to be completed to understand if the NNEG is understated currently; this should be completed before any proposals are adopted;*
- *the retrospective nature of the proposals could give rise to a discrete shock in capital position for some firms. This would be disruptive to the industry, undermine confidence in the UK insurance system and increase the cost of raising capital due to the increased regulatory risk in the UK. This would be exacerbated if only a limited transitional period were available;*
- *furthermore, we do not understand why a change in estimate, which the proposed change in NNEG calibration seems to be, is being implemented as if it were an error;*
- *we suggest firms should be given a reasonable amount of time to prepare for the implementation of any new supervisory statement that follows CP13/18;*
- *the Effective Value Test (EVT) could lead to pro-cyclical behaviour by insurers;*
- *requiring the Individual Capital Assessment (ICA) to use the EVT is a significant departure from the practice assumed by firms in their ICA at the time of transition to Solvency II;*
- *we have some concern that the proposals are overly prescriptive. We would prefer the PRA to set out the principles and standards to be met."*

The authors of this paper are concerned that this response (particularly the fifth bullet) might lead some parties (Private Eye for instance) to misconstrue the IFoA's response as lobbying for the industry.

The PRA then issued PS 31/18 which summarised and largely rebutted the responses received. It did however allow two significant softenings. Firstly the implementation was put back to 31.12.19 with deferment rate of 0% rising to 1% by 31.12.21. Secondly the ICAS regime was not changed.

We also noted this comment (paragraph 2.131 of PS 31/18)

Nonetheless, the PRA does not agree with the principle underlying these comments that the regulatory treatment of insurance assets and liabilities should not change over their lifetime.

Evidently the regulatory view has evolved over time. That is to be expected; if we do not permit change in regulation then new classes of business will never be able to emerge, the effects of developments in actuarial science will be throttled and ultimately business may be either badly over or under reserved. The PRA current position looks reasonable to us. In Ireland where firms are not encumbered with large ERM back books, we would recommend the PRA's approach to the Central Bank of Ireland as a starting point for their deliberations.

As of writing the last stage in this saga was that the research commissioned jointly by the ABI and IFOA was delivered by Professor Tunaru and discussed at an IFoA meeting. The Tunaru paper contained many technical points which we consider in detail when we look at valuation assumptions.

Following the publication of Tunaru's paper, the IFoA published a press release reassuring the public that "valuations arising from insurers' current models and bases are sufficient"

This research contained one rather curious argument. In assessing the deferment rate, some parties give consideration to the rental yield that could be obtained while ownership of the property is being deferred. Tunaru estimates that as 5% but then multiplies that number by 20% because only 20% of the residential property is rented. We see no justification for this. It is apparent from the discussion at Staple Inn that there was little support for it.

His major conclusion is that a rather complex model "ARMA-EGARCH" is better than Geometric Brownian Motion for modelling house prices. We would not argue that this might incorporate features that have been observed in the market. However we are concerned that this may transpire to be an example of Burg Khalifa modelling (The Burg Khalifa is the tallest building in the world. If you jumped off the top of its 163 stories and measured, very quickly, what had happened to you after 162 stories, you would have a lot of data saying everything was going to be alright.)

3 Key Parameters for Valuing Equity Release Mortgages

In this chapter we describe aspects of ERM valuation. We focus on fair value measures. Similar procedures apply for IFRS and for valuation under Solvency II.

While aspects of the methodology for valuing equity release mortgages have been contentious, there is also much common ground. In particular, the concepts of deferment rate, discount rate and implied volatility can apply to pretty much any method. These concepts provide a common language within which we can compare and evaluate different valuation approaches.

3.1 Common Features of Valuation Models

Most valuation methods start by analysing the possible future date of borrower death (including the small proportion who move into long-term care). This follows a relevant mortality table reflecting the age, gender and (sometimes) the wealth of the borrower. The ERM is then a probability-weighted sum of valuations over possible dates of death.

Whatever the model being used, we can refer to a fixed-term ERM as the hypothetical value of an ERM where the borrower's date of death is known. Then the stated ERM value will be a weighted average of fixed-maturity ERM-lets, weighted according to the probability distribution of the date of death. If deaths occur at year ends, the ERM value for a life aged *x* then takes the form of:

$$ERM_{total} = \frac{1}{l_x}\sum_{t=1}^{\infty}(l_{x+t-1} - l_{x+t})ERM_{let}(H_0, e^{Rt}LTV.H_0, t)$$

Here, ERM_{let} refers to the value of a fixed-maturity t, current house value H_0 and initial loan to value ratio *LTV*, accruing compound interest at continuously compounded rate *R*.

For each ERM-let, the value is an increasing function of the value of the house, because the larger the house value, the more secure the ERM-let. Likewise, the ERM-let value is an increasing value of the projected mortgage balance on the maturity date. Most models are homogeneous, in that if both the house value and mortgage balance double, then so does the ERM-let value.

Under many interpretations of IFRS, acquiring an asset should not give rise to a profit on inception. This means that certain pricing inputs are back-solved to price new ERMs at the amount of the sum advanced. Those back-solved assumptions are then used to calculate the fair values of ERMs from earlier vintages.

3.2 Voluntary Early Redemptions and Further Advances

One of the more difficult technical areas in ERM valuation is the modelling of decrements; not just mortality but moving into long-term care and voluntary early redemptions. It is plausible that mortality would not depend on house prices, and while the mechanism for entering long term care could involve financial decisions, this typically applies to a small proportion of borrowers, often only a few months from the end of their lives.

Voluntary early redemptions are a more difficult matter. A borrower paying a high fixed rate on an ERM may consider refinancing if interest rates fall and they can secure a lower rate in the market. Furthermore, refinance involves a reassessment of the house value, leading potentially to a selection effect where customers whose houses have fallen in value are more likely to keep their ERMs.

Allowance for dependence of voluntary redemptions on house prices or interest rates generally breaks closed-form expressions such as Black's formula, and a numerical method such as Monte Carlo simulations may be required. In this note we focus on the more commonly used approaches where voluntary early redemptions are considered, like mortality and entering long term care, as being stochastically independent of house prices. This might most plausibly be the case for ERMs with high or variable prepayment penalties

3.3 Limiting Behaviours and Extreme Ratios

Tools for valuing ERMs need to be able to cope with a range of initial house prices and forecast loan balances. We can describe models in terms of assumptions, but we can also work backwards and define implied parameters from a given ERM valuation model, in order to compare different models.

Let us consider a series of fixed-term ERM-lets, of some term T years, lent on a house with current value S_0 and different mortgage balance $K = e^{Rt}LTV.H_0$. We ask what happens to those ERM values as the mortgage balance becomes very large. Ultimately, the ERM-let value is limited by the lender's ability to recover value out of a finite house.

We can denote the limiting ERM-let value as:

$$\lim_{K\uparrow\infty} ERM_{let}(H_0, K) = e^{-qT}H_0$$

Here, q is the *implied deferment rate*. The value of q might be positive, implying that the right to a house in the future is less valuable than a house now. Alternatively, for some models q might be negative, if house prices are assumed to grow at a rate higher than used for discounting the ERM value. The appropriate choice of q is possibly the most contentious assumption in ERM valuation.

For some (not very good) models that limit might not exist at all. For example, a firm might decide to discount the promised mortgage flows at one rate and the NNEG at a different (lower) rate for prudence. The problem with this approach is that K gets larger, eventually we could find the NNEG is worth more than the mortgage (without NNEG) so the ERM is a liability, not an asset. We need to take special care over discount rates to prevent that happening.

We can also consider a different limit, looking at ERM-let values with a given future mortgage balance *K*, and letting H_0 tend to infinity. In that limit, the mortgage gets more and more secure, as the NNEG tends to zero. We denote the limiting ERM-let value as:

$$\lim_{H_0 \uparrow \infty} ERM_{let}(H_0, K) = e^{-rT} K$$

Here, *r* is the discount rate for the ERM pricing model.

Our arguments that ERM_{let} be an increasing function of both H_0 and in *K* implies the PRA's principle III from CP 48/16.

These quantities are illustrated in the chart below, which for a house with unit value H_0 = 1 shows how the ERM value might depend on the amount lent:

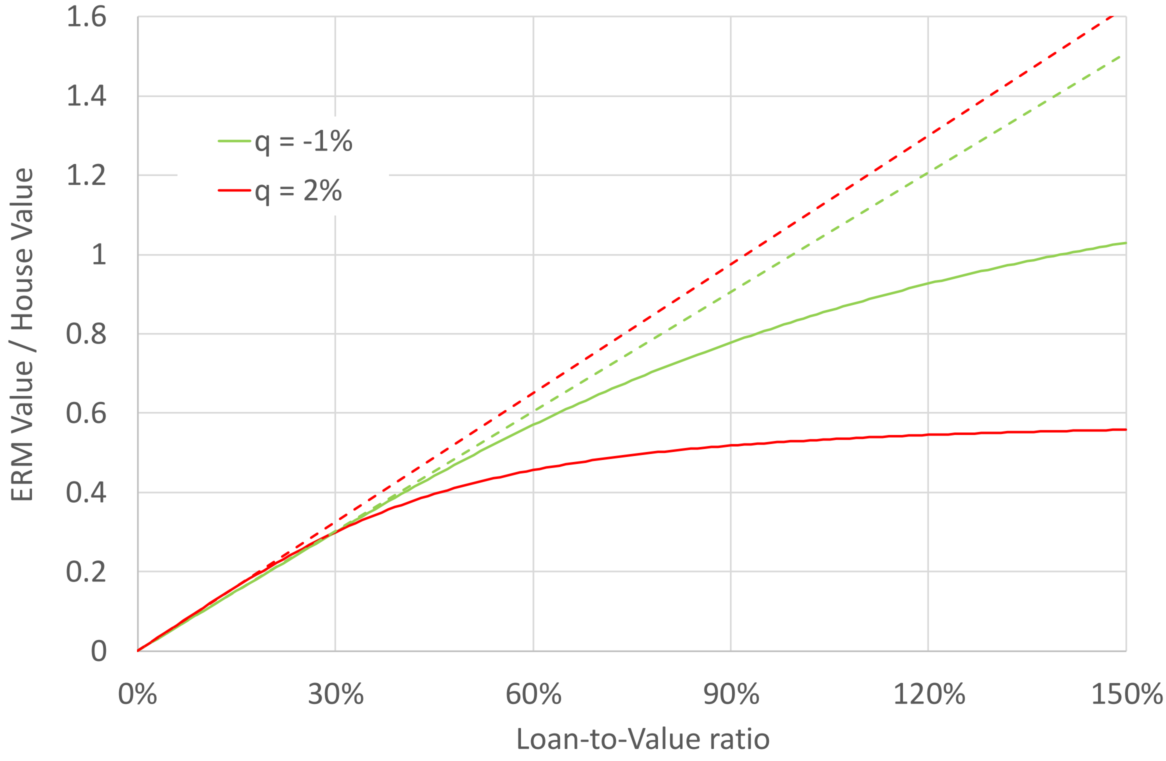

When the LTV is low, then the ERM is very secure to the extent that the price of the house is almost irrelevant. The gradient of the tangent the left (shown as a dotted line) is then $e^{(R-r)T}$ where *R* is the roll-up rate, *r* is the discount rate and *T* is the term of the ERM-let.

We usually want an ERM to be worth its face value for new lending, which in this chart we have assumed occurs at LTV = 30%. As the ERM is a concave function of LTV, that means we will always end up with $r < R$.

The limit to the right is e^{-qT}. Some (including the PRA , principle III of CP 48/16) argue that $q > 0$ always, so that a loan secured on a house is never worth more than the house itself. Indeed, the PRA argued further in CP 13/18 that $q \geq 1\%$. Others argue that *q* could be negative if the house is assumed to grow fast enough relative to the discount rate.

In this paper, *q* and *r* will always refer to the deferment rate and discount rate for an ERM. Where we discuss related concepts, such as rental yields or risk-free rates, we will use appropriate language but will not use *q* or *r* for any purpose other than the parameters implied from an ERM model.

3.4 Option Pricing Formulas

3.4.1 Implied Volatility

We now move to the more complicated general case, where the house is not vastly more valuable than the mortgage balance, nor is the mortgage balance far more than the house is worth.

In that case, we can say the following about feasible bounds for the ERM-let value:

- It is a positive number (because it is not a liability)
- It is less than $e^{-qT}H_0$ (because that is the increasing limit as K tends to infinity)
- It is less than $e^{-rT}K$ (because that is the increasing limit as S_0 tends to infinity)

Now let us consider the valuation formula:

$$ERM_{let} = e^{-rT}K\Phi\left(\frac{\ln(H_0/K) + (r - q - \sigma_H^2/2)T}{\sigma_H\sqrt{T}}\right) + e^{-qT}H_0\Phi\left(\frac{\ln(K/H_0) \ + (q - r - \sigma_H^2/2)T}{\sigma_H\sqrt{T}}\right)$$

Here, Φ represents the cumulative standard normal distribution function Φ. We call σ_H the individual house price volatility.

We notice that the right-hand side is:

- A continuous, decreasing function of the parameter σ_H
- Equal to min{ $e^{-qT}H_0$, $e^{-rT}K$ } in the limit as σ tends to zero.
- Equal to zero in the limit as σ tends to infinity.

Therefore, for any ERM-let valuation model satisfying the feasible bounds for the ERM-let value, there is a parameter σ, the *implied volatility*, for which the valuation formula replicates the model output.

We have no a priori reason to believe that the same σ would apply for different ERMs. If the valuation model is homogeneous (most are) then σ would be a function of the ratio K/H_0. Homogeneity might fail if, for example, property value is used as a proxy for wealth in forecasting mortality / morbidity, or if per-policy expenses are included in the valuation.

If the Black-Scholes model holds, then the same σ would hold for all K, H_0 and T.

3.4.2 Other versions of Black's Formula

There are several equivalent ways of writing this valuation formula. We can write as a mortgage minus a NNEG (put option), that is:

$$ERM = e^{-rT}K - \left[e^{-rT}K\Phi\left(\frac{\ln(K/H_0) \ + (q - r + \sigma_H^2/2)T}{\sigma\sqrt{T}}\right) - e^{-qT}H_0\Phi\left(\frac{\ln(K/H_0) \ + (q - r - \sigma_H^2/2)T}{\sigma\sqrt{T}}\right)\right]$$

We can also write it as a deferred house minus a call option, with the option reflecting excess house value to heirs:

$$ERM = e^{-qT}S_0 - \left[e^{-qT}H_0\Phi\left(\frac{\ln(H_0/K) + (r - q + \sigma_H^2/2)}{\sigma\sqrt{T}}\right) - e^{-rT}K\Phi\left(\frac{\ln(H_0/K) + (r - q - \sigma_H^2/2)}{\sigma\sqrt{T}}\right)\right]$$

The same formulas can written in Black form, where H_0 is replaced by $e^{(q-r)T}F$ and F is the forward value. These representations are equal because Φ(-z)= 1- Φ(z).

The difference *r-q* is sometimes called the house growth rate.

3.4.3 Example Assumptions and their Consequences

Among the different financial assumptions required to value an ERM, the most disputed assumption has been the deferment rate, with most estimates of volatility in a narrow range from 11% to 13%. In 2019, two new pieces of research have produced substantially divergent volatility estimates, which we consider further below.

Source	Deferment (per annum)	Volatility (per annm)
Historic HPI ex Cash	-0.58%	11%
Hosty et al (2007)	0.25% (real world), 3.3% (risk neutral)	11%
Sportelli	4.75%	N/A (we have used 11% in the chart below)
CP 13/18	2% central, 1% lowest acceptable.	10% to 15%, with best estimate 13%.
Buckner & Dowd	2% to 4%, best estimate 3%.	13%
Tunaru	0.66% proposed, calculations use 1%.	3.3%
Just Group December 2018	0.3%	13%
Eumaeus blog 27/02/2019	N/A (we have used 3% in the chart)	16% to 24%

The chart shows these sets of assumptions, together with contour lines corresponding to equal NNEG values, with higher NNEG to the top right.

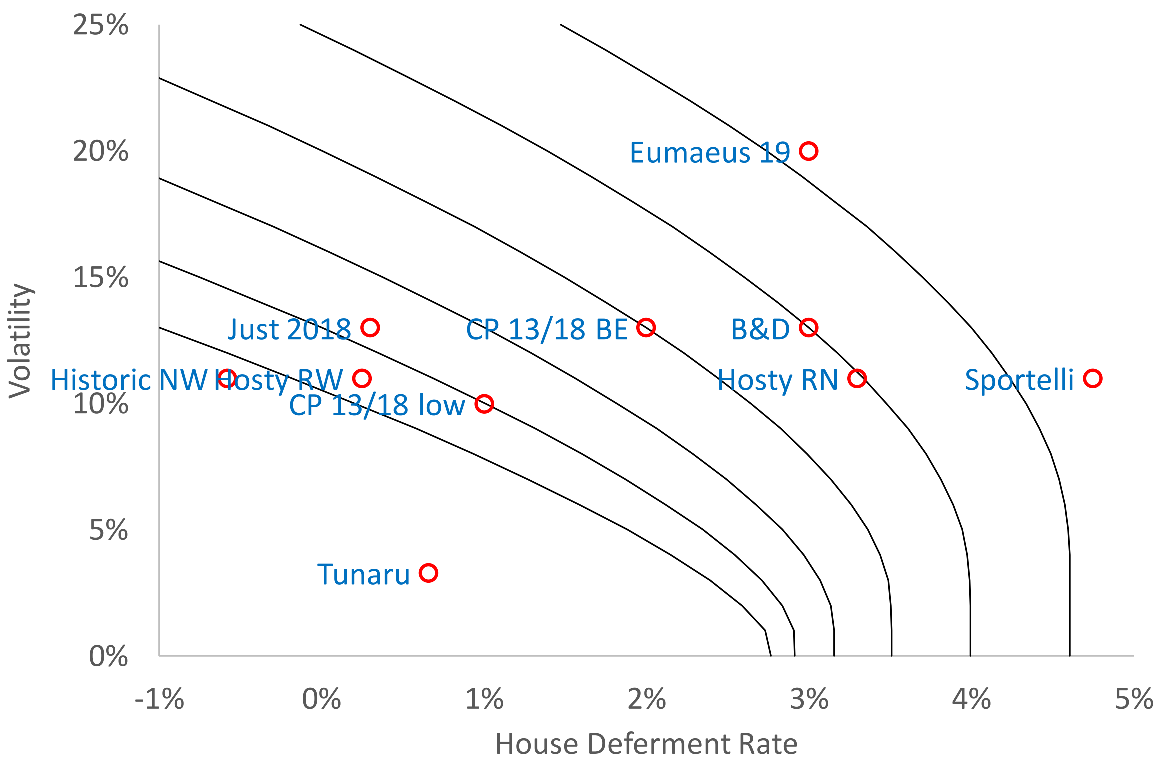

The contour lines appear in slightly different places according to assumptions for the ERM-let. In this picture, we have assumed LTV = 30%, discount = 2.8%, rollup = 5% and term = 25 years.

4 Evidence for Setting Assumptions

4.1 Discount Rates

The discount rate for an ERM is theoretically the easiest part. Aside from longevity risk, an equity release mortgage is a fixed-income instrument, akin to a bond. There are many bonds in issue whose yield can be observed, and we might use one of them to discount an equity release mortgage.

There are several practical points of detail, which we now consider in more detail.

- What bond should we use as a benchmark – what credit risk, and what degree of liquidity?
- How to allow for expected defaults?
- The limiting case of high loan-to-value ratios
- Discounting the NNEG separately from the mortgage itself.
- Zero profit on inception

4.1.1 Choice of Benchmark Bond

Ideally, we could compare an ERM to a similar bond: similar in terms of currency, term, credit risk and liquidity.

The simplest comparison would be government bonds, of similar term to the ERM. There is a deep and liquid market, with prices publicly reported. But government bonds are safer and more liquid than ERMs, so that probably gives a discount rate that is too low.

The main alternative is to use corporate bonds. We might have a measure of credit risk for the ERM, such as the probability of NNEG coming into the money, and then look for a bond whose credit rating would give a similar probability of default. The difficulty is that any bond whose yield we can observe must be more liquid than an ERM. So the best we could do is find a corporate bond matching term, currency, and credit risk of the ERM, recognising that any illiquidity premium appropriate to the ERM might be even higher than that on a quoted corporate bond.

The prevalent industry approach has been described by Hosty et al (2008), and more recently by Kenny et al (2018). Kenny et al report an industry survey in which "The general consensus for calculating a discount rate [for the ERM including the NNEG] was to use a risk-free term structure, plus illiquidity premium". Possible methods for estimating that illiquidity premium include:

- "based on market prices of ERMs to be the rate at which transactions eliminate the day-1 gain
- based on implied liquidity premiums seen in the market for other assets (e.g. bonds)"

4.1.2 High loan-to-value equity release mortgages

Although equity release mortgages may start with a low loan-to-value ratio (perhaps 30% on inception), this LTV ratio can increase with the passage of time, as interest accrues on the mortgage balance, while house prices may not grow as fast as the debt and may even fall. Many ERM writers have existing books of business with considerably higher LTVs than new business.

For a high LTV equity release mortgage, the NNEG comes into the money and the ERM investor essentially has an exposure to the housing market (or more accurately, each ERM has an exposure to the value of one particular house, none of which necessarily follows the market as a whole). It therefore make sense for the discount rate to reflect the assumptions made about house market returns. In particular, if a firm has assumed that house prices will continue to grow faster than inflation, and at the same time rates of rent will be maintained, this implies that investors earn a significant risk premium in the housing market. Arguably, the discount rate should reflect the same risk premium in projection and discounting. This is the essence of principle IV from CP 48/16.

As far as we are aware, it is not industry practice to reflect house market risk premiums in ERM discount rates; nor is it usual to increase the discount rate when the NNEG is closer to the money. The fact that insurers do NOT make these risk adjustments is fundamental to current market practice. To the best of our knowledge, there has been no attempt to defend the lack of risk-sensitive discount rates from a theoretical perspective. The use of 'illiquidity premium' terminology to describe the discount rates has perhaps deflected difficult questions about how much of this is a market reward for taking house market risks. We will give further details later on the liquidity arguments.

4.1.3 Separate discounting of the NNEG

In our presentation we have taken a single discount rate r for both the loan portion of an ERM and also for the NNEG.

It could be argued that discount rates should reflect the risk of individual cash flow components. The loan component is low risk (although illiquid) and might be discounted using yields on comparable illiquid bonds. The NNEG might be valued more like an insurance policy, reflecting the value that homeowners might place on having a guaranteed price at which they can sell their house. This could imply a low or even negative discount rate. We will later see that this is what option pricing theory could imply.

In this calculation it is important to distinguish expected cash flows from promised cash flows. The loan portion of the ERM (without the NNEG) corresponds to the promised cash flows on a bond (with no allowance for defaults). With separate discounting of the NNEG, the appropriate comparison discount rate is a bond yield. If the loan and NNEG are discounted together, then the bond yield minus expected defaults is a more consistent discount rate.

It is not industry practice to discount the NNEG at a lower rate than the loan portion of an ERM.

4.1.4 Zero profit on inception

A common approach in industry is to derive a discount rate such the reported profit on new ERM lending is zero. That same "market implied" discount rate is then used to value older vintages of ERMs. We consider this in (much) more detail later.

4.2 Considerations in the Deferment Assumption

Deferment rates reflect the extent to which a promise to receive a house at some point in the future is less valuable or (exceptionally) more valuable than receiving that house now.

There are several possible ways to estimate this parameter:

- By reference to rented properties: a consideration of rental yields.
- For owner-occupied leasehold properties: an analysis of relativity curves.
- As the difference between a discount rate r and assumed house growth g.
- As an implicit parameter derived from new rates of ERM lending; in this calculation either the discount rate or (less commonly) the house growth rate are fixed, and the deferment rate q is solved to equate a new ERM value to the amount advanced.

4.2.1 Rental Yields

Rental yield information on residential properties is widely available, for example from the ONS. This information typically implies gross yields in the region of 5% at the time of writing, that is, tenants across the UK would be paying rents of the order of 5% of the open market value of the house they are renting.

However, the net rent received by an investor might be smaller than this, for a number of reasons including maintenance and repair, management expenses, voids, taxation and rent arrears. That might reduce the 5% to a 3.5% net figure that an investor might reasonably expect. This investor comparison is relevant because ERM writers are primarily investors; purchasing ERMs is an alternative to direct property investment as a way of gaining exposure to the housing market.

What then can we say about the perspective of an owner-occupier, a category which includes ERM borrowers? They may not consider their house as an investment but rather as a place to live. People do not charge themselves rent, but we can ask what value they put on having a home to live in. The natural estimate of this value is the saved cost of having to rent somewhere else. In that case, assessment of rental yields for ERM valuation would simply consider aggregate rents and aggregate values of rented properties. A minority view (Tunaru, 2019) is that owner-occupied properties should be treated as tenants playing zero rent, in which case the observed yield is diluted, by a factor of five (because, in the UK, roughly 1-in-5 properties is rented).

There is a useful analogy with other commodities; oil for example. A factory operator may know they require a quantity of oil in in one year's time, which they can acquire in the forward market. There is also a spot market for buying oil now. Paying later would enable the operator to earn interest on the price in the meantime, but the difference between spot and forward prices of oil is not explained entirely by reference to the risk-free rate. The balancing item is sometimes called the convenience yield, that is, the value a factor operator places on having the oil in storage and immediately available if needed. The fact that oil produces no direct income does not imply a deferment rate of zero. We cannot expect to estimate house deferment parameters from oil determent parameters but we can read across the principle that an asset producing no income can still carry a convenience yield.

4.2.2 Leasehold Enfranchisement

There is another way to assess the value of deferred house interest compared to current ownership, and that is to look at the value for leaseholds, that is, ownership which terminates at a fixed future date. There were once active markets in leases, which enabled analysts to deduce market deferment rates by comparing leases on different properties. More recently, leaseholders have been granted rights to buy out the lease from the lease owner, effectively acquiring a deferred interest in their own home, on top of their existing time-limited right of enjoyment. The Royal Institute of Chartered Surveyors collates data on the prices at which these transactions take place, publishing leasehold relativity curves from time to time (RICS, 2009). These typically imply a deferment rate of around 4% per annum. This is a complex legal area, and it is not clear whether the enfranchisement prices should be interpreted as market prices of a deferred interest in a house, or if they are a strike at which a lessee has exercised an option (Radevsky & Greenish, 2017).

The *Sportelli formula* for deferment is also based on a leasehold enfranchisement court ruling, and proposes a deferment rate of 4.75%. This is based on estimates of house price growth; we consider those in more detail below.

4.3 Considerations in the House Price Growth

House price growth might be determined in one of three ways:

- By reference to historic house price growth
- Calculated as a discount rate minus a deferment rate
- Computed as an implied parameter from new lending, with either the discount rate or the deferment rate held fixed.

In the remainder of this section we consider historic house price experience.

4.3.1 Summary of Historic Experience

Historic house price indices are published by two building societies: Nationwide and Halifax, the latter now bring part of the Bank of Scotland group.

The longest time series comes from Nationwide. More recent data is provided monthly, but there is a quarterly series of annual house prices starting in Q4 of 1952. At that point, the average house price was £1,891. By Q4 of 2018, this had risen to £214,178, an annual growth rate of 7.43% between 1952 and 2018.

4.3.2 Indexation and Revaluation

For analysing historic house price movements, we can easily allow for the effect of inflation or interest rates because there are good historic records of both. In addition, expectations of future inflation and interest rates can be inferred from market prices of conventional and index-linked gilts. This allows us to modify a house price forecast to allow for expected interest rates or inflation to be different from how they were in the past.

The Office for National Statistics has published the retail prices index monthly since June 1947. The series has been re-based several times since then, most recently in January 1987 where the RPI was re-set to 100. The figure for December 2018 is 285.6. The December 1952 index value, corresponding to the start of the Nationwide house price index, linking several past RPI series from the ONS in the UK, comes out at 10.15, so the average annual rate of inflation was 5.19% pa. The corresponding real house price growth was 2.13% above inflation. We should caution that the RPI itself contains an element of house price inflation within it, albeit a small one. As a result, house prices in excess of RPI would be slightly less volatile that house prices measured relative to an inflation index that did not include housing costs.

For cash returns, one might refer to inter-bank rates such as LIBOR, but these have limited history. For a fairer comparison, we have referred to Bank of England base rates, for which a compete history is available back to the bank's foundation in 1694. At the end of 1952, the base rate was 4.00%, while at the end of 2018 it was 0.75%. The interest is calculated daily at the stated rate divided by 360. The average annual return over the period from 1952 to 2018 was 6.81%. The annual increase of the house price index divided by the cash account was 0.58%.

The table below shows some intermediate values.

Year End	House Price (Nationwide)	RPI (ONS)	Cash Rollup BoE base rate
1952	1891	10.15	100.00
1958	2068	12.40	132.49
1968	4089	16.97	232.54
1978	16823	51.76	572.26
1988	57245	110.30	1876.33
1998	66313	164.40	4523.85
2008	156828	212.90	7370.59
2018	214178	285.60	7749.34

4.3.3 Assumptions at Current Bond Yields

Relative movements are helpful to study, because we have direct, market-based estimates of future inflation and risk-free rates by reference to the markets in conventional and index-linked bonds. If we can determine a good estimate of house price growth relative to cash or inflation, we can then use the bond markets to deduce estimates of house prices.

The Bank of England published the following yields at the end of 2018. These are continuously compounded spot rates:

Years	Nominal	Real	Inflation
5	0.91%	-2.19%	3.10%
10	1.31%	-1.99%	3.29%
15	1.67%	-1.84%	3.51%
20	1.87%	-1.71%	3.58%
25	1.92%	-1.59%	3.51%
30	1.87%	-1.51%	3.38%
35	1.78%	-1.47%	3.25%
40	1.71%	-1.46%	3.17%

We can convert these into (annually compounded) house price growth forecasts based on nominal house price growth, inflation pus a margin, or cash plus a margin. The resulting figures are:

Horizon (years)	Annual spot nominal HPI Growth Forecast, based on …		
	Historic growth	Inflation + spread	Cash + spread
5	7.43%	5.35%	1.49%
10	7.43%	5.55%	1.90%
15	7.43%	5.78%	2.27%
20	7.43%	5.85%	2.47%
25	7.43%	5.79%	2.53%
30	7.43%	5.65%	2.47%
35	7.43%	5.51%	2.38%
40	7.43%	5.42%	2.32%

These columns are strikingly different to each other. This is because future inflation expectations (as implied from gilt markets) are far below historic levels of inflation, and future cash returns are expected to be even further below historic cash returns.

We show the historic data, together with forecast values of for each index (all three shown for HPI) in the chart below. We have re-based the indices to start at 100. Note that the vertical axis is on a log scale:

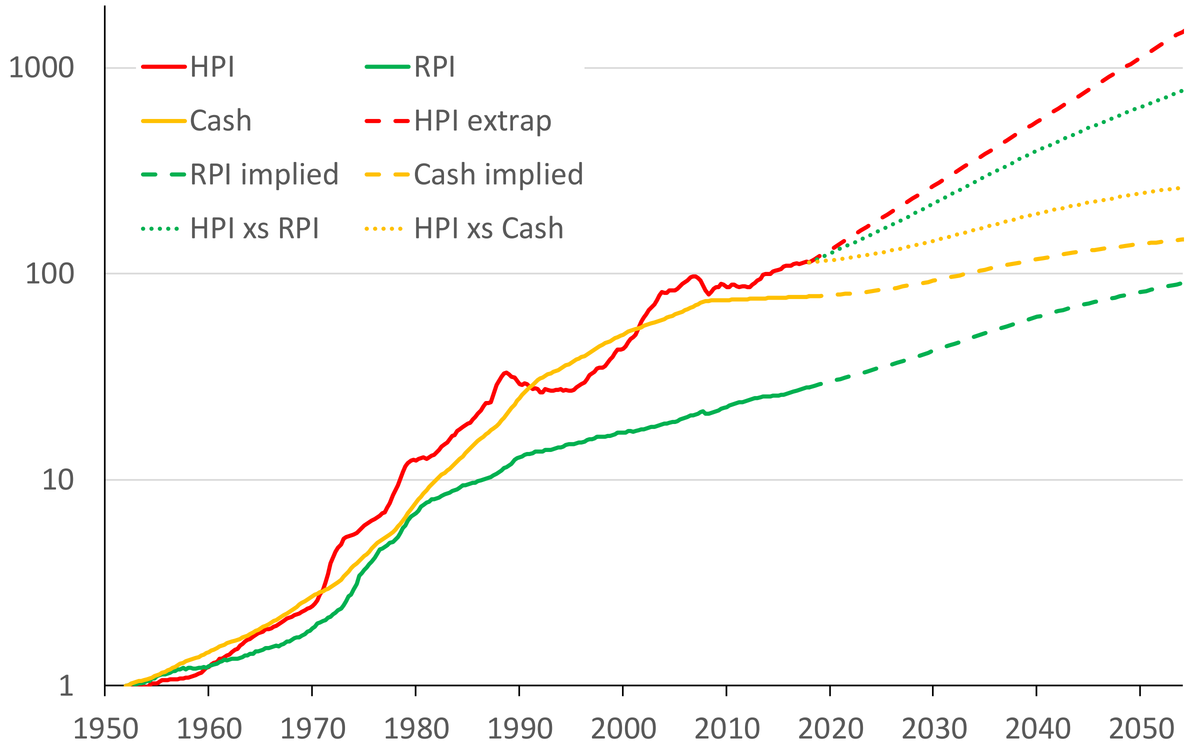

To add one further complicating factor, the difference *r-q* should theoretically represent an *arithmetic* mean. Our analysis has used geometric means, looking at compound returns over *n* years and taking the n^{th} root. The arithmetic mean is always larger than the geometric mean, unless all observations are equal in which case both means are the same. Therefore, something small and positive (a convexity adjustment) should be added to the geometric mean calculation to correct for the difference between arithmetic and geometric means.

4.3.4 Was House Price Growth due to Risk Premiums, or Luck?

The last sixty years have seen more or less steady growth in UK house markets, often outpacing more conventional risk assets such as equities. This raises the question as to whether this period of historic growth is typical, or an anomaly. The period from 1970-2000 was also a period of historically high inflation and interest rates. There has been no armed conflict, revolution or military coup on mainland UK soil over this period. There are socio-demographic trends, including a well-documented rise in middle-class household income as a result of more families with two earners (Rouwendaal & van der Straaten, 2003). The cause and effect are disputed: did house prices rise because families could afford more, or were parents of young children driven more quickly back to work because of unaffordable housing?

Analysis in Hosty et al (2008) shows that house price growth in the UK and Netherlands between 1970-2004, has been stronger than in many other countries. Ireland had relatively low house price growth during the same period. It may be that the UK's growth contained an element of good luck (for people who owned houses); arguably the more modest growth in other countries is relevant in estimating how the UK house market could fare in future.

Some writers of ERMs have published their house price growth assumptions in their annual statements. These typically cite forward-looking considerations and peer comparisons in addition to the historic data. For example, Just Group (2019) states:

"During 2018 the Group reviewed and updated its property growth and volatility assumptions, which are the key inputs into the valuation of the NNEG included in our lifetime mortgage assets. The property growth assumption has been reduced from 4.25% to 3.8% per annum, and the property volatility assumption has been increased from 12% per annum to 13% per annum. In updating these assumptions, the Board took into consideration future long and short-term forecasts, benchmarking data, and future macro-economic uncertainties including the possible impact of Brexit on the UK property market."

4.3.5 Derivation of Sportelli's Formula

An often-quoted court case (Sportelli vs Earl Cadogan) resulted in a deferment rate of 4.75%. This was calculated assuming landlords earn real house price growth of 2.0%pa (roughly in line with historic UK experience), and rental income q. This is equated to a real risk-free rate of 2.25% plus an assumed risk premium of 4.5% per annum (Wilson, 2003). Together, these imply a balancing item q = 4.75%. Under current market conditions, with negative real yields, the implied deferment rate would be lower than 4.75%.

4.3.6 Effect of Dilapidation

It is the nature of physical assets that, without attention, they dilapidate over time. With investment the assets can become more productive and more valuable. This is true for direct investments in property, just as for assets owned by companies. If companies fail to invest in their assets then dividends are higher, but growth lower, than would otherwise be the case. Likewise, if a landlord chooses to reinvest some of the rent in property improvements, their short-term cash income is reduced but long-term growth enhanced. Corporate debt typically has covenants restricting large dividend payments, while ERMs require borrowers to maintain their property. These clauses have the same purpose; to protect lenders against the risk of dilapidation of the asset on which the debt is secured.

When combining historic house price and rental income series, it is important to ensure consistent treatment between income and house price growth. From the rent paid, enough must be deduced to maintain the property in line with the condition underlying the price indices, in order to calculate the effective investment income.

When forecasting house price growth for purposes of ERM valuation, it may be that the population of houses under ERM loans grow faster, or more slowly, than the index as a whole as a result of differential rates of investment in the property. In that case, an allowance would need to be made – a so-called *dilapidation adjustment*. This is simultaneously an addition to the deferment rate and a subtraction of the growth rate, to capture that cash has been diverted from maintenance to boost immediate income, or (if the dilapidation adjustment is negative) that the property has seen a greater investment than on the reference index. This could in theory result in a negative deferment rate. If for example someone buys and empty field with wooden shack, with plans to build an executive home on the site, then an investor will pay more for the home in a year's time than for the shack now. This commonly happens when new-build properties are sold off-plan.

In this note, for simplicity we treat a dilapidation adjustment as a constant rate per unit time. Non-uniform dilapidation might also occur, for example if old age impairs someone's ability to maintain their home. One off spending on home improvement (which is sometimes what borrowers spend the ERM money on) would be another example of non-uniform dilapidation.

Finally, we should be alert to the possibility of selective dilapidation. When someone maintains or improve their home, they not only get the enjoyment of an improved property but also get to realise a more valuable asset than can be sold or passed to heirs. If NNEG in the money then only benefit of home improvement is temporary enjoyment; the increase in value accrues to the lender. It might be thought that, despite covenants to the contrary, owners of homes that have fallen below an ERM balance might be particularly reluctant to keep their property in good condition.

4.4 Considerations in the Volatility

4.4.1 Implied Volatility and Time Series Volatility

Under a Black-Scholes model, the implied volatility of an option is equal to the volatility of the underlying geometric Brownian motion, that is, the standard deviation of log price changes per unit time, which is assumed to be constant.

In a real-world model, we also want to estimate the standard deviation of log price changes, usually over the horizon relevant to the NNEG, that is, over years or even decades.

4.4.2 Historic Index Volatility

To calculate historic volatility over a given holding period, a possible procedure is as follows:

- Collect historic data on log price changes over periods over a chosen holding period length, that is the length of time over which a change in house prices is measured. These may be allowed to overlap, especially for longer holding periods, in order to increase the number of data points.
- A data window must also be selected. For example, we might choose to look at changes over annual holding periods, but with start dates falling quarterly within a ten-year window.
- Given the data window, the standard deviation of holding-period changes is calculated over start dates within the window.
- The annualised volatility is defined as the standard deviation of changes over that holding period, divided by the square root of the length of the holding period (in years).

In our calculations, we have varied the holding period (on the horizontal axis) and chosen the longest possible time window so that we use every available data point. Because of the autocorrelation in house market returns, the use of shorter windows can give significantly lower stated volatilities.

When we apply this to many series, such as equities or currencies, the annualised volatility is approximately constant as a function of the length of holding period. House prices do not behave like this; the annualised volatility shows a marked upward slope.

The chart shows annualised volatility based on the Nationwide quarterly series. These are measured on three different bases: nominal (raw HPI), real (in excess of RPI) and risk premiums (in excess of cash).

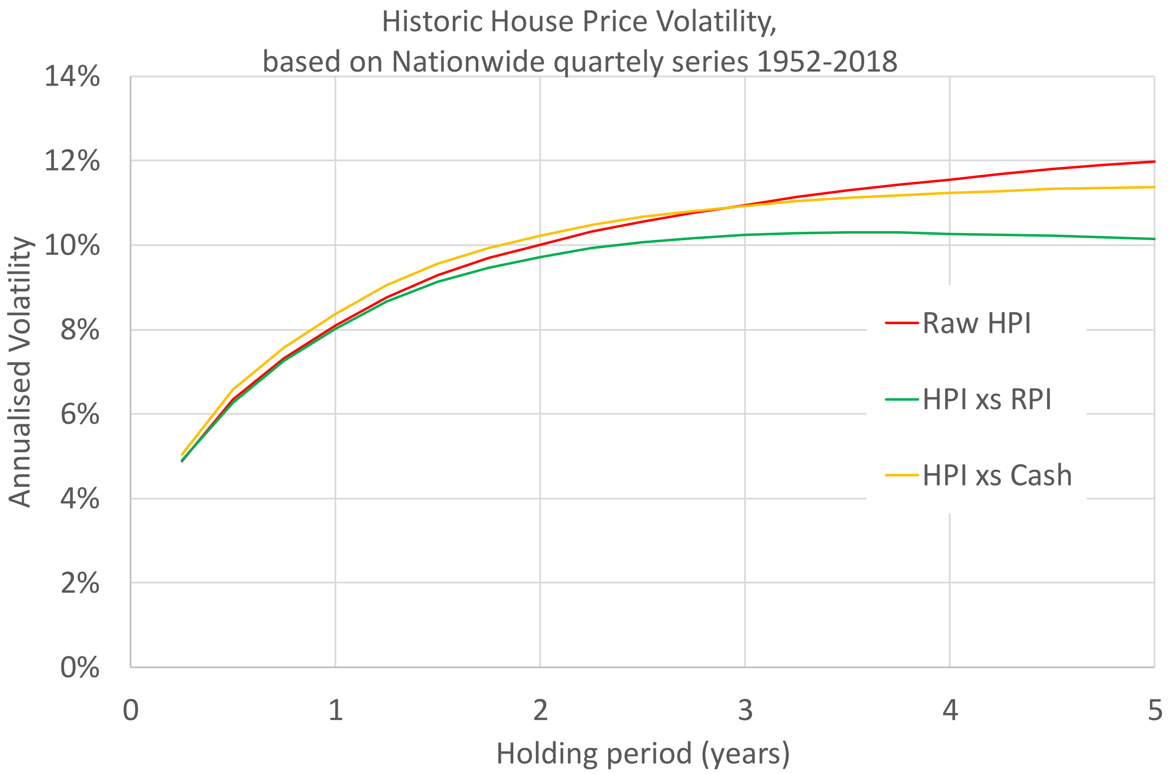

The upward slope leaves us in a quandary regarding the part of the curve to use in the Black-Scholes formula. There is a further complication that the historic volatilities we have shown are unconditional, in that they average many periods of the same length even if the market conditions at the starts of those periods were different. For option pricing, we would ideally use a conditional measure of volatility given current conditions, which would generally be lower than the unconditional volatility, although this may temporarily be reversed in particularly volatile market conditions.

The relevant point of the curve for option pricing depends on the arguments being used to value the option.

- Most authors have argued for using the volatility of the term of the ERM. This implies extrapolating the data to the right, approaching a limit which is often estimated at around 11%.
- In contrast, for dynamic hedging arguments, the relevant term of the volatility is the frequency of hedging. For example, if a hedge portfolio is rebalanced at a frequency of three months, then the relevant volatility is that at three months. In his recent paper, Tunaru (2019) used monthly data and a risk-neutralisation argument to arrive at an annual implied volatility in the range from 3.2% to 3.5%.

4.4.3 Individual House Volatility

We denote the index volatility by σ_I. It is this which we measure from historic house indices. A portfolio of ERMs, however, is a portfolio of options on different properties. This is not the same as an option on a portfolio of properties. A similar distinction exists in option markets between single stock options and equity index options, with the equity index options having much lower implied volatility as a result of diversification within the index. The index volatility is not a simple average of individual house volatility, because diversification between different houses causes a reduction in the index volatility.

Data on individual houses is more difficult to analyse. The Land Registry contains prices of sales, so for homes that have changed hands twice or more in the recent past, this can give an indication of house price changes. More comprehensive data has been analysed in the context of commercial property (Mitchell, 2015).

The trajectory of an individual house relative to an index can have several components. Firstly, there is the idiosyncratic noise of that particular house. Secondly, there may be systematic attributes of a particular house, such as geographic location, number of bedrooms, construction type, which are not representative of the wider index. Lastly, there can be differences in dilapidation because of incentives specifically related to the existence of an ERM.

More useful for ERMs is the data set published by Aviva, in relation to ERMs it has securitised. These are available on the equity release funding web site (Aviva, 2019). These are a series of files that include the initial valuation, indexed house price and realised sale price for all properties generating a loss, that is, in which the NNEG bites. It is remarkable that in virtually all of these cases, there would have been no loss had the house price followed the index. It is the basis risk (or stochastic dilapidation) relative to the index that is generating the losses.

We denote the volatility of stochastic dilapidation by σ_D, and assume this is independent of the index itself. The individual house volatility σ_H then satisfies the Pythagoras rule:

$$\sigma_H = \sqrt{\sigma_I^2 + \sigma_D^2}$$

We have not found much research on the estimation of stochastic dilapidation specifically for residential property, but it seems that plausible estimates of dilapidation volatility are at least as large as the index itself. If, for example, we took $\sigma_I = \sigma_D = 11\%$ that would imply $\sigma_H = 15.5\%$

5 Theories of Illiquidity Pricing

5.1 Valuation Methodologies

5.1.1 The Need for Theory

We have looked at historic evidence for calibrating valuation parameters, including these three: a discount rate *r*, a house rental *q* and a house growth rate *r-q*.

There are debates about the correct values, but under current market conditions for the maturity of a typical ERM, we would be looking at risk-free rates of around 1.8%, historic house price growth of 7.4% and rental yields of around 3%. It is clear that the house price growth plus the rental yield is far higher than the risk-free rate. The debate about ERM methodology has then focused on whether the historic house price growth should be ignored, to be replaced by the risk-free rate minus the yield (sometimes called the risk-neutral approach) or, alternatively, if the actual rental yield should be ignored, to be replaced by the risk-free rate minus house price growth (the so-called 'real world') approach. Debates about the advantages and disadvantages are about theoretical methodology; how we should convert facts to assumptions in a pricing model. The basic facts (historic house price growth, government bond yields, levels of rent) are not seriously in dispute.

The idea of a binary choice between risk neutral and real world, is a false dilemma. If you dislike the risk-neutral approach, many alternatives remain. In this section we outline a variety of theoretical arguments and explain their consequences for ERM valuation.

5.1.2 Broad Range of Theories

The way that empirical data turns into pricing assumptions relies on a valuation methodology. The number of available methodologies is large. The approaches might be divided into 'real world' and 'risk neutral' approaches, but these terms mean different things to different people in different contexts and represent only two points on a spectrum.

Finding the right framework is particularly difficult for ERMs because we have embedded options involving illiquid assets. This is problematic because our best theories for pricing options require dynamic hedging in frictionless markets, while our best theories for pricing illiquid asset relate to fixed income markets where (promised) cash flows are not stochastic. Somehow these assets must be combined, which requires new thinking.

The chart shows a possible decision tree for valuing any instrument. The first question is whether a market exists for that instrument, so the price can be observed directly. If the market price exists then, by definition, that is the fair value of the instrument. We might still want to use a model to come up with a view of fundamental value, that might then inform us if the market is cheap or dear. Some of those fundamental value models are of the same form that we would use for a fair value if we had no market value to inform us.

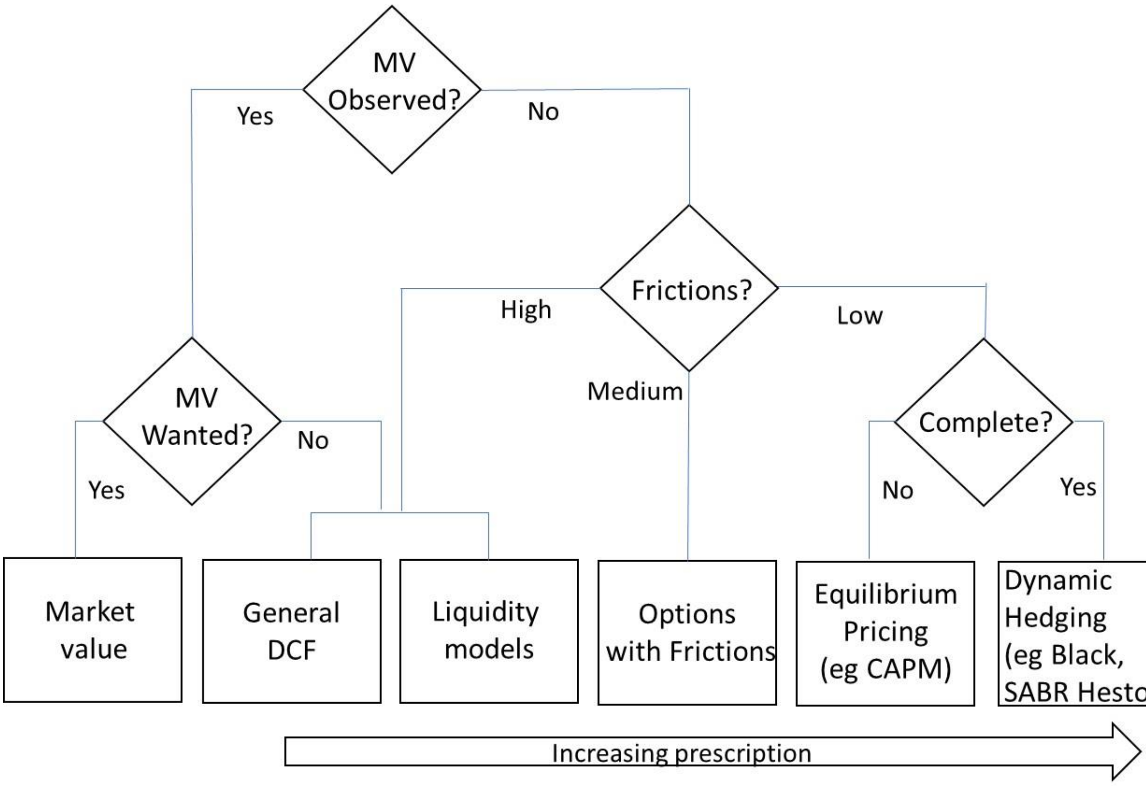

5.2 Describing Different Approaches

We now give brief descriptions of the standard approaches in this tree,

5.2.1 General Discounted Cash Flow

General discounted cash flow models involve forecasting cash flows and discounting them at an 'appropriate' rate. This is also sometimes called 'real world' pricing because of the claimed realism of the project cash flows. Often the cash flow assessment is deterministic, and the appropriate discount rate is a subjective assessment based on how risky those cash flows are thought to be. The discount rate should be consistent between structures with similar levels of risk. This approach, as a method for assessing whether market prices are cheap or dear, goes back at least to Graham and Dodd (1934).

Under a general discounted cash flow model, we might also consider a structure in the context of how it would be used and who would likely hold it. For example, although the cash flow of a put option is highly variable, that does not make a put option a risky investment, because few investors hold put options on a stand-alone basis. Held in connection with the underlying asset, the put option reduces the portfolio risk. As a result, in a general DCF calculation, we might use a discount rate for put options below the risk-free rate, or even negative. Indeed, option prices from models such as Black-Scholes can be re-expressed in terms of general DCF and it is instructive to see how the implied discount rates behave, as we do below.

Liquidity pricing is a special case of general DCF, where discount rates reflect the liquidity of an investment vehicle as well as the cash flows that the vehicle generates. This is an approach which has been widely used but does not appear to have been well documented, so we give it a special section below.

In markets, such as equity indices or foreign exchange, where option prices can be observed, there are two ways to estimate future volatility: by extrapolating past realised volatility or by using volatility implied by option prices. Although some financial models (including Black-Scholes) suggest that realised and implied volatilities should be the same, this is not found to be the case in practice, for a number of well-researched reasons including transaction costs, stochastic interest rates, risk premiums for uncertain future volatility, allowances for credit risk, funding costs and so on. In this context, models based on historical volatility would be called ‘real world’, while ‘risk neutral’ uses market implied. This is the usual terminology in actuarial work involving economic scenario generators, for example.

Another definition of ‘real world models’ is those which project cash flows on a realistic basis, with valuation allowing for risk via a judicious choice of discount rate, or via stochastic discount factors (also known as deflators). This approach is common for corporate valuation, where dividends or profits are forecast and discounted at an estimated shareholder required return. The Capital Asset Pricing model, or more advanced multi-factor models, may be used to assess the return required on risky cash flows. There is no conflict in principle between ‘real world’ approaches and option pricing theory; indeed early derivations of Black-Scholes used realistic projections and CAPM for discounting.

5.2.2 Frictionless Markets

One of the controversies surrounding ERMs is the applicability of the body of mathematics called *asset pricing theory*.

This theory underpins famous option pricing models such as Black-Scholes (1973), Heston (1993), SABR (Hagan et al, 2002) as well as models for the term structure of interest rates, credit spreads and many other financial quantities.

These models are derived based on mathematical assumptions, so-called *frictionless markets*, including:

- Every asset has a single price, for buying and selling. There are no transaction costs, and trades of any side have no impact on market prices.
- Assets can be held in any quantity, integer or fractional, positive or negative. No limits on borrowing at the risk-free rate, or short selling.
- Markets are open for trading in continuous time; transactions can be executed instantly.
- There are no institutional constraints: tax liabilities, margin or capital requirements.

These assumptions do not hold literally, even for equity and foreign exchange markets where the option pricing models are most widely used. The housing market fits these assumptions to a lower degree.

Within this framework, many models make specific distributional assumptions. Looking at equity option models, Black-Scholes assumes constant volatility, while Heston assumes volatility follows a square-root process and SABR volatility follows a lognormal process. There are option pricing models for series that are autocorrelated (as we will see house prices are), as these arise with so-called Asian options based on moving averages of share prices (Kemna & Vorst, 1990).

In the options literature, models such as SABR and Heston are especially useful when observed option prices at different strikes have different implied volatilities (the *skew effect*); the models are then calibrated to observed option prices to interpolate prices for other option-like structures. There has been little success to date in replicating those implied parameters from historic time series analysis. Similarly, in the ERM literature, it is well-known that time series properties of house market indices do not conform closely to the geometric Brownian motion underlying Black Scholes (see for example the ARMA-EGARCH models of Li et al (2010) and Tunaru (2019) which claim to be more realistic. These same papers, however, then produce ERM valuations that conform remarkably closely to Blacks & Scholes' formula. While we cannot rule out that, in future, property options might trade and skew effect become observable, at the time of writing different views on the level of volatility are so strongly divergent as to make detailed discussions of skew effects premature.

5.2.3 Option Implied Discount Rates

Under real-world models, a reference point for setting discount rates is the internal rate of return on trades whose prices are known and the distribution of whose future cash flows can be estimated. Equity options would be examples of this. Let us take a simple example, of a share with current value of 100, paying dividends at a (continuously compounded) rate of 3% per annum. Let us suppose that the risk-free rate is 2%, continuously compounded. We take a risk premium assumption of 8% per annum (arithmetic mean) which is optimistic for equities but is close to that implied by some of the more optimistic 'real world' house market models. We take an implied volatility of 12%, also consistent with values often proposed for the housing market; equity volatilities would usually be higher.

For a range of strikes, we can calculate the expected future values of 5-year options, both calls and puts. We can also estimate a market value of those options using Black-Scholes' formula. The 'real world' discount rate is the internal rate of return for those options.

We show some sample calculations below. "For a range of strikes we can calculate the expected future payoffs of -year options, both calls and puts, under 'real-world' GBM, including the risk premium. We can also estimate a market value of those options using Black-Scholes' formula, which is independent of the risk premium assumption. The 'real-world' discount rate is that necessary to equate the 'real-world' expected payoff with the Black-Scholes valuation." Discount rates are continuously compounded.

	Call Options			Put Options		
Strike	Expected	Market	Discount	Expected	Market	Discount
50	91.9070	40.8746	16.21%	0.0003	0.0456	-103.20%
80	62.0729	16.9841	25.92%	0.1661	3.3003	-59.78%
100	43.3396	7.3772	35.41%	1.4328	11.7901	-42.15%
120	27.5496	2.7574	46.03%	5.6429	25.2671	-29.98%
150	11.8601	0.5301	62.16%	19.9533	50.1849	-18.45%
200	2.1392	0.0280	86.72%	60.2324	94.9247	-9.10%
250	0.3153	0.0014	107.70%	108.4086	140.1400	-5.13%

Discount rates for call options are above the expected return on the shares, while discount rate for put options are below risk-free and typically negative. In financial theory, these results are in no way controversial, but people not familiar with option pricing theory might not anticipate that such extreme discount rates would arise.

Taking the case of the put options, how can we rationalise these negative discount rates? Why would an investor even consider an asset that is expected to lose money, let alone one as risky as a put option which has a chance of expiring worthless, losing everything? The answer is that few rational investors hold a portfolio 100% in a put option. Rather, a put option is a form of insurance held in connection with other assets. An investor in shares can, sometimes with a modest outlay, acquire a put option that substantially mitigates losses in a market crash. The willingness to accept a negative expected return on the put option reflects the reduction of risk to the portfolio as a whole. This is the same reason that buyers of household or motor insurance would not expect (or hope) to make a profit on that insurance.

The point can be made that option pricing theory uses concepts of dynamic hedging, which do not so easily apply to illiquid assets. So maybe options on illiquid assets would be priced differently? This turns out to be the case; option writers do charge in their pricing models for the higher costs of dealing in illiquid stocks, resulting in prices higher than implied by Black-Scholes. From the point of view of the investor in an illiquid asset, a put option (with physical delivery) provides protection not only against price falls but also against a lack of liquidity, because the option writer can be forced to pay cash in exchange for the asset. Option pricing allowing for illiquidity would logically use more extreme (ie more negative) discount rates for put options compared to methods in our table that assume perfect liquidity, because the option is mitigating not just market risk but also illiquidity risk.

This has important implications for the valuation of ERMs. The PRA's 4th principle from CP 48/16 was that "the compensation for risks retained by a firm as the result of the NNEG must comprise more than the best estimate cost of the NNEG. " The implied discount rates from options clearly reflect this principle. Industry practice for "real world" valuations is to use the same discount rate for ERMs regardless of NNEG moneyness. This might be justified if all the spread in the discount rate were due to illiquidity, but if some of it represents market risk from property prices than a more careful approach is needed.

5.2.4 Estimation of Volatility in a Frictionless Model

Under models of frictionless trading, option pricing follows from a replication argument. The relevant volatility is that faced by a dynamic hedger, who is rebalancing frequently. The annualised volatility over longer periods is irrelevant. Tunaru (2019) proposes a risk-neutralisation approach consistent with this argument; only the short-term volatility counts for option pricing.

Tunaru's approach does not, however, assume geometric Brownian motion. Instead, he fits a time series model in which annualised volatility increases with holding period, as, of course, the historic data shows.

Upward sloping volatility is consistent with positive autocorrelation in housing market returns. After a month of strong growth, the next month is more likely also to grow strongly, and the month after that. This implies an inefficient market; momentum traders could in theory get rich by waiting for the market to rise and then buying to profit from the continued rise. Conversely, if the market starts to fall, momentum traders could exit near the top, leaving someone else holding the house as the value subsequently tumbles.

An option replication strategy involves holding houses according to the delta of an option. An ERM writer wishing to hedge their exposure would need to take a small short position for high LTV ERMs, but would need to sell further as the market falls. When the market starts to rise, the delta moves closer to zero and the hedger starts to buy into the market again, unwinding some of their short position. These actions of an option hedger are similar to a momentum trader. Tunaru's low implied volatility (and hence low NNEG price) is taking credit for these momentum trading profits. This raises the question of whether these profits really exist, and if so, do they really reduce the cost of a NNEG? If we were dealing with equity options, the answer would be to use the short volatility, and yes, profits from trading the hedge would contribute to the option price – but of course equity returns are not strongly auto-correlated in the way that house indices are. Could it be that Tunaru is trying to have his cake and eat it – taking the frequent hedging argument from the equity analogy but then using the volatility slope from housing markets?

The index construction is important. In a sale of any asset, there is a gap between the striking of a bargain and the settlement when cash changes hands and ownership is legally transferred. For shares, that gap is typically one business day. For houses, at least in England and Wales, the gap may be up to one year. A house transaction typically starts with an indicative ask price, a potential buyer makes a non-binding bid, there can be a period of haggling while surveys are undertaken and financing is put in place. A forward contract is then agreed ("exchanging contracts") and the sale completes weeks or even months later. Equity indices report prices on the date of the bargain, while property indices report trades at the date of completion. As a result, estate agents can forecast house price index moves in the short term, not by time series analysis but by inspecting the deal pipeline before the transactions enter the index. It is not easy for arbitrageurs to profit from those index forecasts because, for example in a rising housing market, to trade at the current index price, the investor would need to get in a time machine and start the process six months earlier.

If we interpret that long settlement periods and other index construction/reporting effects act to boost reported short-term autorcorrelation and dampen short-term volatility, then it follows that short-term index volatility cannot be a justifiable basis for option pricing. Instead, we should look at the option pricing literature for options when the price of the underlying is artificially smoothed. Such options are known as *Asian options*, which have been studied for example by Kenma & Vorst (1990).

5.2.5 Option Pricing with Frictions

There is a literature on the impact of transaction costs on option prices. The standard references are Leland (1985), Hodges & Neuberger (1989) and Davis et al (1993). The approach in all of these papers is that transaction costs comprise part of the cost to the writer of providing an options service, while the option buyer is meeting some other need and is not trying to buy the option solely in order to hedge it away.

The option writer faces some difficult optimisation problems to establish when to run unhedged risk, versus when to incur transaction costs to reduce the risk. In contrast to the case of frictionless models, with transaction costs the asset risk premiums and utility function of the writer must also be estimated. The conclusion is that transaction costs are borne by the bearer, in the form of a larger option premium. Early papers, including those we list above, consider transaction costs as being a proportional bid-offer spread; more recent work has also looked at models where hedging trades can themselves impact market prices, typically in a direction that is costly for the hedger.

Black & Scholes can be criticised on many grounds, including the lack of an allowance for transaction costs. However, it would be wrong to suggest that transaction costs invalidate Black-Scholes, to the extent that general DCF is the only fall-back position. Indeed, while (depending on the assumptions), general DCF may give higher or lower option prices than Black-Scholes, more specific attempts to incorporate transaction costs invariably increases option values because those costs are ultimately borne by the option buyer.

5.3 Pricing Theory for Illiquid Assets

Liquidity pricing theory is a relatively new discipline which focuses on investors' attitude to illiquidity and how this affects asset prices. In a market with significant frictions, many investors may be able to choose between different instruments whose cash flows are similar, but some are more liquid than others. In this situation, investors naturally prefer the more liquid assets; this demand causes more liquid assets to trade at a higher price than otherwise similar illiquid assets.

Pricing theories based on liquid markets, by definition cannot explain the pricing differences between liquid and illiquid assets.

5.3.1 Bond Illiquidity Premiums

Illiquidity premiums are most visible in fixed income markets, as they emerge as a component of bond yields.

There is an emerging literature on the decomposition of bond spreads. The yield on a bond might be expressed (Webber, 2007; van Loon et al, 2015) as a sum of the following parts:

- A liquid risk-free rate
- Expected defaults, that is annual probability of default times expected loss given default
- A risk premium to compensate investors for uncertainty in defaults (sometimes called unexpected defaults)
- An illiquidity premium, that is, the difference between the observed bond yield and the yield on a hypothetical perfectly liquid bond with the same credit risk.

The illiquidity premium is there to compensate investors for illiquidity costs, that is, the contingent costs of having to liquidate an illiquid asset.

In the context of liquid market theories, thus raises a puzzle. How could we have two bonds with the same cash flows, but different prices? Does this not create an arbitrage for traders issuing the liquid bonds and collecting the illiquidity premium on the illiquid bonds?

To see why this arbitrage does not work, consider an illiquid but default-free bond. Suppose a trader borrows at the liquid risk-free rate to invest in this bond. The liquid risk-free borrowing must be liquid from the perspective of the lender, who can call in their loan at any time. If the lender calls in the loan early, the borrower must sell their asset into an illiquid market. The achievable price in a forced sale may be less than amount borrowed, so the lender now faces default losses. Lenders, anticipating the possibility of such losses, would not lend at the risk-free rate, and so the whole arbitrage unravels.

5.3.2 Illiquidity Costs

Against the illiquidity premium should be considered costs of illiquidity. For insurers, there is a long list of potential triggers that can cause illiquidity shocks, potentially forcing asset sales. In the face of an illiquidity shock, insurers naturally sell first their most liquid assets, but for a sufficiently large shock, illiquid assets might also have to be turned into cash. Smith & Spivak (2012) give a list of illiquidity triggers; some (but not all) of these could also apply to annuity writers.

Illiquidity costs, however, have some particular features that are different in nature to default losses, and these differences merit a different treatment in pricing theory. If a bond defaults, liquidators are appointed and bondholders recover a certain proportion of the bond value. All investors recover the same proportion, and if you hold twice as many bonds you recover twice as much cash. Illiquidity costs do not behave like this. Illiquidity shocks are specific to a financial institution, according to how often the capital providers demand cash returns. Furthermore, illiquidity costs are convex in the proportion of illiquid assets held in a portfolio. If an institution sells liquid assets to buy more illiquid assets, they reduce their capacity to absorb illiquidity shocks, and illiquid assets will more often have to be realised. This is in contrast to the illiquidity premium, which being an inherent property of the bond, applies equally to different investors and scales linearly by asset proportions.

To give an example calculation, let us suppose that a firm's expected illiquidity costs per annum take the form ch^p, where c is a constant depending on the institution, h is the proportion of illiquid assets in a portfolio and p is an exponent; to reflect the convexity of illiquidity costs, we set $p>1$.

The *average* illiquidity costs per illiquid asset is ch^{p-1}, that is, the illiquidity cost divided by the quantity of illiquid assets held. The *marginal* illiquidity cost is obtained by looking at small changes in h, giving (on differentiation) cph^{p-1}. As $p > 1$, the marginal cost exceeds the average cost. This is just a simple example; firms' illiquidity cost curves are difficult to calibrate (as they require a model of liquidity shocks and their sizes). A power law may not be the right function; our argument requires only that the function is convex.

In a portfolio optimisation, a rational investor increases their holding h in illiquid assets until the marginal illiquidity cost is equal to the illiquidity premium. This means that institutions with a larger value of c will hold less in illiquid assets; insurers with long-term liabilities would have a lower c and so optimally hold illiquid assets. What is in common to all investors, however, is that their average illiquidity costs are a proportion $1/p$ of the illiquidity premium (because the illiquidity premium is their marginal illiquidity cost). Investors then lose a proportion $1/p$ of the illiquidity premium in illiquidity costs, but get to retain the remaining portion $(p-1)/p$ of the illiquidity premium. As their low value of c is attributable, at least in part, to the stable funding arising from long-term liabilities, we might take the retained portion of the illiquidity premium onto the liability discount rate, resulting liabilities discounted at a rate higher than (liquid) risk-free.

5.3.3 Implications for Liability Discounting

Given that different bonds have different yields, how do we pick an appropriate yield for valuing insurance liabilities? The conventional, liquid market, point of view is that risk premiums in the bond are rewards for risks (chiefly default) risks borne by bond holders. If you hold the bond, then you bear the risk and (expect to) earn the risk premium over the time. This line of reasoning excludes double-counting any of the bond risk premium, both earning it over time and also up-fronting the premium as an uplift to liability discount rates.

The argument is more subtle if we accept that different bonds could have different yields, because (for example) of liquidity, even if the credit risks are the same. A theory of liquid asset pricing cannot help us decide between yields representing different liquidity levels, because according to that theory, the differences should not exist in the first place.

We can regard insurance policies as a form of borrowing on the part of the insurer. Where policyholders can ask for their money back, for example in the form of a surrender value, then this borrowing is liquid (from the policyholder perspective). However, some long-term policies, particularly annuities, cannot be surrendered or can be surrendered only on terms unfavourable to the policyholder. The insurance industry has argued that as long-term liabilities are predictable and stable and thus "illiquid", their discount rate should also contain a market liquidity risk premium. (Perotti et al, 2011).

These arguments have been reflected in Solvency II's matching adjustment. Roughly, the *fundamental spread* corresponds to the expected default and the default risk premium (the latter being estimated by a theoretically unconventional 'cost of downgrade' calculation); the yield spread minus the fundamental spread, that is the liquid risk-free rate plus the illiquidity premium, is then used for discounting liabilities. There is a long list of eligibility criteria (Article 77b - Directive 2009/138/EC) designed to ensure that the liabilities are sufficiently illiquid and that the assets match the liabilities so that the insurer is unlikely to need to liquidate those assets in order to meet policy liability outflows. They might still need to liquidate assets for other reasons, such as keeping within default risk appetite following bond downgrades, and others discussed in Smith & Spivak (2012).

For the avoidance of doubt, we are not holding up the MA calculation as a shining example of fair value; we suggest that in future, fair value accounting arguments may use some of the illiquidity concepts developed for Solvency II.

5.3.4 Information Asymmetries

The rationale we have stated here for illiquidity premiums is not universally accepted within the academic community. An academic panel retained by EIOPA to advise on regulatory use of illiquidity premiums concluded in a seldom-quoted report (Perrotti, 2011):

"In our opinion, this argument [for illiquidity premiums in liability discounting] has no sound scientific basis. … Using a liquidity premium to discount liabilities is in essence a fudge discount rate that is financially unsound and economically indefensible."

The fundamental thought experiment underlying the illiquidity premium is of a liquid bond with cash flows equivalent to an illiquid bond. This treats illiquidity as an incidental feature of a bond; for some unexplained reason, dealers charge larger spreads on some bonds that on others. Academic studies such as Glosten & Milgrom (1985) explain illiquidity in terms of more fundamental concepts involving information asymmetries. Dealers widen spreads not only because they are greedy for margins, but also to avoid adverse selection when other parties are better informed. Low risk bonds, such as government bonds, have the least scope for information asymmetries and also tend to have the tightest dealing spreads. It is debatable whether the concept of two bonds identical, aside from liquidity, makes economic sense.

Despite numerous attempts, it has proved difficult empirically to disentangle illiquidity from credit risk, with different methods producing different answers. However, in practical terms, the consensus that a split exists in principle is a step forward. It gives different parties: regulators, industry, government, consumer groups; a single number to haggle over so that at least we have a common language to articulate what is inevitably a political settlement.

5.3.5 Illiquidity Premiums in Housing

While illiquidity premiums in bonds have been widely studied, in principle there could be illiquidity effects in other asset classes too.

Let us consider residential housing. This is plainly an illiquid asset class; it is difficult, costly and time-consuming to buy or sell residential properties. Just as with illiquid bonds, we can try to tease out the concept of an illiquidity premium by comparing a house investment (let's say a deferred interest, to avoid complications of accounting rents) to a theoretically perfectly liquid investment linked to the price of the same house on the same deferment date. We know how to price the illiquid investment – using a deferment rate of index rental yield plus dilapidation.

Just as with bonds, we should expect the perfectly liquid structure to trade at a higher price than the less liquid structure, or, equivalently, the direct house investment should carry an illiquidity premium in terms of its return.

Just as with bonds, we have to check that the concept of market illiquidity pricing does not create arbitrage opportunities. The theoretical arbitrage would be to buy physical property, collecting the illiquidity premium, financed by liquid (and default-free) housing linked bonds. The arbitrage does not work because the illiquidity of the house could prevent the owners of the liquid assets cashing in their investment; the default-free bond is not default-free after all.

This does not mean that financial firms can never profit from making property investments, and issuing the market risk in more liquid form, such as unit-linked policies or unit trusts. However, in order to avoid default risks, such funds must have the right to pass illiquidity risks to the investor. Several UK property funds exercised this right in the wake of the Brexit referendum result (Goodley & Treanor, 2016).

There are vehicles where property illiquidity risk is genuinely reduced, by turning the property into tradable securities such as investment trusts or exchange traded funds. It is well-attested that these instruments do not always trade exactly at the value of the underlying assets, with differential liquidity a possible component of the difference (Adams, 2000), alongside management expenses and asset selection skill. In the same way, social housing bonds (which do exist) probably do reflect liquidity premiums in the underlying houses, although this is likely masked by a long list of other effects.

Putting together the illiquid and liquid case, we can regard the returns in physical property investment, capital plus rental income, of being composed of four elements:

- A liquid risk-free return
- Plus risk premium (to compensate for market risk)
- Plus illiquidity premium (to compensate for illiquidity costs)
- Plus or minus random noise (the market risk)

We could argue that, as with option pricing theory, the market risk premium should cancel out in the projection and discounting process for producing a current value. However, maybe as for corporate bonds, part of the illiquidity premium does not self-cancel, but instead justifies an uplift to a price growth assumption. To put this differently, the relevant house price for putting into a structure such as an ERM under a risk-neutral model should not be the current house price (which is reduced by illiquidity) but instead the (higher) hypothetical price of a fully liquid house-linked bond.

Some possible assumptions are shown, for bonds and for direct housing investment, in the chart below. In this chart HTM (hold to maturity) denotes a return before illiquidity costs, while the net return is after illiquidity costs. The average illiquidity costs were estimated using a power law with p=4.

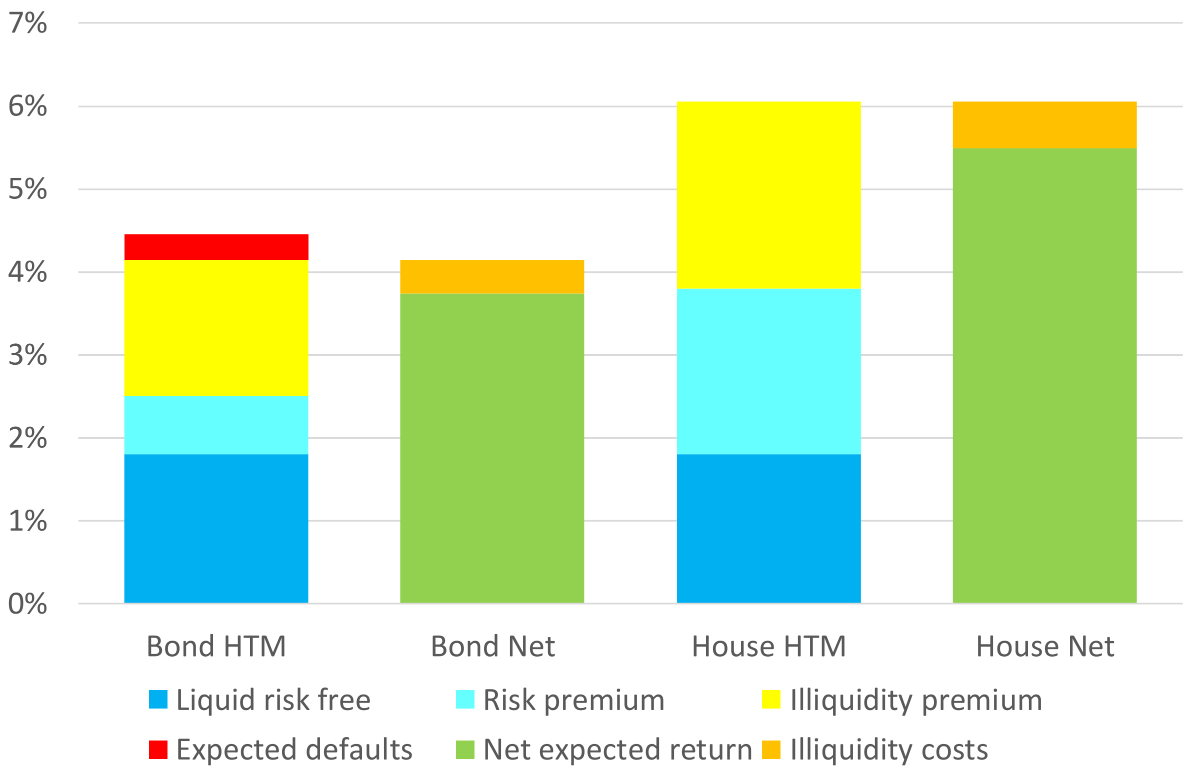

We are not aware of published studies quantifying illiquidity premiums on residential property. However, much more has been written about commercial property liquidity premiums. There are several measures of volatility; some based on sector measures such as volumes traded, or size of inflows / outflows from the market. Others, such as the Time on Market measure, require data about the sale process, not just the final transaction. Studies such as Marcato (2015) examine the impacts of these different liquidity measures on property returns.

5.3.6 Illiquidity Premiums on Housing-Backed Bonds

We have discussed how yields on illiquid bonds can be elevated (or prices depressed) relative to liquid bonds, because illiquidity imposes real costs on (some) bond investors. We have also argued that the house prices are lower than they otherwise would be, because of the illiquidity of the housing market. This raises the question of whether we would be double-counting illiquidity premiums, if we added an illiquidity premium to an ERM discount rate and at the same time used as input a house price whose value was depressed due to an illiquidity premium.

Equity release mortgages and other debt structures secured on houses, are not perfectly liquid. They have their own illiquidity premiums. It is the nature of illiquidity premiums that they depend on the vehicle, and not on the underlying risks. Indeed, it is almost the definition of illiquidity premiums; while risk premiums pass on through restructuring, the illiquidity premium is the portion of the return that restructuring can change.

For the final step in the deferment rate argument, we therefore have to add the illiquidity premium for the bond.

To summarise the arguments so far for the deferment rate, we have proceeded as follows:

- Start with the equivalent rental yield in an index
- Add an allowance for expected dilapidation relative to the index
- Subtract the illiquidity premium for physical investment in houses
- Add back the illiquidity premium for the bond structure.

The full set of assumptions, allowing for illiquidity premiums, is tabulated below.

ERM Parameter	Construction
Discount *r*	Liquid risk-free rate + bond illiquidity premium
Deferment *q*	Index rental yield + dilapidation – house illiquidity premium + bond illiquidity premium
Growth *r-q*	Liquid-risk free rate – index rental yield – dilapidation + house illiquidity premium
	House expected total return – index rental yield – dilapidation - market risk premium

In terms of the impact on ERM values, we see that the house illiquidity premium acts as negative dilapidation. The two (equivalent) expressions for the growth rate *r-q* in an ERM are obtained by subtracting *q* from *r*. We can illustrate the growth assumption in the bar chart:

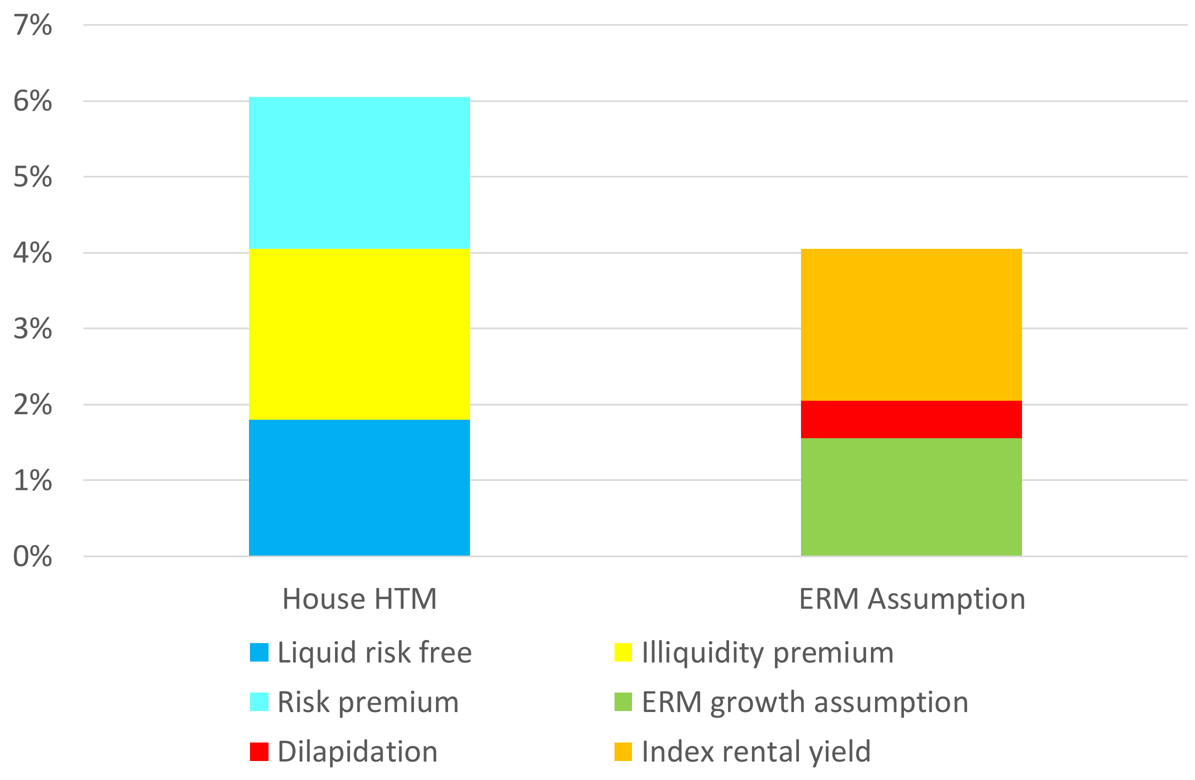

5.4 Applying the Zero Profit at Inception

While bond yields, house price growth and house price volatility might be important for pricing ERMs, we can directly observe the prices of ERMs at the point of lending, where the lender advances cash and, in return, receives an ERM.

There is unfortunately no observable secondary market in ERMs, so insurers writing ERMs will typically observe the price of new lending and use this to calibrate a model which is then also used to re-value older lending. That means that if, for example, increased competition leads to reduced roll-up rates on new lending, other things being equal, a larger value will be attributed to older cohorts of ERMs. It also means that new ERM lending shows zero profit on inception.

The model matters; given that the rollup-rate for ERMs typically exceeds recent UK house price growth, older cohorts of ERMs typically have higher LTVs than current lending. The impact of the model assumptions is the way they differentiate between different LTV levels. Lenders with a large book of historic ERMs are much more sensitive to model choices, compared to newer entrants to the ERM market whose back book more closely resembles current lending.

5.4.1 Alternative Ways to Mark to New Lending

The marking of back books to new contracts is widespread in accounting, not just for ERMs. For example, in option markets, it is usual to observe the prices of option transactions and work backwards to the *implied volatility*. This implied volatility then calibrates models used for re-valuing other positions that have not recently traded. The marking of models to recent transactions is a property of international financial reporting standards (IFRS 9 for financial instruments, which in turn refers to IFRS 13). See Kenny et al (2018) for a fuller discussion of the application of accounting standards in ERM valuation.

The implied volatility of an equity option is a relatively easy calculation, when the risk-free rate and the dividend yield (or forward price) are readily observable. There is one equation to be solved, in one unknown. The option price is an increasing function of the volatility, and the range of the option price covers all arbitrage-free option values, so there is a unique implied volatility. We exploited this property in our initial discussion of ERM parameters.

For ERM, apart from the non-financial assumptions (mortality, early redemptions etc) there are three financial assumptions, all of which may be disputed: the discount rate, the deferment rate and the volatility. Unlike the situation for equity options, an implied volatility approach is seldom used for ERMs. Instead, firms estimate the volatility from historic data, and one of the other parameters is solved from the new lending rate.

As well as solving for *q* given *r*, and vice versa, there is a third alternative which is to specify the implied growth rate *r*-*q* and then solve for the discount rate *r*. This, the most popular method in practice, has the computational advantage of requiring only an internal rate of return calculations, as the feasible parameter choices do not change the moneyness of the NNEGs.

The chart below shows the alternative ways of marking to new lending. In this chart, we have assumed that new lending is at a 30% LTV with 5% rollup rate and 13% volatility. The dotted line shows the feasible combinations of deferment rate and discount rate which price a new ERM at the initial loan balance (ignoring acquisition costs and other expenses).

Maybe by some independent theoretical process, the actuary has derived a preferred deferment rate of 2.8% and deferment rate of 1.25%. This clearly does not price the initial lending at par; it would give rise to a substantial profit at inception. To remove the profit on inception, we can do any of the following:

- Project onto the dotted curve at 45 degrees, taking the implied house growth of 1.55% as fixed.
- Project horizontally onto the dotted curve, taking the discount rate of 2.8% as fixed and disregarding the theoretical deferment rate.
- Project vertically onto the dotted curve, taking the deferment rate of 1.25% as fixed and disregarding the discount rate.

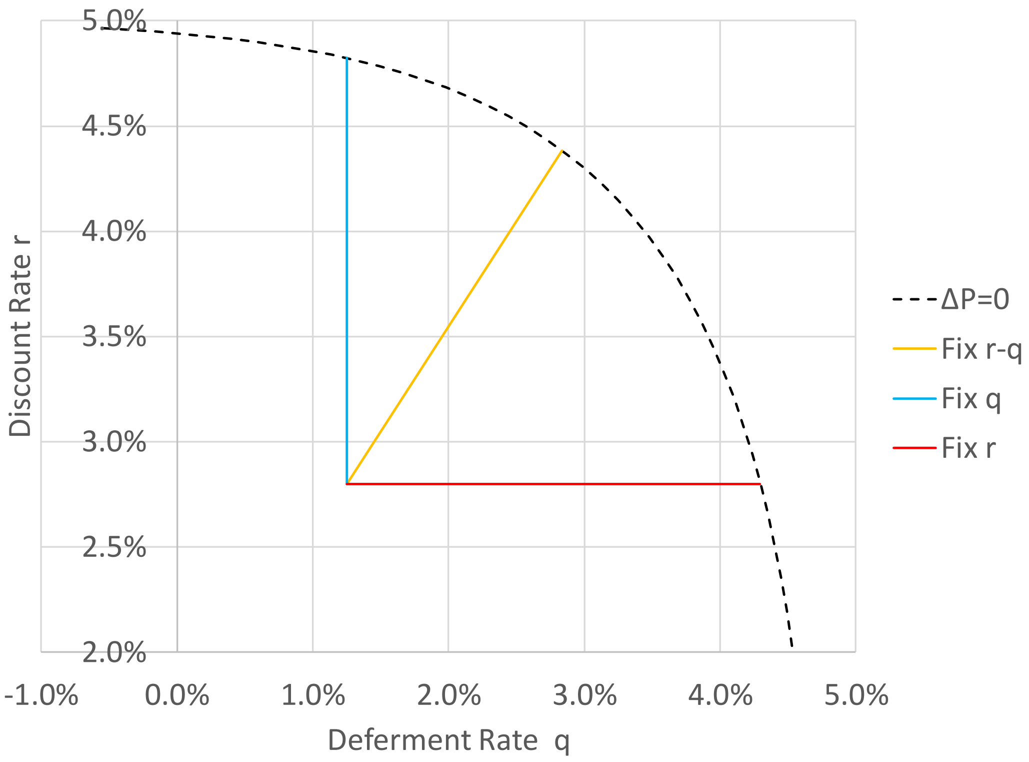

This chart also explains some of the industry nervousness about PRA's proposed minimal deferral rates. Only with the blue method does the insurer choose the deferral rate. In the other cases, *q* is a result of marking the model to new lending. If new lending spreads were to tighten considerably, a situation could arise where firms' estimated *q* falls below any given threshold, not as a deliberate choice but as a result of market moves and a consistent methodology.

5.4.2 Illiquidity Premium Interpretation

Marking to new lending by changing model parameters could be misconstrued as estimating parameters that we than have to admit were wrong. This is not usually how the calculations are articulated. Instead, the difference between the theoretical value of new lending is supposedly explained by an illiquidity premium. This is usually in the context where bonds are assumed to have illiquidity premia but houses do not.

If we move to a model recognising illiquidity premia on both bonds and houses, then the marking of new business is more problematic; we have one equation but two unknowns. We have previously expressed the mark-to-model choices in terms of the parameters to be solved, but we can instead express these in terms of illiquidity premia, in the table below.

Method	Parameters to Fix	Illiquidity premiums to fix
IRR for fixed cash flows	Fix *r-q*, solve for *r* (and hence *q*)	Fix house illiquidity premium; solve for bond illiquidity premium
Solve for deferment	Fix *q*, solve for *r*	Fix the difference between house and bond illiquidity premia, solve for house illiquidity premium (and

		hence deduce bond illiquidity premium, or vice versa.)
Solve for discount	Fix *r*, solve for *q*	Fix bond illiquidity premium, solve for house illiquidity premium.

5.4.3 Effect of Profit Zeroisation on ERM Values

We have already remarked that, because of the zero-profit constraint on inception, model choice has no impact on new lending, and little impact on past lending if they have similar terms and LTV to new lending.

We can ask what happens to the value of past lending under the different modelling options. These are shown in the chart below. For comparison, we also show the implied volatility calculation:

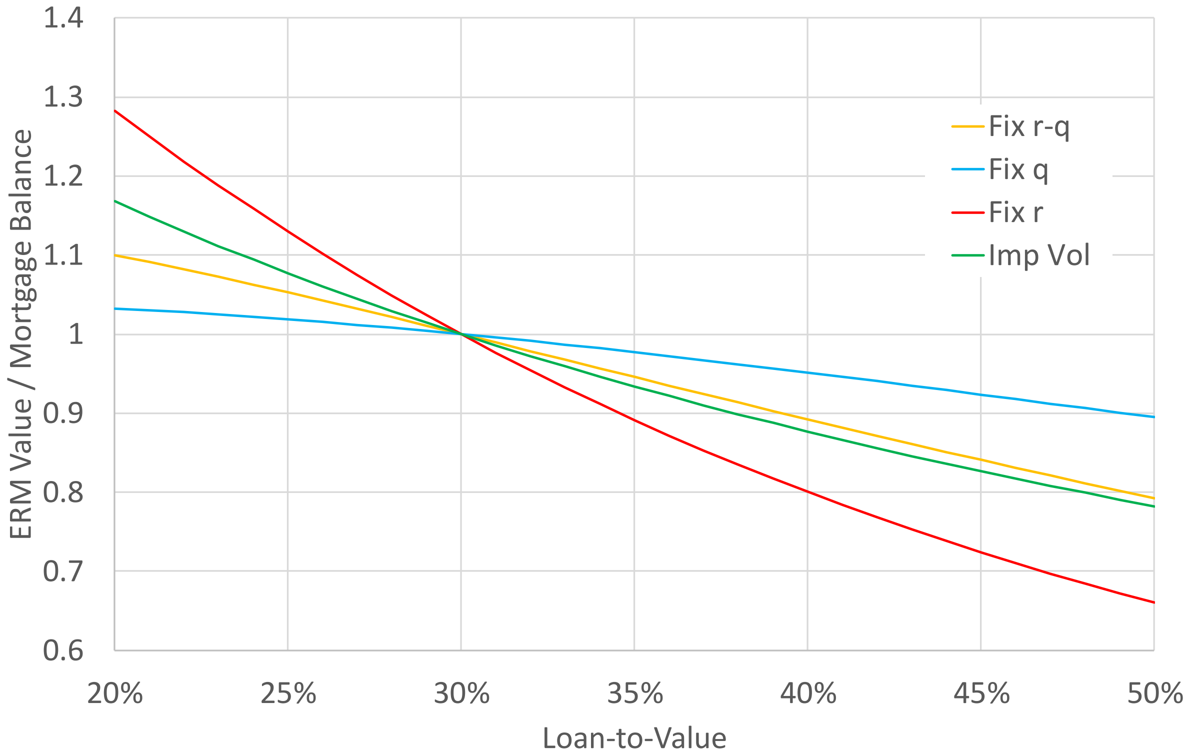

These results (apart from the implied vol which is not obvious) are easy to explain. The higher LTV mortgages are more sensitive to the deferment rate *q*, with higher deferment rates leading to lower ERM values. We get the highest deferment rate projecting horizontally, the lowest deferment rate projecting vertically while the 45-degree projection produces values in the middle.

5.4.4 Zero Profit on Inception and the EVT

The accounting requirement for zero profit on inception is usually interpreted that the value of new ERM lending should be the amount advanced (perhaps minus some initial expenses). However, this is not strictly a zero-profit test; it is a zero-change-in-assets test.

As the PRA has pointed out with its economic value test, the inception of an ERM can also give rise to a reduction in liabilities via the matching adjustment. The valuation of liabilities under IFRS is still being interpreted, but it is possible, even likely, that under the so-called ‘top-down’ approach to discount rate construction, ERM lending could likewise reduce stated liabilities under IFRS.

You can see what is coming next. Suppose that the zero-profit criterion were to be re-interpreted more literally, to include changes in assets and liabilities. If that were set to zero, then ERMs would look much less attractive from an accounting point of view. A small change in interpretation of accounting standards could have wide implications for ERMs,.

6 Portfolios of ERMs

6.1 Estimating the Initial House Value

We have developed option pricing formulas under the assumption the initial house value is known. However, we generally do not know the house value at the time of valuing the ERM. We know the house value earlier, and we might update this in line with an index.

6.1.1 ERM Value with Known House Price

To develop this further, we have to separate the date on which the house was last valued, from the date of valuing the ERM-let. Specifically, suppose the house was last valued at time 0, while the valuation date is t and the ERM-let expiry date.

If the house price H_t is known at time t, then the Black-Formula for the ERM gives:

$$ERM_{let} = \left\{ e^{-rT} K \Phi\left(\frac{\ln(H_t/K) + (r - q - \sigma_H^2/2)(T-t)}{\sigma_H\sqrt{T-t}}\right) + e^{-qT} H_t \Phi\left(\frac{\ln(K/H_t) \ + (q - r - \sigma_H^2/2)(T-t)}{\sigma_H\sqrt{T-t}}\right)\right\}$$

6.1.2 Dilapidation Uncertainty

If the house price is not observed, though, the best we can say is that it probably depends on the following:

- The initial house price H_0
- The change I_t/I_0 in the index since the house was last valued
- Expected dilapidation μ_d in the house relative to the index
- Volatility σ_d in the dilapidation relative to the index.

Putting these together, we could use the distribution:

$$\ln H_t = \ln\left(\frac{I_t}{I_0} H_0\right) - \mu_d t + \mathcal{N}\left(-\frac{\sigma_D^2 t}{2}, \sigma_D^2 t\right)$$

Now we need to combine the ERM-let price with this distribution of H_t.

We might be tempted just to take the expected house price $I_t H_0/I_0$ and plug it into the valuation formula. However, that would be

The necessary integral is tedious to evaluate, but the answer is straightforward. The expected ERM value over the distribution of H_t is:

$$\mathbb{E}(ERM_{let}|I_t) = \left\{e^{-r(T-t)}K\Phi\left(\frac{\ln(I_tH_0/(KI_0)) + (r-q-\sigma_H^2/2)(T-t) - \mu_D t - \sigma_D^2 t/2}{\sqrt{\sigma_H^2(T-t)+\sigma_D^2 t}}\right) + e^{-\mu_D t - q(T-t)}\frac{I_t}{I_0}H_0\Phi\left(\frac{\ln(KI_0/(I_tH_0)) + (q-r-\sigma_H^2/2)(T-t) + \mu_D t - \sigma_D^2 t/2}{\sqrt{\sigma_H^2(T-t)+\sigma_D^2 t}}\right)\right\}$$

6.1.3 Revaluation Effects

Although we might model a house price assuming (as a best estimate) that it follows an index minus dilapidation, there are times when we get to find out the true value of the house. At least this happens when the borrower dies. It may happen earlier as, for example, a valuer inspecting the house of a new ERM applicant may be asked to pass by other mortgaged properties that are in the area. At this point the distribution of dilapidation collapses from a lognormal distribution to a single point. The ERM value also jumps from its conditional expectation to its value given the no-observed house price. The sizes of these jumps have zero mean (because the value before the jump is the conditional expectation of the value after the jump).

6.2 Portfolios of ERMs

6.2.1 The value of an ERM portfolio

Let us suppose now we have a portfolio of ERMs on a large number of houses, each of which has dilapidation relative to an index I_t with mean μ_D and volatility σ_D. All houses are based off the same index but the dilapidations, we will assume, are independent.

Between house revaluations, the conditional expected ERM values (given the index) are functions of the index. House revaluations give rise to independent zero-mean jumps, but for a diversified portfolio these jumps (by the central limit theorem) are small compared to the impact of index moves. So to all intents and purposes, the sum of conditional expectations *is* the sum of the ERM values of the form:

$$ERM_{let} = \left\{e^{-r(T-t)}K\Phi\left(\frac{\ln(I_tH_0/(KI_0)) + (r-q-\sigma_H^2/2)(T-t) - \mu_D t - \sigma_D^2 t/2}{\sqrt{\sigma_H^2(T-t)+\sigma_D^2 t}}\right) + e^{-\mu_D t - q(T-t)}\frac{I_t}{I_0}H_0\Phi\left(\frac{\ln(KI_0/(I_tH_0)) + (q-r-\sigma_H^2/2)(T-t) + \mu_D t - \sigma_D^2 t/2}{\sqrt{\sigma_H^2(T-t)+\sigma_D^2 t}}\right)\right\}$$

6.2.2 Rewriting as an option on an index

We have argued that an individual ERM should be valued as an option on an individual house, with volatility $\sigma_H = \sqrt{\sigma_I^2 + \sigma_D^2}$. We now show how this can also be viewed as a different sort of option on the index, not by any profound economic argument but just by rearranging the algebra.

We start by manipulating the variance terms:

$$\sigma_H^2(T-t) + \sigma_D^2 t = \sigma_I^2\left(\frac{\sigma_H^2}{\sigma_I^2}T - t\right) = \sigma_I^2(\tilde{T} - t)$$

Here, the *modified strike date* is given by an extension formula:

$$\tilde{T} = \frac{\sigma_H^2}{\sigma_I^2} T$$

Likewise, we define the modified deferment rate as:

$$\tilde{q} = q - \mu_D$$

Recall that in our construction of q, we started with the index yield, added the dilapidation rate, subtracted the house illiquidity premium and added the bond illiquidity premium. Therefore, subtracting μ_D simply cancels the dilapidation we added in the first place, so the modified deferment rate is the index yield plus the bond illiquidity premium minus the house illiquidity premium.

We also need to define a *quantity*, that is the effective number of euros per index point, by

$$\tilde{Q} = e^{(q-\mu_D)\tilde{T}-qT} \frac{H_0}{I_0}$$

The H_0/I_0 term should be obvious – this is the number of units in the index represented by the initial house price. That ratio then has to be adjusted for expected dilapidation and extension to the modified strike date.

Finally, we define the modified strike price by:

$$\tilde{K} = \frac{e^{r(\tilde{T}-T)} K}{\tilde{Q}}$$

This is a rescaling so that the strike is expressed in terms of index points rather than euros, together with a compound interest adjustment to the new term.

The ERMlet now takes the simpler Black-Scholes form (for all $0 \leq t < T$):

$$ERM_{let} = \tilde{Q} \left\{ e^{-r(\tilde{T}-t)} \tilde{K} \Phi \left(\frac{\ln(I_t/\tilde{K}) + (r - \tilde{q} - \sigma_I^2/2)(\tilde{T}-t)}{\sigma_I \sqrt{\tilde{T}-t}} \right) + e^{-\tilde{q}(\tilde{T}-t)} I_t \Phi \left(\frac{\ln(\tilde{K}/I_t) + (\tilde{q} - r - \sigma_I^2/2)(\tilde{T}-t)}{\sigma_I \sqrt{\tilde{T}-t}} \right) \right\}$$

On substitution, it can be seen that this combination of modified parameters gives the same ERM price as the expectation formula. We have re-expressed the ERM value as an option on the index I_t with volatility σ_I rather than on an individual house with volatility σ_H. The difference in volatilities is compensated by artificially lengthening the term of the option, in proportion to the squared volatility ratio, and making corresponding modifications to the strike.

Therefore, to those who argue that the relevant volatility for ERMs is the index volatility rather than the individual house volatility, we say: you have a point, the ERM value can be expressed in that way, but you should then extend the strike date, and to do so correctly still requires an assessment of the volatility of individual houses.

6.2.3 Replication with rental portfolios.

An individual ERM contains an embedded option on an individual house. For options on equities or currencies, dynamic hedging is an in important concept in pricing. This relies on shares and currencies being available in multiple identical units, so that long and short positions can combine to make hedges. Such offsetting trades cannot work for houses because no two houses are identical.

And, of course, it is not possible to trade fractional participations in an ERM'd house under the nose of its owner.

However, we can compare the return on an ERM portfolio with that on direct investment in a residential property portfolio. Both portfolios depend on an index, and there can be multiple ways to construct portfolios to track a given index, provided the mix of geographic location, construction type and so on are controlled appropriately.

Comparing ERMs to direct property investments, according our valuation formulas, both earn the market risk premium in proportion to their house market price exposure. However, while the direct property investment earns a property illiquidity premium on top of the property risk premium, an ERM earns a property risk premium plus a bond illiquidity premium. The difference in illiquidity premiums compensates for the difference in liquidity. This relation has to hold in order to answer the question – if insurers can access high property returns via ERMs, why do they not just invest in property?

There is another reason why insurers do not generally invest in residential property, at least in the UK and Ireland. The market rent on a property index would include allowances for voids and arrears experienced by a typical landlord. Insurers, arguably are not typical landlords because their branding emphasises looking after the vulnerable: widows, orphans and so on, by paying insurance claims in their hour of need. Insurers are only too aware that the business of letting residential properties sometimes involves the nasty business of evicting tenants in arrears, who may often turn out to be – guess what – widows and orphans. If an insurer were to try to enter the private rental market, public relations considerations may dictate more generous treatment of unprofitable tenants, so the insurer would not be able to earn the same market rent as other landlords (Ashurst et al, 2008).

An attraction of ERMs from an insurer's perspective is the ability to gain exposure to the housing market and earn a market rent (as this is baked into the ERM price) without the (financially and reputationally) costly business of evicting living tenants.

6.2.4 Tranche Securitisation

Some firms writing ERMs have securitised those portfolios externally, separating them into senior and junior tranches which are sold to external investors. More recently, firms have applied internal restructuring, again separating into senior and junior notes but, now, these notes are all held by the insurer, albeit often in different funds.

Modelling of securitisations can be complicated, because each ERM depends on a single house. It may seem that all these houses have to be modelled individually within a huge multivariate structure. However, we know that diversified ERM portfolios can behave like options on indices.

If the junior and senior notes were split on a fixed date, we could decompose the value of an ERM portfolio into junior and senior portions using the formula for valuing options on options (Geske, 1979). Where, as is more often the case, there are multiple opportunities to fund senior note payments over time, the modelling of an ERM portfolio as a function of a single index still offers a substantial computational simplification even if closed form solutions no longer apply.

7 Stress Testing ERMS

While the sections above deal with the market consistent value or best estimate, consideration also needs to be given to the capital required or SCR.

7.1 Basic Stresses

For the purpose of basic stresses, we believe that the one that needs special treatment is that for residential property prices. This dominates other risk drivers, such as actuarial decrements which can be treated in ways familiar from other insurance businesses.

We would like to make a number of assertions.

Firstly whatever model is used to derive the best estimate, it is not a given that the same model will be suitable for deriving a 99.5 percentile (or that it will not be suitable). Separate justification is required as extremal values are unlikely to be accurately derived by looking at data from periods when all is well (even if you have lots of that data). The lessons of the financial crisis surely mean that we need not justify this statement.

Secondly that the nature of residential property crashes means that there is very little trading on the downward slope of a crisis. The shocked property owners will be very reluctant to sell, the delays in settlement and processing of house sales mean that the depth of the fall is not seen for a little while. The graph below shows that it took five years for the bottom to be reached in the measured house price index but actually the stresses were there to be seen already in 2008. Any company which had taken possession of a house in 2008 would have been extremely lucky to get out at the index as shown by that graph.

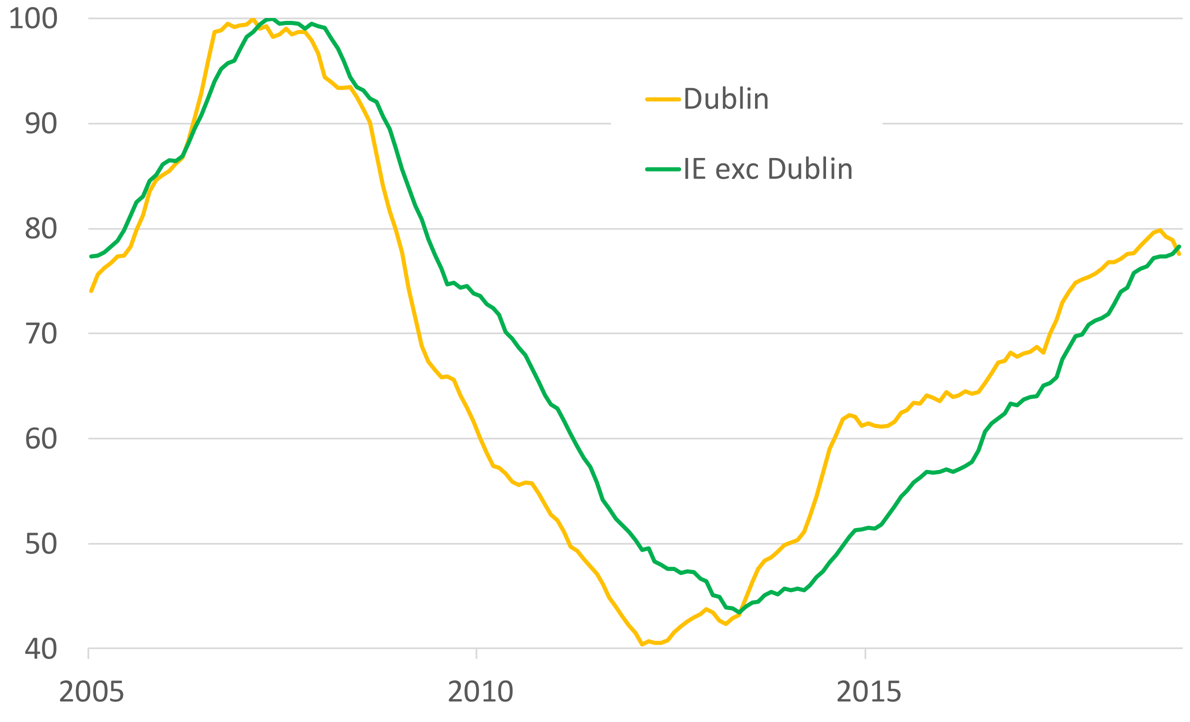

The series shown are Dublin-all residential properties, and National excluding Dublin – all residential properties, downloaded from the CSO. We have scaled them so that the maximum pre-crisis was 100.

We also believe that for ERMs there is a valid case for arguing that they “deviate(s) significantly from the assumptions underlying the standard formula calculation” as Article 110 of the Solvency II directive prescribe. We believe that were ERMs to become a significant factor in the solvency of an insurance company, we think that a stress of 25% would be woefully inadequate. That stress may have been calibrated against European property markets where they have less importance to the economy and to the nation's psyche. It is not reasonable here.

Therefore we believe that an internal model is necessary for companies with significant amounts of ERMs backing annuity liabilities.

In 1995 a SoAI paper (Demographic Margins for Prudence – Jeffery & Quinn, 1995) suggested that a valid approach to setting margins was to consider how it would look to a public with the benefit of hindsight if it has gone wrong, noting that with the clarity that hindsight brings can be harsh.

The paper then went on to enunciate 3 common sense principles that it believed should be applied:-

1. If something has happened before, it can happen again
2. If it has happened elsewhere, it can happen here
3. If it happens, when it happens it will happen faster than last time.

For these reasons we would strongly recommend that the stress test should be as least as strong as what has been experienced here since 2007. To put this in number terms we suggest that, in Ireland, a suitable test could be an instant fall of at least 55% which lasts for 5 years and then recovers at the rate of 5% per annum after that. The second principle implies that not just the Irish should be thinking about this.

7.2 Correlations

Internal models as constructed by UK life companies generally use a statistical approach with proxy models using copulas. These are generally regarded as fit for purpose and many have been approved. It should be noted that the use of a proxy model should be validated by out of sample testing and that this can be particularly difficult with with-profit portfolios.

However an issue that needs attention is that of casual links. If an extremal value of one variable leads to that of another then the statistical relationship derived from across the range of the variable normal range may not be valid.

There are two instances that we would note that need attention.

Firstly, interest rates do have an effect upon house prices. If they rise then those with mortgages may find it harder to meet payments and new purchasers less able to buy at the current price. The brief property price falls in the UK in the early 1990's are thought to be related to higher interest rates. This led to a double-whammy on the non-life insurance companies who had written domestic mortgage guarantee business. On the other hand, high interest rates may be accompanied by high rates of inflation and that this may push up house prices.

Secondly, careful consideration needs to be given at the macro level to how interest rates, house prices (including dilapidation) could be correlated with improvements in longevity.

We recommend that something akin to the Irish house price experience from the 2008 crisis should be the stress test for residential property prices in Ireland.

8 Benefits of ERM to lenders, borrowers and society

8.1 Are ERM's a suitable product for backing annuities?

In the David Rule letter of 2nd July 2018, after strengthening the regime for valuing ERMs he says

We continue to believe that restructured ERMs are an appropriate asset to back annuities as part of a diversified portfolio.

But we have not seen any clear evidence why this should be so and feel that the question warrants some exploration. It is clear that one reason why companies might want to back annuities with ERM is that they give lots of yield. The MA (spread you can use above risk free rate) is highest of any asset class, as shown in this table (Source PRA Dear CEO letter, 2018)

Asset class	Spread above Risk-Free	Matching Adjustment
Sovereigns – UK	0.55%	0.55%
Corporate bonds	1.85%	1.25%
ERMs	3.50%	2.00%
Infrastructure	2.10%	1.50%
Social Housing	2.10%	1.60%

But that cannot be the sole criterion.

8.1.1 Other factors to consider

There are several factors that should be considered in assessing whether ERM's are in fact really a good asset to back annuities. We would suggest that these should be

a) Are ERMs a good longevity hedge (as has often been stated)?
b) Annuities demand cash to the policyowners, how does the liquidity of ERMs compare?
c) What are the risks in ERMs?
d) What are macro-economic implications of holding ERMs to back annuities?

8.1.2 Longevity Hedge

Using the projection model described in the Appendices the cash flow to the ERM provider can be projected and then discounted back to the present to calculate the value to on different longevity assumptions. The table below gives the values for our starting longevity position but with different assumptions for the rate of longevity improvement.

	Rate of Longevity improvement			
Joint Age	0.5%	1.5%	2.5%	3.5%
60	50,529	48,252	44,902	39,008
65	52,562	51,204	49,012	45,116
70	52,255	51,944	50,992	48,921
75	49,852	50,270	50,330	49,745

So this table raises several issues. Firstly it gives values much greater than the loan, because it is accumulating at 5% (until NNEG bites) and we discount at 3%. We have assumed a rate of property growth of 1% p.a. (more on this below)

More interestingly the supposed longevity hedge does not exist, except at old ages for low longevity improvements. Why is this the case? Because for the older people with low improvements the increase in the longevity is mostly increasing the period when the loan is accumulating and does not exceed the value of the property. For all other cases the extension of life has most effect when the property is the loan.

So this appropriate asset aspect simply is owning the property and relying on property values not falling (or more onerously at least equalling the discount rate of the insurance company).

So if we pump the house price growth rate up to 3%, then we get a different picture.

	Rate of Longevity improvement			
Joint Age	0.5%	1.5%	2.5%	3.5%
60	66,944	71,013	78,135	73,166
65	60,780	63,833	67,780	69,994
70	55,346	57,481	60,404	63,509
75	50,666	52,082	53,994	56,451

We can show this in a chart. These figures relate to a join age of 70, different levels of house price growth and of mortality improvement.

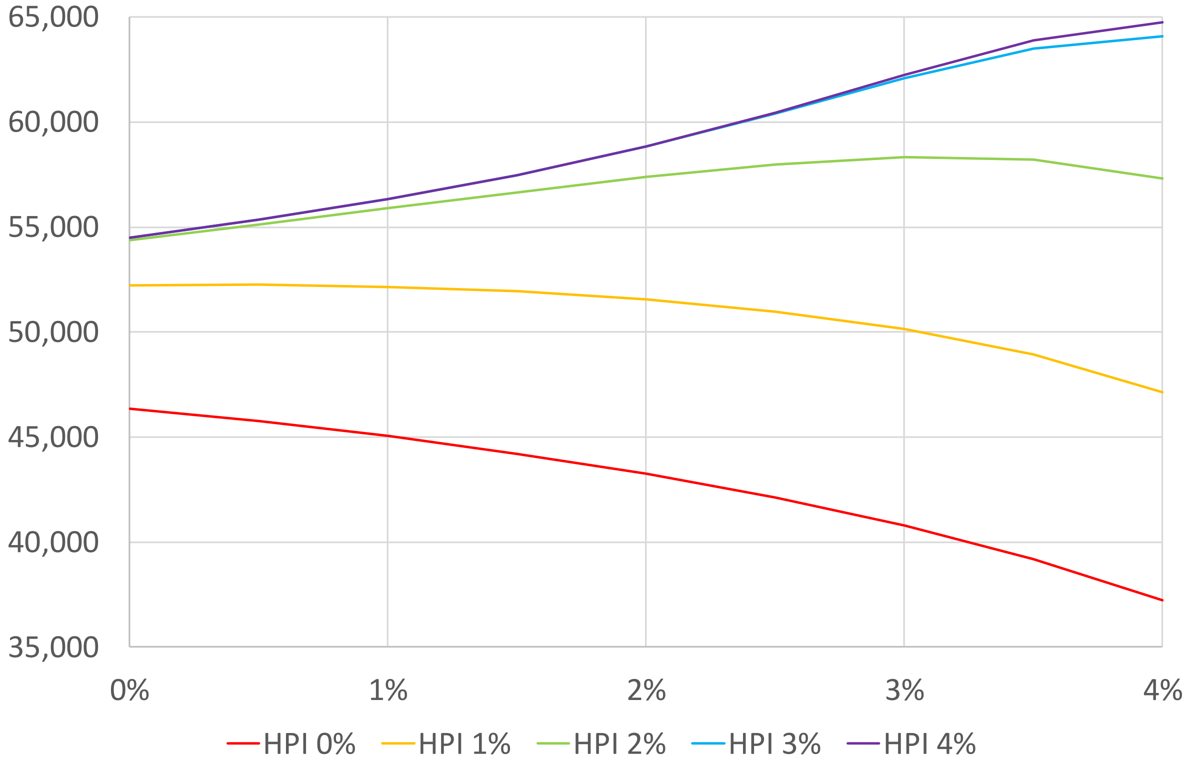

These lines slope upwards (longevity hedge) if we assume a high rate of property growth, but slope down (longevity risk concentration) if property growth is lower.

Is this the right approach? To say that the ERM is a longevity hedge because the property values are going up?

If we accept the deferment rate being positive (which we do) then we have a quite different picture. Now instead of comparing the rolled up loan with an increasing "real world" assumption we have to applying a deferred value decreasing by the deferment rate of at least 1% per annum and then rolling it up by the risk free rate.

The widely held view that ERMs offer a longevity hedge for annuities emerged at a time when most insurers were using ‘real world’ NNEG valuation methods, with negative deferment rates. As deferment rates used in the market have reduced, the hedge effect has also diminished and even may have changed sign.
Any claim of longevity hedge should be re-assessed with each revision to deferment rates.

We have performed these calculations with deterministic projections, in order more clearly to indicate how mortality sensitivities can change. We recognise that calculations are much more complicated in practice. Firstly, of course firms would take account of the time value of the NNEG option, which our calculations here ignore. Secondly, if one firm changes their mortality basis, this doesn't necessarily change the market rollup rate on new lending. So, in order to eliminate day zero profits, that firm's implied ERM discount rate would then change as a result of their new mortality basis. The new discount rate offsets the mortality change on new lending, but not on the back book. And then as a result of the change in the discount rate, so the MA in turn changes and thus also the PV of any annuity liabilities.

8.1.3 Liquidity

Superficially it is hard to think of a worse asset than ERM to match the liquidity profile of annuities. ERMs pay out only when somebody dies and annuities pay out while people live. Of course, they need not be the same people. A portfolio of ERMs on older people may generate a steady stream of income as the loans are repaid out of property sales, but let us look at that more closely.

To start with we must dismiss the idea that liquidity can be considered as flowing from portfolio sales of ERMs to other companies. Recently there have been sales of portfolios of ERMs. However they are only taking place between insurance companies in the UK. Insurance companies in the UK, benefit from being allowed to take credit for the Matching Adjustment and in addition have had a regulatory regime applied that appears to be no longer acceptable. But the key is that liquidity that only flows in good times is not liquidity. To be reliant on another company for your cash when times are tough is to put yourself at the wrong end of the negotiating table.

So what about the flow of cash that comes from the ERM loan unwinding at the death of the borrowers? To get the cash the property must be sold, so the liquidity of the ERM cannot be better than that of residential property. In fact when the NNEG nearly bites, the need to negotiate with heirs can make the sale process more cumbersome than when the NNEG is deeply in or out of the money, when only one vendor is involved.

Actuaries with experience of the crisis in Ireland will know that for a long while property simply did not change hands.

ERMs are therefore to be considered as highly illiquid. Does that matter? Well the annuities cannot be called in by the policyowners (or anybody else) so the only issue is: can the company meet the nearer regular payments as the fall due from other sources? For this reason, many of the ERM restructuring arrangements include liquidity buffers, with a pool of cash or other liquidity facility earmarked for this risk. This is particularly critical more recent ERMs on younger borrowers, especially joint life cases.

When it comes to cash flow it is now clearly the case that improving longevity hits both the need for cash and its availability. A longevity stress should be incorporated into any liquidity projection.

We cannot quantify this for a company without lots of specific data but it appears to us that the holdings of ERMs on younger lives in any quantity may give serious liquidity issues.

We would expect any company investing in ERMs to have thought carefully about liquidity risks

8.1.4 Property Risk

ERMs are a hybrid between a property asset and a loan. Loans are considered as suitable assets to back annuities whereas property is definitely not. So how do we distinguish between the property and the loan portions? A simplistic approach might be to apply a stress to the value of property and say that if the loan is greater than the stressed value of the property then it is more property like and if not then it is more bond like.

Based on this approach and a stress of 20% an ERM starting at 35% LTV loses its bond nature completely at the following durations (for a couple 60 years old).

Stress	Loan Duration
0%	27
15%	22
25%	19
35%	16
45%	11
55%	7

But this misses the point, at a stress of 35% while the 15 year loan is not under water, it will be in a year's time. So instead the table below gives the percentage of loans that would have the NNEG applying if there was a stress of X% applied at the start of the loan. So example for the 35% stress every payment after year 16 would be purely property.

This table therefore shows what percentages of loans are going to be limited by the value of the property after the stress (again for an age 60 couple)

Stress	Loan Duration	Property Percentage
0%	27	87%
15%	22	95%
25%	19	98%
35%	16	99%
45%	11	100%
55%	7	100%

Which basically says that for that particularly case these are not investments in loans, they are in residential property.

Of course taking these figures are very dependent on the age, the LTV and the rate at which property increases in value. It could be argued that a more reasonable scenario stats at age 70 and allows 3% per annum house price increases. This gives the table below.

Stress	Loan Duration	Property Percentage
0%	54	0%
15%	46	2%
25%	39	9%

35%	32	28%
45%	23	39%
55%	13	95%

But think about that 3% for a moment. If properties are going to return 3% per annum, you could simply invest in property and have done. You would then, in the event of a house price crash, at least get some of you money back immediately. So the table below is for 70 year olds assuming no house price growth.

Stress	Loan Duration	Property Percentage
0%	21	72
15%	18	82
25%	15	90
35%	12	95
45%	9	98
55%	5	100

It is clear that ERMs, whilst not by any means identical to property, are very exposed to property risks. We would not wish our own pensions backed by them unless the LTVs were low and ages high.

8.1.5 Dilapidation Risk

As people get very old, they sometimes find it harder to look after their property and indeed themselves. This can be because they are not physically up to it, they have insufficient money or because they lose interest in dong so. There are also a perception effects. What somebody perceives as being necessary in a home depends on what they are used to. What people buying a home perceive as being necessary depends on what they are used to. If you have always had round pin plugs why would you want to pay good money to go to install square? We also fail to notice gradual changes in our own environment e.g. need for repainting or replacing.

All this means that there can be a substantial difference in the value of a property because of its state of repair. In theory borrowers of ERM are required to keep their property in good repair but this can be difficult to enforce.

The Eumaeus Project (post dated 18th February 2019) looks at some data from Aviva

It turns out that if all properties had followed the index, no NNEG would have been exercised, and all properties would have been safely out of the money. The exercise was in all cases due to the underperformance, often a dramatic underperformance, of the properties used as collateral.

As an extraordinary example, consider the property that caused the large blip in 2016. It was originally valued at £1.2m, with an estimated LTV of 45%, i.e. a loan value of about £540,000. (Aviva do not provide an explicit loan rate, but I estimate about 7% based on redemptions and loan amounts at exit). The loan value at exit was £1.4m, but the sale price of the house was only £625,172, leaving a NNEG loss of £763,225.

In other words, while the Halifax index went up 70%, with the indexed house value being over £2m – easily enough to cover the loan value at exit of £1.4m – the property not only failed to follow the index, but actually fell in value (by about 50%). And so it was with 44% of the properties where the

NNEG was exercised. I.e. nearly half the properties used as collateral for equity release not only failed to match the index, but were worth less than when they first collateralised the loan.

We have not attempted to investigate this issue any further but would simply point out that there is a risk which should not be modelled by merely knocking off a small percentage from the ultimate sale price because there is optionality at play due to the NNEG.

The optionality effect of dilapidation needs modelling

8.1.6 Macro-economic Implications

The 2008 crisis demonstrated that at the time the Irish economy was over dependent on property. Imagine that if at the same as the banks were needing rescue, the insurance companies were also reporting trouble from their annuity books being under water. As it was a number of companies had problems with unit linked property funds where delays to redemption on the units of 6 months were proving inadequate (all such companies did successfully manage the issue it should be noted).

We believe that Ireland does not want the solvency of its insurance industry to be property dependent. Therefore a strong test of solvency should be required

8.1.7 Investment issues – summary

We believe that the main advantage of ERMs as an investment for annuity writers, lie in their exceptionally high yield. Anticipating some of that yield in best estimate liability calculations may be justifiable – we have outlined the illiquidity pricing arguments for this. However, we should not be blind to the risks specific to ERMs. Therefore, either the percentage given over to ERMs in an annuity portfolio should be small, or the company must have lots of capital to back it, or both.

8.2 Are ERMs a good idea from the purchasers' point of view?

8.2.1 Promotional Claims

If you play FreeCell in Islington (and who doesn't?), then you will be besieged with clickbait from the Daily Telegraph telling you that ERM's are wonderful and that you should buy one (via the Daily Telegraph) right now. Alternatively, you might watch daytime TV in the UK and see happy grandparents making everything wonderful for their children and grandchildren, going on cruises and building extensions all with the help of an ERM.

You might wonder if this is all too good to be true. So do we.

8.2.2 Is equity really released?

It is however, not surprising that the reality is somewhat different. The first and most important consideration is that the term ERM is itself potentially misleading. The Equity is not released at all, it is borrowed against. As house prices move up and down (!), the loan remains unchanged in value. It is highly unlikely that geared equity products would be deemed suitable investments for people in the disinvestment phase of their lives, why then should ERM be?

8.2.3 Does the NNEG Mitigate Consumer Risk?

Well, of course there is the NNEG. That does slightly mitigate risks but in the event that an ERM of 30% had been taken out just before the Irish crash, the residual value would be 15% immediately after. If the retirees are depending on their home for their future income, they would have no more money to come.

8.2.4 Financial Advice

Money Saving Expert on the subject says

Equity release isn't something to be taken on lightly, so before you dive right in, first evaluate whether downsizing your property could be an option. If you can sell up and move on to a smaller home, and live off the excess cash you have made, great.

8.2.5 Factors for Assessing Suitability

So what are the factors that need to be taken into account in assessing whether an ERM is correct for a person? They are, not surprisingly cost/value, risk and alternatives. But the situation is made much more complicated by the question of long-term care which we have already discussed in section 6 above.

So what exactly is the cost of an ERM? From the simplistic view a loan at a rate of around 5% is much lower than the cost of bank loans or credit card loans. But this is because the borrowing is against the value of the asset of the home.

If we take a risk neutral view of the world and assume that property will roll up at the risk free rate, then the projections in Appendix c show that the loan will pretty much swallow the property. The borrower in this case should consider that they have effectively sold their property for 35% of its value plus the right to stay in it until joint life last death, which may not look attractive.

So one has to take a more optimistic view to consider doing this (if one has a choice). There are other alternatives. Which is better taking out an ERM? Or downsizing? Or to put it another way, by how much does the house value go up in order for it to be better to ERM and gear rather than down size? Well, obviously in the long run the house price would have to keep up with the rate of loan. The most optimistic forecasts indeed show this happening, but even the ‘real world’ house growth assumptions now common in the industry are below the rollup rates on ERMs.

Downsizing is not a panacea, either. Moving house can be stressful, especially for older people who may be less able to adapt to new surroundings. The change can be particularly difficult if the downsizing involves moving to a new area in order to find a cheaper property. Downsizing incurs costs such as legal fees and stamp duty.

These are not the only two alternatives. Borrowers should also consider a 'reverse mortgage facility' which could be dynamic and offered only in a context of professional advice with regular reviews.

In many cases the stated benefits of ERMs may be achieved just as well by downsizing or reverse mortgage facilities. From a consumer perspective, ERMs should be evaluated by careful comparison, both quantitatively and qualitatively, to other options, all of which have advantages and disadvantages.

8.3 Are ERMs in the Public Interest?

Throughout the discussion in the UK, has been the assumption that they are a GOOD THING. This does not seem to have received much critical examination. We give below some points in favour and against ERM from the Public Interest point of view.

Points in Favour

1. ERMs permit asset-rich cash-poor borrowers to get hold of money. It means that people can fund their retirement from their own home. This has some merit indeed, for those who either do not want to downsize or cannot (possibly because they have already downsized as far as they can go).
2. They permit higher annuity rates. ERMs are believed to be highly profitable. This can be verified without any financial or actuarial knowledge but simply using marketing theory. How to tell products with high margins – look for those that are heavily promoted. Some of that margin may be passed on to consumers in the annuity rates offered, although it may instead get competed away.
3. They permit annuity buy outs on better rates. This is definitely a good thing. Some pension schemes in Ireland are not as healthy as one might like. Allowing them to access better buy-out rates is welcome.
4. Raising cash from homes to allow deep retrofit may help Ireland meet its carbon targets

Points against

1. There is an acute shortage of family homes in Ireland and particularly in Dublin. This could be eased to some degree if empty-nesters were encouraged to down-size and move into smaller homes.
2. Care of more fragile people may be easier and cheaper in retirement villages or other purpose-built arrangements.
3. The housing stock may be adversely affected by dilapidation.
4. Younger occupants of family homes will have longer time horizons and therefore may be more willing to commit to the deep retrofit that is necessary for Ireland to meet its carbon commitments.
5. We have grave concerns if the ERM market makes the solvency of insurance companies dependent on a historically cyclical residential property market.

We believe the case is finely balanced. We find it hard to agree that the arguments are overwhelmingly in favour of ERMs in a majority of cases.
We would also suggest that there is an easy win for the Irish government to level the playing field from a tax perspective, either by making downsizing exempt from stamp duty or by charging an equivalent tax on ERMs.

9 Conclusions

The aim of our paper was to list what we think should be considered, to help the actuaries develop a position on the possible reintroduction of ERMs into Ireland and in particular the relationship to writing of annuities. The path in the UK has been long and winding and there are many lessons to be learned.

The reader will also conclude that a product that has been painted in a flattering light elsewhere, also has some shortcomings and difficulties which need to be weighed carefully.

From a technical point of view, ERMs raise difficult questions about how to assign a value to a product that contains embedded options on illiquid assets. This requires us to synthesize option pricing theory and illiquid asset pricing theory. While the valuation formulas are not complicated (by option theory standards), inputs proposed by different authors span wide ranges. We have outlined the theoretical arguments for the various positions.

We feel that actuaries should be careful to apply their skills of financial analysis and should not uncritically amplify industry claims. If we, as actuaries, fail to provide sufficient challenge, then we open ourselves up to charges of being captured, as has been alleged in the UK.

We hope that this work stimulates the debate in Ireland. Ours are not the final words nor should they be.

Let a hundred flowers bloom, let a hundred schools of thought contend.

10 References

Adams, Andrew T., 2000, Excess volatility and investment trusts, Working paper, University of Edinburgh.
Ashurst, R., Blundell, G., Booth, P., Cumberworth, M., Morell, G., Pugh, R. & Smith, A.D. (1998). Institutional Investment in Residential Property. Institute of Actuaries Investment Conference.

Aviva (2019) Equity Release Funding Reports. https://www.erfunding.co.uk/.

Black, F. (1976), 'The pricing of commodity contracts', *Journal of Financial Economics* **3**, 167–179.

Booth, P. & Marcato, G. (2004), 'The measurement and modelling of commercial real estate performance', *British Actuarial Journal* **10**(1), 5–61.

Bulley, A., H. Jupe and U. Choudry "Equity Release Mortgages – where Does the PRA Stand? Centre for Regulatory Strategy, Deloitte. September 2017.

Davis, M. H. A., Panas, V. G., Zariphopoulou, T., 1993. European option pricing with transaction costs. SIAM Journal of Control and Optimization 31 (2), 470–493

Dimson E, Staunton M & Marsh P (2002). Triumph of the Optimists. Princeton.

Dowd, K. (2018), Asleep at the wheel: The prudential regulation authority & the equity release sector. Adam Smith Institute.

Equity Release Council (2018), Autumn 2018 market report, Technical report.

Geske, R (1979) The Valuation of Compound Options. Journal of Financial Economics. 7 (1979) 63-81.

Graham, B and Dodd D (1934) Security Analysis.

Goodley, S & Treanor J (2016). Property funds halt trading as Brexit fallout deepens. The Guardian, 5 July 2016. https://www.theguardian.com/business/2016/jul/05/aviva-halts-trading-in-its-property-fund-brexit-standard-life

Glosten, L. R., and P. R. Milgrom (1985): Bid, ask and transaction prices in a specialist market with heterogeneously informed traders, Journal of Financial Economics, 14(1), 71–100.

Hagan, P S; Kumar, D; Lesniewski, A S and Woodward, D E (2002). Managing Smile Risk. Wilmott magazine.

Heston, Steven L. (1993). A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options. The Review of Financial Studies. 6 (2): 327–343

Hodges, S. D., Neuberger, A., 1989. Optimal replication of contingent claims under transaction costs. Review of Futures Markets 8, 222–239.

Hosty, G. M., Groves, S. J., Murray, C. A. & Shah, M. (2008), 'Pricing and risk capital in the equity release market', *British Actuarial Journal* **14**, 41–109.

Ji, M., Hardy, M. & Li, J. (2012), 'A semi-Markov multiple state model for ERM terminations', *Annals of Actuarial Science* **6**(2), 235–257.

Kemna, A. G. Z. & Vorst, A. C. F., 1990. A pricing method for options based on average asset values, Journal of Banking & Finance, vol. 14(1), pages 113-129.

Kenny, T, C. Golding, G. Craske, A. Dobinson, S. Gunter, O. Griffiths, N. Hayes, A. Mockridge, S. Robertson, R. Saundh, J. Thorpe (2018). Actuarial management of equity release mortgages. Institute and Faculty of Actuaries.

Leland, H., 1985. Option pricing and replication with transaction costs. Journal of Finance 5, 1283–1301.

Li, J., Hardy, M. & Tan, K. (2010), 'On pricing and hedging the no-negative-equity guarantee in equity release mechanisms', *Journal of Risk and Insurance* **77**, 499–522.

Marcato G (2015) Liquidity Pricing of Illiquid Assets. IPF research.

Mitchell P (2015). Individual Property Risk. IPF research.

Prudential Regulation Authority (2018), Solvency II: Equity release mortgages, Supervisory Statement PS31/18, Bank of England.

Prudential Regulation Authority (2018), Solvency II: matching adjustment – illiquid unrated assets and equity release mortgages, Supervisory Statement SS3/17, Bank of England.

Pastor L. and Stambaugh R.F. (2003). Liquidity Risk and Expected Stock Returns. Journal of Political Economy, 111, 642–685.

Perotti E, Danielsson J, de Jong F, Laux C, Laeven R & Wüthrich M (2011). A prudential regulatory issue at the heart of Solvency II. Vox EU.

Radevsky A & Greenish, D (2017). Hague guide to Leasehold Enfranchisement. 7th Edition. Sweet & Maxwell.

Rouwendal, Jan, van der Straaten, Willimijn, 2003. Dual Earners, Urban Labor Markets and Housing Demand. Tinbergen Institute Discussion Paper 2003-084/ 3.

Royal Institute of Chartered Surveyors (2009). Leasehold reform – graphs of relativity.

Smith A D & Spivak G (2012) Liquidity and Investment Strategy. Life Insurance Convention. Institute of Actuaries.

Van Loon, P R F, Cairns A J G, McNeil A J and Veys, A. (2015). Modelling the liquidity premium on corporate bonds. Annals of Actuarial Science, Vol. 9, part 2, pp. 264–289.

Webber, L. (2007). Decomposing corporate bond spreads. Bank of England Quarterly Bulletin, 47(4).

Wilson, J. (2003). Valuing the Freeholder's Interest. Royal Institute of Chartered Surveyors.

Appendix: Projection Basis for Deterministic Calculations

To investigate ERMs we need to make projections of how they will turn out under different assumptions.

The base longevity tables used are PNMA00 and PNFA00 (collectively known as PNXA00). The prime reason for choosing these is that they are publicly available, so anybody can use them.

It seems likely that an insured lives table would be more appropriate than a population study. All prospective borrowers have to be homeowners, and in general we believe would be better off homeowners in order to have equity worth releasing.

The ILMI study is the only recent study of Irish insured lives mortality both assured and annuitant lives. Although the main body of the study did not compare the experience against PNXA00, that comparison is given in an appendix and is 84% for males and 92% for females.

We have assumed that the lives of partners are independent of each other. It is well known, of course, that they are not but we are not aware of any data quantifying this effect. To bring this feature in would mean reducing mortality of both lives while both are alive and increasing it in widow(er)hood. We suspect the effect on our results would be small.

For improvements, we have simply taken the CSO assumed long term rate of improvement. This is 1.5% for both genders. The CSO has greater short-term assumptions which it blends into the longer-term rate over a number of years. Using the simple assumption keeps the spreadsheet maths simpler. We also have some doubts whether the improvements of the beginning of the century are going to continue. We have used this rate to improve the longevity from the central year of the ILMI study until the present and then into the future.

We have done projections on a joint life basis assuming the couple is composed of one male and one female of similar age. The major sensitivity would be if one the partners was very much younger but we would expect that to be specifically underwritten.

We have assumed that the ERM ends on death, not on entry to long term care. Whether it would be the case that it becomes common for old very ill people to burn their boats and sell up on entering a nursery home remains to be seen.

It is worth drawing attention to how long an ERM will last. Joint life significantly extends duration. People have probably become attuned to the idea of retirements lasting 20 plus years but this table shows the median and 95% boundary for the length of ERM on our base table for various ages at inception.

Joint age at inception	Median Duration (years)	95% Percentile Duration (years)
60	36	51
65	30	45
70	25	39
75	20	33

For simplicity we have assumed that the risk-free rate is a flat 1% p.a., that the rate of interest charged on the loan is 5% and that the Matching Adjustment is 2% so that life companies can discount at 3%. For this purpose, we have assumed that the whole portfolio receives the MA. This will not be possible in practice but again we are trying to keep the maths down.

We have not allowed for any frictional costs in our projections. These are far greater for property than for such things as equities and gilts. Bringing them into the projections would make ERMs look worse as investments to both borrowers and lenders.

We have not allowed for delays in selling properties after death of borrowers. We would expect these to make ERMs look worse for lenders.

We have not allowed for ERMs repaying early due to entry of last survivor into long term care. We would expect these to make ERMs look better for lenders.

We have not allowed for other early repayments. We would expect these to make ERMs look worse for lenders as they are likely to only happen when house prices have not collapsed so some optionality will have been lost.

We have not allowed for dilapidation effects. Anybody who has been house hunting in Ireland will be familiar with homes that are being sold after the death of the elderly occupant. We would expect this to make ERMs worse for lenders, especially as lone aged occupants are far less likely to move if the ERM has consumed all the value of the property.

Results of projections of Value to Borrowers

Using the longevity basis we have outlined, it is easy to do some projections of how an ERM turns out from a borrower point of view. The other assumptions needed are the rate of accrual of loan which we have taken to be 5% per annum and the rate of house price growth. This last factor is critical as will be seen. For simplicity's sake we have ignored dilapidation costs and transaction costs. It is not hard to see how bringing these into account would affect figures.

Our base line for house price growth is a risk neutral approach. That means that you roll up at the risk-free rate which for simplicity we have set at 1% per annum. At this rate the picture is fairly stark. The mean and median amounts passed on to the heirs from a €100,000 property subject to 35% loan is, naturally age dependent

Joint age at inception	Median Inheritance (€)	Mean Inheritance (€)
60	0	3,388
65	0	8,391

70	11,908	17,225
75	30,761	28,817

These inheritance values are not present values but absolute amounts. Present values are of course smaller. Perhaps 1% may be too pessimistic. If we take a typical inflation targets for Central banks which might be 2%. Then we get

Joint age at inception	Median Inheritance (€)	Mean Inheritance (€)
60	15,917	4,102
65	31,694	28.225
70	45,752	41,352
75	56,484	51,932

If we go to 3%, then everything in the garden looks lovely

Joint age at inception	Median Inheritance (€)	Mean Inheritance (€)
60	87,720	78,270
65	91,525	85,247
70	90,628	87,026
75	87,327	85,416