THE EUMAEUS GUIDE TO EQUITY RELEASE VALUATION
Restating the Case for a Market Consistent Approach

Dean Buckner and Kevin Dowd
[This version: 15 July 2019]

Executive Summary

- The UK equity release sector is permeated by poor valuation practice: as far as we are aware, not a single equity release firm is valuing its No-Negative Equity Guarantees (NNEGs) in a scientifically valid manner.

- This NNEG under-valuation problem is on a large scale and implies correspondingly large over-valuations of Equity Release Mortgages (ERMs).

- The Discounted Projection or 'Real World' approach used by the equity release industry is inherently flawed and produces valuations that violate bounds that are known to be inviolable.

- The only scientifically valid valuation approach is the Market Consistent approach, which is also the only approach compatible with accounting principles and technical actuarial standards.

- Market consistent valuations cast doubt on the profitability of ERM loans especially to younger borrowers.

Acknowledgements

We thank the many people who have kindly contributed to this report through helpful discussions on the topics covered here or through comments on earlier versions. We thank in particular: David Blake, Chris Cundy, Jonathan Ford, Charles Goodhart, Tony Jeffery, Howard Mustoe, Gordon Kerr, Cavin O'Driscoll, Andrew Smith, John Skar, Craig Turnbull, members of the PRA's CP 7.19 team, seminar participants at the London School of Economics and Queen Mary University of London, and various correspondents who have offered comments in private. The usual caveat applies: we alone are responsible for any views expressed here, and for any remaining errors.

Disclosure Statement

This report has not been funded by any research organisation or commercial body, and we have not sought any such funding.

Dowd is currently involved as a co-investigator in the Actuarial Research Council's "Modelling, Measurement and Management of Longevity and Morbidity Risk" research programme. For more details, see:
https://www.actuaries.org.uk/learn-and-develop/research-and-knowledge/actuarial-research-centre-arc/research-programmes/modelling-measurement-and-management-longevity-and-morbidity-risk/research-team

Chapter One: Introduction

> "I do urge you all to keep an open mind … because we [actuaries] don't always - sometimes we've got a very narrow way of thinking, but I do think that there's more to go [on No-Negative Equity Guarantee valuation], and this may go on for longer than the Brexit discussions."
>
> Gina Craske[1]

In the Equity Release Council's Spring 2018 *Market Report*, its chairman David Burrowes struck a reassuring tone:

> Annual lending activity by our members has surpassed £3 billion for the first time and customer numbers reached 67,000 in 2017. Property wealth is increasingly recognised by people as a safe and sought-after source of retirement finance, with the market attracting twice as many new customers as it was five years ago. …
>
> The range of product options available to equity release customers has grown 25% year-on-year, providing more choice to underpin a robust and competitive market.
>
> Looking forward, we expect the need for new sources of income in retirement will continue to grow as many people will be unable to rely on pressured pension pots. (Equity Release Council, 2018, p. 2)

Mr. Burrowes omitted however to mention an issue that had been causing some worries in Equity Release Mortgage (ERM) circles for a little while now. The problem is that firms are under-valuing the No Negative Equity Guarantees (NNEGs) that are a standard feature of most ERM products. This under-estimation seems to be on a large scale too.

These concerns received some publicity with the publication on 7 August last year of reports by BBC business journalist Howard Mustoe[2] and the Adam Smith Institute[3] on the issue, and with the airing that evening of a BBC Radio 4 programme, "The Equity Release Trap."[4] Since then, there has been considerable public discussion of the NNEG valuation issue. We provide a commentary on our blog, The Eumaeus Project (eumaeus.org/), made a presentation on the subject to the London School of Economics in October and published a second analysis in our Johns Hopkins University discussion paper of November 2018.[5] Another report on NNEG valuation by Radu Tunaru was published by

---

[1] Staple Inn event transcript 28 Feb 2019.
[2] H. Mustoe, with additional reporting by M. Keyworth (2018) "Home Equity Release May Cost Pension Firms Billions." *BBC news* website (7 August).
[3] K. Dowd (2018) *Asleep at the Wheel: The Prudential Regulation Authority and the Equity Release Sector*. London: Adam Smith Institute (7 August) (a)
[4] https://www.bbc.co.uk/programmes/b0bd8h78
[5] D. Buckner and K. Dowd (2018) "Equity Release: Another Equitable in the Making," Johns Hopkins Institute for Applied Economics, Global Health, and the Study of Business Enterprise, *Studies in Applied Economics* No.129, November 2018.

the Institute and Faculty of Actuaries in February this year and given extensive coverage at its Staple Inn launch.[6] A further report on NNEG valuation was presented to the Society of Actuaries in Ireland by Tony Jeffery and Andrew Smith in March, and we gave another seminar on the subject at QMUL in June.[7] It would be fair to say that there is a wide range of views and the issue is now more controversial than ever.

Consider the following 2017 quotes from UK equity release firms discussing the methodologies they use to value their No-Negative Equity Guarantees (NNEGs). As you do so, ask yourself what they all have in common:

> "When calculating the value of the no-negative equity guarantee on the lifetime mortgages, certain economic assumptions are required within the variant of the Black-Scholes formula. […] In the absence of a reliable long-term forward curve for UK residential property price inflation, the [firm] has made an assumption about future residential property price inflation. … This results in a single rate of *future house price growth* of 4.25%."[8]
>
> "[The value of the NNEG] is calculated using a variant of the Black Scholes option pricing model. The key assumptions used to derive the value of the no-negative equity guarantee include current property price, *property growth* and property volatility."
>
> "Stochastic modelling is used to capture the expected cost of [the NNEG], which will depend on the expected rate and volatility of *future house price growth* …
>
> "Equity release and securitised mortgage loans … are valued using an internal model. Inputs to the model include primarily *property growth rates*, mortality and morbidity assumptions, …."
>
> "The fair value of the guarantee is determined using a stochastic model. The fair value of the loans is determined using assumptions for interest rates, *future house price inflation* and its volatility …"

The answer is that they are all using incorrect valuation approaches. They all use property growth assumptions in their NNEG models, but no correct option pricing models include property growth variables. Their use of an irrelevant variable then indicates that they are not valuing their NNEGs properly.

To their credit, the PRA have been aware of this problem for some time. Referring to the results of an earlier survey, CP 48/16 states (p. 25):

---

[6] R. Tunaru and E. Quaye (2019) "UK Equity Release Mortgages: a review of the No Negative Equity Guarantee," Actuarial Research Centre/Institute and Faculty of Actuaries (19 February).
[7] T. Jeffery and A. D. Smith "Equity Release Mortgages: Irish & UK Experience," presentation to the Society of Actuaries in Ireland, 28 March 2019
[8] See F. Black and M. Scholes (1973) "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy* 81: 637–654.

> Many respondents mentioned a version of the Black-Scholes formula known as 'Black 76', where the underlying price is the 'forward price' of the property. This version uses the current price of a forward contract. Some respondents appeared to *conflate* this with the forecast future price of the property, but provided *no justification* for why house price inflation was relevant to the current price of a forward contract. (Our italics)

The key word is "conflate". The reason why these correspondents provided no justification for using projections of future house price inflation to value these guarantees is because no such justification exists.

To spell it out: some firms say that they are using *assumptions about future house price growth*, but the PRA correctly says that this is *obviously wrong*. From which it follows (1) that some firms are using a method wholly at odds with the one endorsed by the PRA and (2) that the PRA would not be bothering to state this point at all, particularly through a protracted consultation period if it had not experienced substantial pushback from firms. We can then infer (3) that firms with equity release exposure have been undervaluing their no negative equity guarantees. We can make this inference because the PRA would not be publishing on the subject or seeking industry consultation if they thought that these guarantees were correctly valued. Consequently, some firms are presumably undervaluing them. Also (4) by a similar logic, if firms are dedicating substantial resources to pushing back, they must think that the valuation of guarantees is a material issue.

In fact, we are not aware of a *single firm* that has *demonstrated* that it is valuing its NNEGs using a defensible methodology. Our impression is that they are all getting it wrong.

Equity Release and the Ghost of Equitable Life

We have seen this movie before. A couple of decades ago, there was a scandal surrounding Equitable Life. The world's oldest mutual insurer, Equitable Life was founded in 1762 and pioneered age-based premiums based on mortality assumptions. In the middle of the 20th century, it also pioneered Guaranteed Annuity Rate (GAR) options that offered guaranteed fixed returns. At its peak in the 1990s, it had 1.5 million policyholders with funds worth £26 billion under management. However, it failed to value these options properly, and in some cases, it didn't value them at all. As a result, it failed to provide for them properly. Equitable came to grief in 2000 when it was no longer able to keep its promises. There was then a big outcry and the insurance regulatory system was overhauled to make sure that an Equitable-style fiasco never happened again. So the problem was that the company had been under-valuing opaque and apparently innocuous long-term guarantees and the undervaluation of these guarantees eventually brought it down.

In both cases, there was a toxic combination of intellectual error and short-term thinking. In the Equitable case, there was an underlying presumption that the guarantees in question, Equitable's GAR options, didn't really matter and that any problems that they might entail were well into the future anyway. In the equity release case, the intellectual error involves a profound misunderstanding of option pricing theory by professional actuaries, combined with a mindset on the part of industry leaders that puts short-term profitability and 'competitiveness' ahead of notions of long-term sustainability. When it comes to NNEG valuation, this mindset prioritises *low* NNEG valuations over *sound* NNEG valuations and the rest is obvious.

The intellectual error centres around the underlying variable in the option pricing formula. A NNEG involves a portfolio of put options and we are dealing with puts on forward contracts. For example, if a customer takes out an ERM at the age of 70, there is a NNEG put for the possibility that the ERM loan might end when the customer is 71, another NNEG put for the possibility that the ERM loan might end when the customer is 72 and so forth. Each of these put options is issued now, but has a horizon (or decrement) of one, two, etc. years in the future. The price that enters into each put option pricing equation is the forward price of the underlying, and the deliverable is a house. So for the put option that ends in future year t, the underlying is the forward house price for year t, the price agreed now for the house to be delivered and paid for in year t. This approach is based on standard option-pricing theory and is exemplified by Black (1976).[9] In actuarial circles, this approach is often referred to as the 'Market Consistent' approach to NNEG valuation.

The problem is that a number of practising actuaries in the UK equity release sector have convinced themselves that the underlying price that is relevant for put option pricing is not the *forward* house price for year t but the *future* house price or expected future price for year t. However, forward and future prices are *very different* and to confuse the two is to commit a major error. This error is a *big deal* because inputting the expected future house price into the option-pricing equation gives very low NNEG valuations, whereas inputting forward house prices into it gives much larger NNEG valuations. This second, incorrect, approach is commonly referred to in actuarial circles as the 'Real World' or 'Discounted Projection' approach.

A difference however between the Equitable Life and equity release cases is that when Equitable started issuing GARs in the 1950s, the valuation of options was not well-understood. The option pricing breakthrough only occurred in 1973 with the publication of the famous articles by Black, Scholes and Merton (Black and Scholes, 1973; Merton, 1973), followed shortly afterwards by Black (1976). Both the principles and the nuances of option valuation have been well known for decades and are taught in universities all over the world.

It is curious, too, that the UK actuarial professional association, the Institute and Faculty of Actuaries (IFoA), has yet to speak out against these unsound NNEG valuation practices. On the contrary, it is on the record as endorsing a number of misconceptions about NNEG valuation.

Despite copious protestations to the contrary, the UK actuarial profession appears to have learned nothing from the lessons of Equitable Life, and so welcome to Equitable 2. Those who fail to learn the lessons of the past will have to take the class again.

---

[9] See F. Black (1976) "The Pricing of Commodity Contracts," *Journal of Financial Economics* 3: 167-179.

Purpose of this Report

The purpose of this report is to set out the issues involved in the valuation of NNEGs and Equity Release Mortgages (ERMs), i.e., it provides a how-to handbook for practitioners working on NNEG and ERM valuation issues. In doing so, it examines both valid (i.e., Market Consistent) and invalid (e.g., Discount Projection or Tunaru) approaches, and explains why the former are valid but the latter are not.

We emphasise that this report focuses only on the valuation issues and does not address the broader issues raised by NNEG under-valuation, such as the implications for public policy. We intend to address those issues in a follow-on report.

Organisation of this Report

This article is organised as follows. Chapter 2 provides an introduction to equity release: it explains the products and the significance of the NNEG, and gives an overview of recent developments in the sector. Chapter 3 explains the basics of NNEG and ERM valuation. The next 14 chapters examine key inputs and other factors relevant to NNEG and ERM valuation: the loan-to-value ratio (Chapter 4), the risk-free rate (Chapter 5), the loan rate (Chapter 6), the theory and calibration of net rental and deferment rates (Chapters 7 and 8), dilapidation (Chapter 9), volatility (Chapter 10), mortality (Chapter 11), long-term care (Chapter 12), delayed possession (Chapter 13), credit spreads (Chapter 14), the impact of drawdown facilities (Chapter 15), prepayment (Chapter 16) and fees and charges (Chapter 17). Chapter 18 examines ERM-related scenario analyses and stress tests and Chapter 19 discusses the PRA's ERM Good Practice Principles. Chapter 20 sets out the Market-Consistent approach and Chapter 21 refutes some common actuarial misconceptions about it, Chapter 22 debunks the 'Discounted Projection' approach, Chapter 23 examines the recent Tunaru report and Chapter 24 discusses Just Group's NNEG valuation model. Chapters 25 and 26 examine technical actuarial and accounting standards, and Chapters 27 and 28 offer some recommendations for good valuation practice and its governance.

Chapter Two: Introduction to Equity Release

The Home Economics of Equity Release

An ERM a type of loan collateralised by a property ('house'), and the type of ERM we are interested in goes as follows.[10] The loan is taken out by a customer late in life who owns the property they live in. The customer uses the loan to supplement their income, help their children get on the property ladder or whatever. Unlike a normal loan, this loan has no fixed end date and involves no regular interest payments. Instead, the loan ends when the customer exits the house, either by death or by going into a nursing home, and the amount owed on the loan accumulates over time until the loan is repaid.[11] At the time of exit, the lender takes possession of the property and sells it to repay the loan. If there are any proceeds left over, these are returned to the customer or to their estate.

The term "equity release" is misleading, however. As Jeffery and Smith (2019, pp. 6, 55) point out:

> 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. .... As house prices move up and down (!), the loan remains unchanged in value.

Jeffery and Smith are right, but the term "equity release" is so widely used that we are stuck with it, and the American alternative ("reverse mortgage") has issues of its own. So "equity release" it is.

There are two main types of ERM. The first, known as a Lifetime Mortgage (LTM), is a straightforward mortgage loan, where the lender hands over the loan amount at the time the contract is made. The second is a drawdown ERM, in which the contract provides for a maximum possible loan amount, but the borrower has discretion over how much to draw down against this maximum and when to do so, subject to the constraint that the total amount drawn down cannot exceed the stipulated maximum. Typically, the borrower in a drawdown ERM would make a drawdown when the contract is made, and then make more in later years, as and when he or she feels the need to do so.

In this report, we focus mainly on the valuation issues relating to LTM ERMs. These are simpler, but we discuss drawdown in Chapter 15. Suffice to note that the valuation issues are much the same, except for the additional complications introduced by a drawdown facility and how that might be used by the borrower.

---

[10] We are not concerned here with other types of equity release product such home reversions, in which the borrower sells all or part of their property at less than its market value in return for a tax-free lump sum, a regular income, or both, but stays on in their home as a tenant who pays no rent. Nor are we concerned with ERM loans that do not incorporate NNEGs, suffice to note that most incorporate NNEGs and that all ERMs issued by members of the Equity Release Council do.

[11] In some cases, the loan can also end by early repayment, but we defer that issue to Chapter 16 below.

The ERM loan will be taken out as some low proportion of the property value – 40% is typical for a 70 year old, but Loan to Value ratios (LTVs) tend to be lower for lower ages and higher for higher ones – and the lender is protected against any risk of loss for as long as the loan value is below the value of the house.

The loan rate will be fixed at the inception of the loan.

The value of the collateral, the house, will vary with the house's market price. Typically, house prices have risen in recent years and we might (or might not!) expect them to continue to rise, but we would not usually expect the house price to rise at a rate exceeding the loan rate. In any case, house prices are uncertain and sometimes fall, so expectations of future house prices are unlikely to be exactly realised.

A typical case is shown in Figure 2.1:

Figure 2.1: Loan Equity in a Typical Equity Release Mortgage

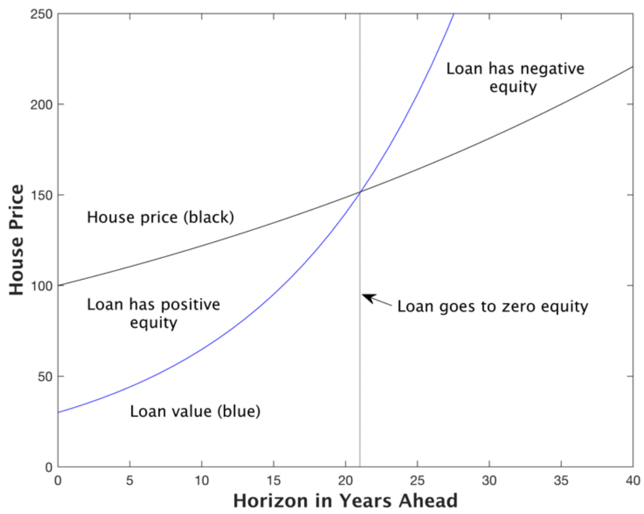

We would certainly expect the loan amount (shown in blue) to rise over time, and we would usually expect the house price (in black) to rise too, but even so, we would expect the loan amount to rise at a faster rate and eventually, if the customer lives long enough, the blue loan amount line will cross over the black house price line. Thereafter the loan amount will exceed the value of the house, i.e., the loan will go into negative equity.

If the customer exits the house before the point of negative equity (which is 21 years in Figure 2.1), then the lender would be repaid in full.

If the customer exits after that point, the loan would expire in negative equity, i.e., the value of the property collateral would not be enough to cover the accumulated loan amount. In the absence of a NNEG, the lender could sue the borrower or their estate, but there might have few assets left, especially if the borrower was moving into a retirement home and any remaining assets were being used to finance their care. Most ERMs incorporate a NNEG, however, and in such cases the negative equity becomes a loss borne by the lender.

Another way to think about the ERM-with-NNEG contract is that it gives the lender the minimum of the house price (black) and loan amount (blue) lines. The fact that the lender gets the minimum of two values indicates that the lender is granting a put option to the borrower.

The lender's potential loss with the NNEG in place is illustrated in Figure 2.2, and let's henceforth assume for the moment that exit is due solely to death (although we will revisit this assumption later):

Figure 2.2: ERM Loan Expires in Negative Equity

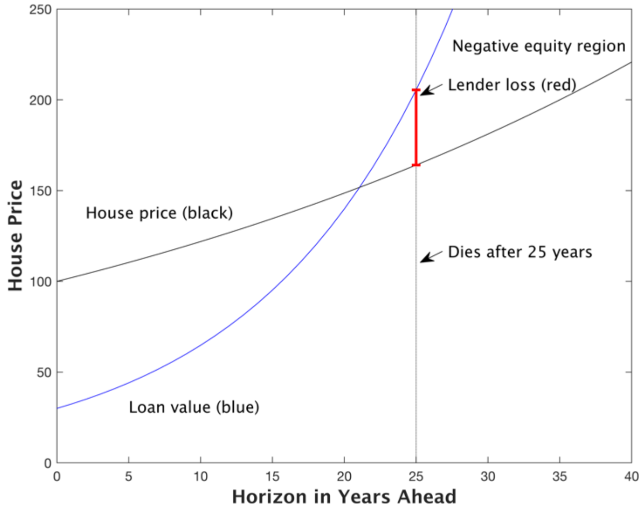

In this case, the borrower dies after 25 years and the lender makes the loss given in red, the difference between the loan value and the house price after 25 years, relative to what the lender would have received had there been no NNEG and the lender been repaid in full.

Naturally, this loss (and whether any loss occurs at all) is uncertain before the event. The timing of death is uncertain and if the customer dies early then there would be no loss to the lender. But if the customer dies later the lender suffers a loss that depends among other factors on the timing of death. Thus, the NNEG potentially exposes ERMs to longevity risk – the risk that the customer might live too long.

ERMs are also exposed to house price risk. This risk is illustrated in Figure 2.3:

Figure 2.3: The Impact of a Fall in House Prices on Negative Equity

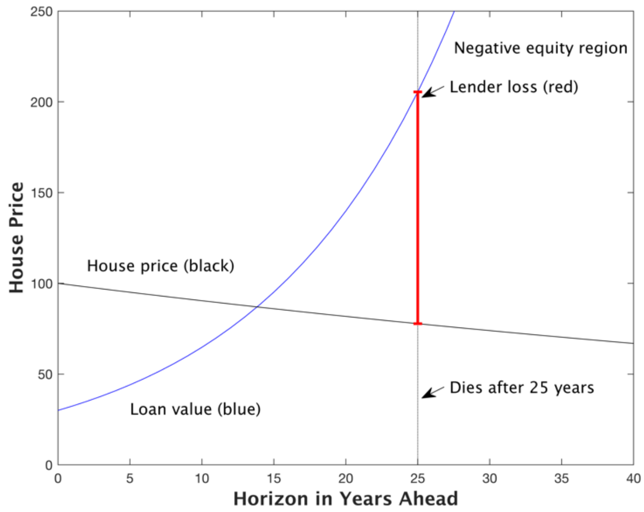

The house price might be lower at the time of death than the lender expected it to be. Figure 3.3 shows a case where the house price declines instead of rising. If the customer dies after 25 years, then it is clear from a comparison of Figures 2.2 and 2.3 that the lender will suffer a larger loss due to the house price fall. ERMs are thus subject to house price as well as longevity risk.

From the borrower's perspective, taking out an ERM loan might be a suitable choice for an older individual or couple who are asset rich but cash poor, e.g., they might have a need for cash or wish for a higher standard of living in retirement. One can also imagine additional circumstances in which an ERM might be suitable, e.g., because their children may be affluent or because they don't want to leave their children any inheritance, or because they may have no children and don't want to leave their house to a cats' home. For such people, a regular mortgage would not normally be practical because they would no longer be working and therefore not have the income to repay such a mortgage.

From the lender's perspective, an ERM loan offers a high loan rate and is highly collateralised, at least to start with. Its main downside is the impact of the NNEG, the valuation of which is the core focus of this report.[12]

To help form an intuition, consider that an equity release mortgage can be broken down into two components. The first component is the loan amount which rolls up at a high loan or roll-up rate, which currently averages at about 5.25% for new ERM loans, but there is considerable disparity around that average. Now this loan is collateralised by the house, and the loan at inception will be some fraction of the house value. For example, for a 70 year old, the loan to (house) value ratio will be around 40%. The high collateralisation of the loan means that we can regard the loan as close to risk-free: when the borrower exits the home, the lender will take possession and use the home to get its loan repaid. There is no danger of default. The second component is the NNEG. So the

[12] It is often claimed that ERM portfolios (or exposure to ERM firms) are suitable for pension funds because they are long-term assets that are correlated with longevity risk. However, those who make such claims often overlook the exposure of ERM portfolios to housing risk. We will have more to say on this issue in a later report.

value of the ERM loan to the lender is the value of the loan (which would include expected profits on the loan, and which is not to be confused with the loan amount) minus the value of the NNEG.

Suppose now that the NNEG value is very low. Then it follows that the value of the ERM to the lender will be close to the value of the loan component of the ERM, because the NNEG value is very low and so doesn't really matter. The ERM is then close to being a high interest loan that is certain to be repaid in full, which is an attractive proposition to the lender. Thus, ERMs are highly profitable to the lender *if* the NNEG value is very low. At the other extreme, if the NNEG value is very high, then an ERM loan might not be profitable at all. A correct valuation of the NNEG is then essential to determine the profitability of the loan.

Equity Release is a Rapidly Growing Sector

The growth of the equity release sector is apparent in the following Figure, which plots the amounts lent since 2000:

Figure 2.4: Growth of the Equity Release Sector

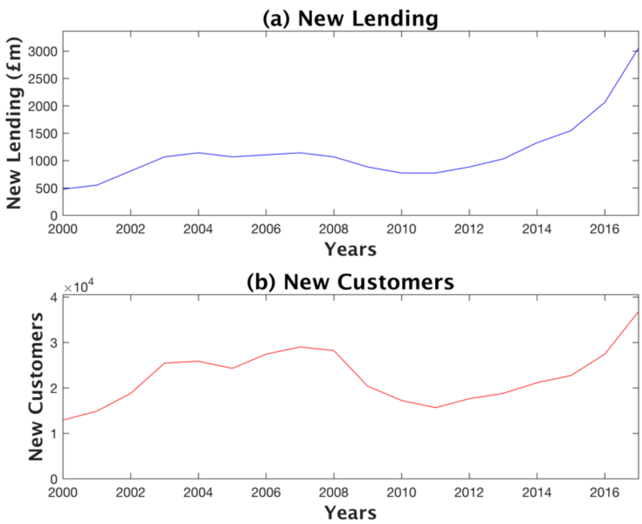

Source: Equity Release Council: http://www.equityreleasecouncil.com/document-library/equity-release-market-report-autumn-2018/equity-release-market-report-autumn-2018.pdf

The blue line plots the amounts lent and the red line plots the number of new customers. The latter should be treated with suspicion as it includes returning customers, i.e., those with drawdown facility who subsequently drew down.

Thinking from a macro perspective, the lending broadly reflects changes in UK house prices, i.e., the 2008 peak is clearly visible with subsequent decline and then recovery.

The final new lending amount is well above the previous peak in 2008. We could attribute the post-crisis growth to a number of factors including possibly stronger marketing and Treasury 'approval' of the product, or perhaps more naturally to increased demand for the product as boomers reach retirement age and realise that they can't afford a new washing machine but have a house worth £500k.

Recent growth in the sector has been remarkable:

> Since 2016, the equity release market has been growing at the average rate of 7.1 per cent each quarter. The amount of equity unlocked from people's homes more than doubled from £514m in the second quarter of 2016 to £1.08bn in the final quarter of 2018. [13]

The Equity Release Council's latest (Autumn 2018) *Market Report* provides further details:

> A total of 38,912 households aged 55 and over used equity release products from members of The Council to access some of their property wealth during H1 2018. This included 21,490 new plans agreed by customers, up by 28% from 16,805 a year earlier.
>
> A further 15,709 returning drawdown customers made withdrawals from their agreed reserve funds between January and June, up 25% year on year. …
>
> *The number of new plans agreed in [2018] H1 exceeded the entire size of the market in 2014* and represented an 81% increase since H1 2016. Activity among new customers has increased during every half-year period in the intervening two years, as equity release and housing wealth have taken up a mainstream position among the products and assets that form part of modern retirement planning.
>
> Two thirds of new customers (65%) opted for drawdown lifetime mortgages in H1, while 35% chose lump sum lifetime mortgages and a small number (<1%) agreed home reversion plans. (ERC 2018, p. 6, our italics)

The ECR report offers further detail about the growth of product choices:

> The growing base of equity release customers in recent years has been met with a greater number of product choices and flexibilities – helping to meet homeowners' increasingly complex needs in later life. As of August 2018, 139 product options were available to consumers, more than double the number (58) seen two years ago. The range of product options has grown by over 78% in the last year alone from 78 in August 2017.
>
> Today's equity release products also offer greater flexibilities thanks to ongoing competition and innovation in the sector. Four in five (80%) product options offer consumers the choice to make ad-hoc, penalty-free voluntary or partial repayments of their loan, up from 68% a year ago. There has also been

[13] L. Warwick-Ching "Dramatic Rise in Equity Release by Over-55s to Fund Retirement." *Financial Times* 4 April 2019.

> an increase in products offering fixed early repayment charges (ERCs), from 49% in August 2017 to 51% in August 2018.
>
> Almost half of product options (45%) offer downsizing protection which allows customers to downsize to a smaller property and repay the loan, without incurring any ERC. Inheritance protection, which allows customers to ring-fence a section of their housing wealth as a guaranteed minimum amount to pass on to the next generation, regardless of the total interest accrued – is offered by 46% of products. (p. 7)

And about customer profiles, e.g.:

> The average age of new customers during H1 2018 was the closest seen to date across the two main categories of lifetime mortgages. At just over 68 years old, the average new lump sum customer was broadly in line with that seen over the last three years. The average drawdown customer was almost two years older and just short of their 70th birthday. (p. 9)

So the sector is doing great, but there is still the issue about NNEG valuation.

Chapter Three: The Basics of NNEG and ERM Valuation

1. Introduction

This chapter explains the basics of NNEG and ERM valuation.

2. Exit Probabilities

The NNEG valuation model has two key ingredients: a set of expected house-exit probabilities and a put option pricing model.

The house exit probabilities (or exit probabilities for short) refer to the probabilities that the borrower will exit the house (and hence terminate the loan) over each of the next 1, 2, 3, … etc years. For the time being, let us assume away the possibility of early repayment of the loan and assume that the borrower is a single male who is expected to exit the house in a box. Under these conditions, the exit probability for year t is equal to

(3.1) $$exit\ prob_t = q_t \times S_t$$

where q_t is the mortality rate for year t and S_t is the probability that an individual alive now will survive to year t. Note that $S_0 = 1$ and $S_t = (1 - q_{t-1})S_{t-1}$ for all $t > 0$.[14]

The exit probabilities for a male just turned 70 are shown in Figure 3.1:

[14] The 'q' terminology for the mortality rate is standard and a little unfortunate in the NNEG context, where q is sometimes used to refer to the net rental rate or deferment rate. The reader should bear this ambiguity in mind, but the context should make it clear whether it is the mortality rate or the deferment rate that we are referring to, and it is mostly the latter.

Figure 3.1 Exit Probabilities for Males Currently Aged 70

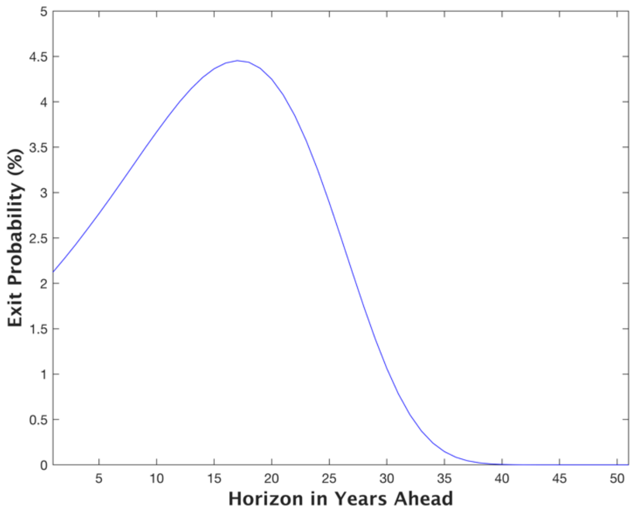

Notes: Exit probabilities are based on CBD-M5 model (Cairns *et alia*, 2006, 2009) cohort mortality rate projections using male England & Wales deaths rate data estimated over ages 55:89 and years 1971:2017. Source: Life & Longevity Markets Association.[15]

The left hand (low t) exit probabilities are close to the low t mortality rates and reflect the early high survival probabilities (i.e., that people aged 70 have a high probability of living at least a few years), and the later (high t) exit probabilities primarily reflect the fact that the probabilities of living to extreme old age are low and approach zero in the limit.[16]

3. Valuation Issues and the Put Pricing Model

The present value ERM of the Equity Release Mortgage loan can be considered to be the present value L of a risk-free loan, one which is guaranteed to be repaid in full, minus the present value $NNEG$ of the NNEG guarantee:

(3.2) $$ERM = L - NNEG$$

The original loan amount grows at the loan rate (sometimes called the rollup rate) l from its current amount until the time when the loan ends. Therefore L is given by

(3.3) $$L = \sum_t [exit\ prob_t \times current\ loan\ amount \times e^{(l-r)t}]$$

---

[15] See A. J. G. Cairns, D. Blake and K. Dowd (2006) "A Two-Factor Model for Stochastic Mortality with Parameter Uncertainty," *Journal of Risk and Insurance*, 73(4): 687-718; and A. J. G. Cairns, D. Blake, K. Dowd, G. D. Coughlan, D. Epstein, A. Ong, and I. Balevich (2009) "A Quantitative Comparison of Stochastic Mortality Models Using Data from England & Wales and the United States," *North American Actuarial Journal* Volume 13(1): 1-35.

[16] We have more to say on the derivation of the exit probabilities in Chapter 11.

where $exit\ prob_t$ is the probability of exiting the house in period t and r is the risk-free interest rate.[17]

The valuation of L is straightforward.

$NNEG$ is the sum of the products of the exit probabilities for each future time t and the present value of the NNEG guarantee for each future time t:

$$(3.4) \qquad NNEG = \sum_t [exit\ prob_t \times NNEG_t]$$

where $NNEG_t$ is the present value of the NNEG guarantee for period t.

The question is then how to value each of these individual $NNEG_t$ terms and thence the NNEG guarantee.

Recall that the NNEG gives the customer (or the person acting for the customer) the right to repay the loan by paying the lender the minimum of the loan value or the house price at the time of death.

The right to repay the minimum of two future values (one of which, the future house price, is uncertain) at some given future time implies a European put option granted by the lender to the borrower. Since the time of exercise is uncertain, we can think of the NNEG as involving a portfolio of such put options.[18]

In the case of our put options the underlying variable is a residential property ('house') or more precisely, a forward contract on a house, and we should think of a house as an asset that bears a continuous yield in the form of a (net) rental rate. This net rental yield reflects the use benefit of living in the house or the rental income we might get by renting the house out.

A natural option pricing model to use in these circumstances is the Black '76 model (Black, 1976). Black '76 is an appropriate pricing model when the underlying is a forward contract with a maturity coterminous with that of the option itself. This model is a near-relative of the famous Black-Scholes model (Black and Scholes, 1973; Merton, 1973).[19]

The Black '76 formula for the price p_t of a European put option with maturity t on a forward contract on a commodity bearing a continuous yield q is given by the formula:

[17] Note the implicit distinction here between the loan amount or rolled up loan amount, on the one hand, and L, the (economic) value of the loan, on the other. The former is the amount loaned plus the interest accumulated since the inception of the loan, whereas the latter is the value of the loan to the lender, including the expected profit on the loan. A concrete example of the distinction between the two is given in Table 3.1. Note too that the economic value of the loan is not to be confused with the accounting book value of the loan, which is another issue again.

[18] One might alternatively model the NNEG as a single American option with an early exercise feature but there would be no point in doing so. American options get interesting only when the option is exercised early in the self-interest of the option holder, but the decision to exercise early makes no sense in this context, because such a decision would be tantamount to the borrower taking their own life to get back at the lender, presumably burning down their house in the process.

[19] The choice of option model is discussed in more detail in the appendix to this chapter.

$$p_t = e^{-rt}[K_t N(-d_2) - F_t N(-d_1)] \tag{3.5}$$

where r is the risk-free rate of interest, K_t is the strike or exercise price for period t, F_t is the forward house price for period t, the function $N(\ldots)$ is the value of the cumulative standard normal distribution at the value specified in brackets, and d_1 and d_2 are given by:

$$d_1 = [ln(F_t/K_t) + \sigma^2 t/2]/(\sigma\sqrt{t}) \tag{3.6}$$

$$d_2 = d_1 - \sigma\sqrt{t} \tag{3.7}$$

where σ is the volatility of the forward house price.[20]

The strike price K_t is then the rolled up or accumulated loan amount by period t:

$$K_t = current\ loan\ amount\ \times e^{lt} \tag{3.8}$$

and the forward price F_t, the price agreed now to be paid on possession in period t, is:

$$F_t = current\ house\ price\ \times e^{(r-q)t} \tag{3.9}$$

where q is the *deferment rate*, namely the discount rate applied to the current house price to give the *deferment price*, the price we would agree to pay *today* to take possession of the house in t years' time. Thus the deferment house price R_t is given by:

$$R_t = current\ house\ price\ \times e^{-qt} \tag{3.10}$$

The difference between the forward house contract and the deferment house contract is that with the forward we settle when we take possession *in t years' time*, but with the deferment contract we settle *today*.[21] Therefore, the deferment house price R_t is the present value of the forward price, where the present value is obtained by discounting at the risk-free rate r.

It is important to note that the deferment house price will be less than the current house price S_0 because the deferment rate $q > 0$.

The forward house price F_t should not be confused with future house prices or expected future house prices:

- Forward prices for future period t are known (or can be approximated) now and we need to be able to value options using information available now.

---

[20] Compared to the original Black-Scholes equation (Black and Scholes, 1973), we replace the spot underlying, the current house price, with the forward house price. A point sometimes overlooked, the r term in the classic Black-Scholes formulas for d_1 and d_2 also drops out because the underlying contract is paid for at maturity and not at inception, and we assume that a rational seller would require compensation for growth on the sum of money they would have received if paid up front. This assumption should not be confused with the common misconception that the model assumes that the underlying grows at the risk-free rate.

[21] See PRA SS 3/17 (p. 12, note 2).

- Options cannot be valued using future house prices because future house prices are currently unknown.
- Options should not be based on expected future prices because expectations of future prices, e.g., future house prices, do not appear in the Black '76 option pricing formula.

We should also keep in mind that although the original Black '76 article discussed options on futures, futures prices are the prices of *futures* contracts, a form of forward contract, *not* actual or expected *future* prices of any sort.

A mistake to be particularly avoided – the one common among UK ERM actuaries – is to confuse forward and expected future prices. This mistake typically manifests itself in the inputting of an assumed expected house price inflation rate into (3.9) instead of the forward rate $r - q$.

To repeat: it is not the *future* or expected future value of a contract for *immediate* possession that we use in the option pricing equation, but rather the *current* value of a contract for *future* possession.

The approach set out here is an example of what is known in actuarial circles as a 'market consistent' approach, which gives 'market consistent' valuations. We would define a market consistent valuation as a 'fair value' valuation based on the IFRS definition of a fair value price, namely

> The price that would be received to sell an asset or paid to transfer a liability in an orderly transaction between market participants at the measurement date[22]

An alternative (and for our purposes practically equivalent) definition is that provided by Tim Gordon: he defines a market consistent valuation as one which is consistent with modern finance theory as the term is used in Exley, Mehta and Smith (1997).[23]

We shall take the validity of this approach as given for the time being, but we will provide a justification for it in later chapters.

4. A Valuation Example

We now build an ERM and NNEG valuation model based on plausible input parameter calibration values.

The baseline parameter inputs are:

[22] See, e.g., https://www.iasplus.com/en/standards/ifrs/ifrs13.

[23] See T. Gordon (1999) "The Price of Actuarial Values." Paper presented to the Staple Inn Actuarial Society (16 February) and C.J. Exley, S.J.B. Mehta and A. D. Smith (1997) "The Financial Theory of Defined Benefit Pension Schemes," *British Actuarial Journal* 3(4): 835-966.

- Current age of customer = 70, a typical age for ERMs.[24]
- Loan to value ratio = 40%[25]
- Risk-free rate r = 1.5%.
- ERM loan rate $l = 5.25\%$.
- Deferment rate $q = 4.2\%$.
- Volatility $\sigma = 14.8\%$ for males aged 70.

All rates are in % p.a.

We will discuss these calibrations in later chapters.

We assume an illustrative house price of £100 which, combined with the assumed loan to value ratio of 40%, implies a loan amount = £40.

The death/exit probabilities are derived from projections of future mortality rates obtained using the M5 version of the Cairns-Blake-Dowd mortality model (see Cairns *et alia*, 2006, 2009) calibrated on England & Wales male mortality data for the period 1971 to 2017 and spanning ages 55 to 89. The data are taken from the Life and Longevity Markets Association database (llma.org). The M5-CBD model is particularly suitable for old age projections and its goodness of fit and performance evaluation are assessed in Cairns *et alia* (2011) and Dowd *et alia* (2010a,b).

Our baseline NNEG valuation results are shown in Table 3.1:

Table 3.1: Baseline ERM/NNEG Valuations

Current House Price	***Loan Amount***	L	$NNEG$	ERM
£100	£40	£74.84	£32.19	£42.66

Notes: L is the present value of the loan component of the Equity Release Mortgage, $NNEG$ is the present value of the NNEG guarantee, and ERM is the present value of the Equity Release Mortgage. Based on the baseline assumptions: male aged 70, LTV=40%, r=1.5%, l=5.25%, q=4.2% and σ=14.8%. Exit probabilities are based on M5-CBD model projections using male England & Wales male deaths rate data spanning years 1971:2017 and ages 55:89.

Given the age of the customer, the expected present value L of the perfectly collateralised loan is £74.84. $NNEG$ is valued at £32.19 and so the value of the ERM, ERM, is equal to £74.84 – £32.19 = £42.66.

It is sometimes convenient to report these results in terms of the ratios of L, $NNEG$ and ERM to the loan amount as in Table 3.2:

Table 3.2: Valuations Divided by Loan Amount

L / ***Loan amount***	$NNEG$ / ***Loan amount***	ERM / ***Loan amount***
187.1%	80.5%	106.6%

Notes: As per Table 2.

[24] Implicitly, we are assuming a single male just turned 70. In the case of a single female, we would expect death/exit to occur somewhat later, which would increase the value of the NNEG. In the case of a couple, we would expect even later exit, when the longest living member of the couple exits the house.

[25] A 40% LTV ratio for a 70-year old appears to be approximately in line with current industry practice for new ERM loans. We will have more to say on this subject in the next chapter.

We see that ERM, for example, is 106.6% of the loan amount.

5. Sensitivities of Valuations to Key Input Parameters

Table 3 shows the sensitivities of L, $NNEG$ and ERM to changes in key parameter inputs. These are expressed in elasticity form, i.e., where the elasticity of the relevant output with respect to a change in an input is the % change in the output divided by the % change in the input.

Table 3.3: Sensitivities of Valuations in Elasticity Form

Elasticity wrt	$\boldsymbol{L}$	$\boldsymbol{NNEG}$	$\boldsymbol{ERM}$
r	-0.27	-0.54	-0.07
l	0.94	1.89	0.23
q	0	0.57	-0.43
σ	0	0.25	-0.19
LTV	1	1.74	0.44

Notes: L is the present value of the loan component of the Equity Release Mortgage, $NNEG$ is the present value of the NNEG guarantee, and ERM is the present value of the Equity Release Mortgage. Based on the baseline assumptions: male aged 70, LTV=40%, r=1.5%, l=5.25%, q=4.2% and σ=14.8%. Exit probabilities are based on M5-CBD model projections using male England & Wales male deaths rate data spanning years 1971:2017 and ages 55:89.

These results indicate that NNEG valuations are highly sensitive to changes in the l and LTV input parameter calibrations. It is also interesting to note that the ERM valuations are much less so, because of the offsetting impacts on the loan value and NNEG valuations.

Appendix to Chapter Three: The Choice of Option Pricing Model

In Chapter Three we suggest that Black ’76 is an appropriate put pricing model when the underlying is a forward contract with a maturity coterminous with that of the option itself.

Why not use Black-Scholes (BS) instead?

A shallow but correct response is that BS is not appropriate because it is based on the assumption that the underlying does not bear any yield, whereas Black ’76 is appropriate because it allows for such a yield. The yield in this context would be the net rental rate.

However, this deficiency of BS is easily fixed. It is well-known that BS can be adapted for an underlying that pays a continuous dividend yield q if we replace S in the BS formula with Se^{-qt}.[26]

If we make this adaptation, then we can use BS because the underlying can also be interpreted as a spot asset that pays a continuous dividend yield. In this case, BS and Black ’76 will be equivalent.[27]

To demonstrate this equivalence, start with the relationship of the forward to the spot:

(3A.1) $$F = Se^{(r-q)t}$$

(3A.1) implies

(3A.2) $$S = Fe^{(q-r)t}$$

Now consider the Black-Scholes put price:

(3A.3) $$p = e^{-rt}K \times N(-d_2) - e^{-qt}S \times N(-d_1)$$

where

(3A.4) $$d_1 = [\ln(S/K) + (r - q + \sigma^2/2)t]/[\sigma\sqrt{t}]$$
(3A.5) $$d_2 = d_1 - \sigma\sqrt{t}$$

Substitute $Fe^{(q-r)t}$ for S to obtain:

(3A.6) $$p = e^{-rt}K \times N(-d_2) - e^{-qt}F \times e^{(q-r)t}N(-d_1) =$$

---

26 See, e.g., J.C. Hull (2003) *Futures, Options and Other Derivatives*, fifth edition, p. 268.

27 We could also price our puts by tweaking the Garman-Kohlhagen foreign currency option model (Garman and Kohlhagen, 1983) or by using an appropriate special case of the Margrabe option, the option to exchange one asset for another (Margrabe, 1978), but these approaches are also equivalent to Black ‘ 76. See M. B. Garman and S. W. Kohlhagen (1983) “Foreign Exchange Currency Options,” *Journal of International Money and Finance* 2(3): 231-237; and W. Margrabe (1978) “The Value of an Option to Exchange One Asset for Another,” *Journal of Finance* 33(1): 177-186.

$$e^{-rt}[K \times N(-d_2) - F \times N(-d_1)]$$

(3A.7) $$d_1 = [\ln (Fe^{(q-r)t}/K) + (r - q + \sigma^2/2)t]/[\sigma\sqrt{t}] = [\ln (F/K) + (\sigma^2/2)t]/[\sigma\sqrt{t}]$$

(3A.8) $$d_2 = d_1 - \sigma\sqrt{t}$$

which is the Black 76 model!

The argument has been put to us that since the forward is more volatile than the spot (on account of the volatility of r and q impacting the forward rate), we should prefer Black Scholes, which is priced off the volatility of the spot. The Black Scholes price will then be cheaper, and the NNEG will be less expensive.

This objection is incorrect. Let's assume that by BS we are referring to BS with the adjustment for a continuous dividend yield, i.e., BS with Se^{-qt} as the underlying rather than BS with S as the underlying.[28] In that case, our response would be that this objection must be incorrect because both models are equivalent. Consider also the fundamental mechanics of option pricing. The option price reflects the payoff at expiry less the cost of hedging, and the cost of hedging results from the constant rebalancing driven by changes in delta, but the option deltas are the same and driven by the d_1 term:

(3A.9) BS delta = B76 delta = $-e^{-rt}N(-d_1)$

For both models, the d_1 term is driven both by changes in the asset price, *and by changes in r and q*. In the Black Scholes model this impact is explicit, whereas in Black 76 this impact is implicit in the definition of F.

[28] As discussed already, if we are *not* making this assumption, then we are dealing with a version of BS that is not appropriate.

Chapter Four: Loan-to-Value Ratios

The Loan-to-Value ratio is the ratio of the loan amount, or the original loan amount plus rolled up interest, to the price of the property that has been mortgaged.

The Relationship of LTV to Age

ERM lenders have to decide how much to loan against the mortgaged property, and a key factor to be considered is the age of the borrower. Given that the loan rate is likely to exceed any future house price inflation (hpi) rate, then the longer the loan lasts, the greater the likelihood that the loan will go into negative equity. Therefore lenders will offer lower loans to younger borrowers.

An example of a rule governing this age-dependency is the 'age – 40' rule set out by Hosty *et alia* (2007, p. 31) and used in their NNEG analysis: "Maximum initial loan to value ratio (MLTV) 15% at age 55 increasing by 1% for each year of age to 50% at age 90 (younger age for joint life cases)."

One forms the impression from Hosty *et alia* that this rule was in use in the period before they wrote. (After all, if it was not, then why would they cite it?) If this view is correct, then it would appear that LTVs/age ratios have risen since then and that lenders now routinely offer larger LTVs than they used to, to lenders of the same age.

Our evidence comes from two sources. The first is from a recent speech by David Rule, the PRA's Executive Director of Insurance Supervision (Rule, 2018). Mr. Rule makes the following statement:

> Chart 5 plots loan-to-value ratios against age for equity release mortgages *sold by life insurers in 2017*. The pink swathe shows insurers' risk limits. The great majority of mortgages are within risk appetites, but some loans exceed the limits. This may or may not be a problem. For example, it could be explained by medical underwriting of the mortgage that justifies a higher loan-to-value ratio for a younger customer if their life expectancy is impaired. (Our emphasis)

The Chart 5 he refers to is reproduced below:

Chart 5

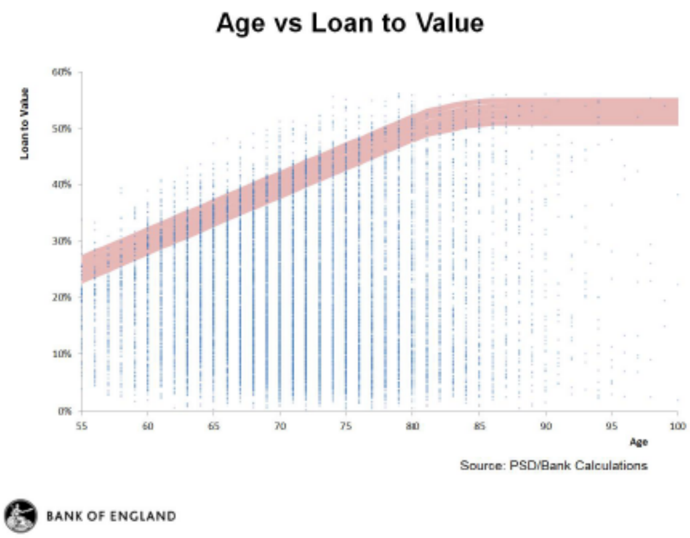

He doesn't define "risk limits" but we might interpret these limits as those suggested by a good practice rules i.e., a recommended ltv/age rule with some discretion around it. The key phrase is here is "sold by life insurers in 2017," so these are newly minted ERMs and not historical ones

The relatively small number of points above the bounds seem to us to be those with impaired lives, who would be offered better LTV terms than those in normal health. The large number below would seem to us to reflect the impact of un-drawdown balances in drawdown ERMs.

The pink risk bounds suggest to us that firms are using something close to a 'modified age – 30' rule that goes as follows:

- For ages 55 to 80: LTV = age/100 – 30.
- For ages 81 to 85: ΔLTV = Δage/100.
- For age 86+: ΔLTV=0.

We have also seen examples from current industry practice that are consistent with this 'modified age – 30' rule and lead us to believe that this rule (or something fairly close to it) is not a bad approximation to the LTVs in current fresh minted ERMs. This point made, there can be considerable variation in LTVs across products (see, e.g., Legal and General (2019) or Tunaru (2019, p. 63, Table 12)). [29]

The Impact of Rollup and Past House Price Growth on Historical LTVs

[29] Legal and General "Flexible Lifetime Mortgage" (2019): https://www.legalandgeneral.com/adviser/files/retirement/literature-and-forms/product-summary/Q0052431.pdf

One must also consider that after an ERM loan is made, LTVs will evolve afterwards in line with the loan rate and subsequently realised house price inflation. If $LTV(0)$ is the initial LTV, after 1 year the LTV will be

(4.1) $$LTV = LTV(0) \times e^{l-hpi(1)}$$

where l = (assumed constant) loan rate and $hpi(1)$ = house price inflation after the first year. After 2 years, the LTV will be

(4.2) $$LTV = LTV(0) \times e^{l-hpi(1)} \times e^{l-hpi(2)}$$

and so forth.

We can form some idea of how the LTVs from loans made in the past would have moved by a simple historical simulation. Let's assume that ERM loans were made in the past in each of the years from 2000 to 2017, and let's assume that these were made on the basis of a 30% LTV and a loan rate of 7%, noting that these calibrations would be reasonable for this period. The next figure shows how the LTVs would have subsequently 'performed'.

Figure 4.1: Historically Simulated LTVs: 2000 to 2017

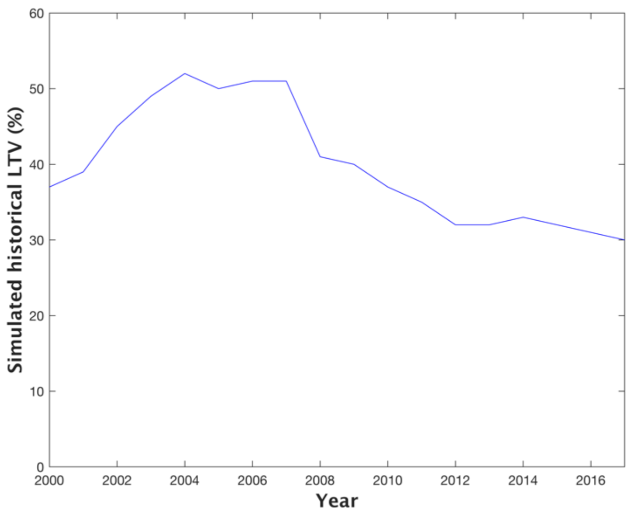

Notes: Based on loan rate = 6% and LTV = 40% when the loan is taken out.

To interpret the figure: an ERM with a 30% LTV taken out in 2000 would now have an LTV of about 39% if the loan were still active in 2017; a 30%-LTV loan taken out shortly before the crisis would have an LTV of about 51% if the loan were still active; and (naturally!) a 30%-LTV loan taken out in 2017 would still have an LTV of 30% in 2017. We see that the loans most at risk (in LTV terms) were those taken out shortly before the crisis.

The driving factor behind these LTV movements was the hpi: the stronger the hpi, the lower the growth of the LTV.

We should however bear in mind that past performance is no guarantee of future performance.

Future LTVs Depend on Future House Price Growth

For any given LTV, its future path will depend on future house prices. Remember too that if the LTV exceeds 100% when exit occurs, then the NNEG will be triggered and the lender will lose the difference between the rolled-up loan amount and the house price at that time.

Figure 4.2: Future House Price Growth and the Path of LTVs

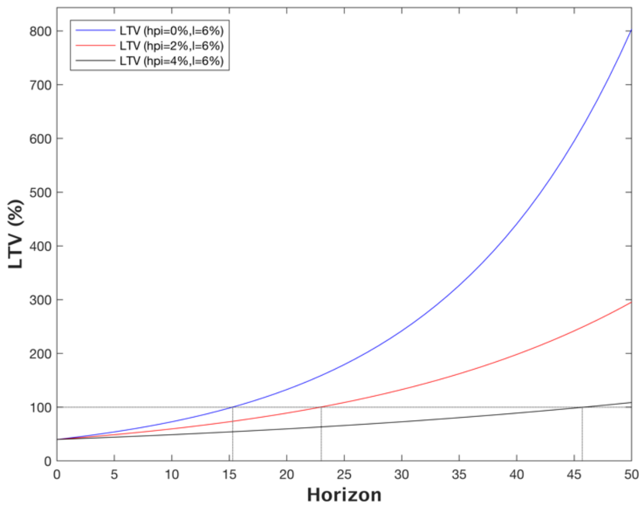

Notes: Based on loan rate = 6% and LTV = 40% when the loan is taken out.

The rapid rates of growth of the top two LTVs plots (i.e., those for hpi = 0% in blue and for hpi = 2% in red) are striking. In the first case, LTV hits 100% in just over 15 years and the second case it hits 100% in 23 years. In the third (hpi = 4%, black) case, LTV hits 100% in 46 years.

The message is that lenders need a high future hpi rate to keep their potential NNEG losses down. For example, if house prices remain flat and the borrower is 70 when he takes out his loan, then there is a good chance that he will live the 15 or more years required to trigger the NNEG. By contrast, if house prices grow at 4% p.a., then the same borrower is vanishingly unlikely to live the 46 years needed to trigger the NNEG. Consequently, an ERM loan is highly exposed to house price risk.

The exposure of ERMs to house price risk is further illustrated by a simple stress test. Suppose that the same borrower takes out the same ERM loan, and then the next day house prices fall 40%. The LTV is 40% on the day the loan is taken out, but the 40% house price fall implies that the LTV jumps to (1/(1-0.4))*40% = 66.7% the next day! If house prices then grow at zero percent, the LTV will hit 100% in 8 years; if they grow at 2% the LTV will hit 100% in 12 years and if they rise at 4%, the LTV will hit 100% in 22 years.

The chances of such an outcome are presumably small, but the point is to illustrate the risk exposure. A more likely adverse outcome would be a Japan type scenario in which house prices do not fall suddenly, but instead fall secularly over a long time. Under this Japan scenario, LTV would hit 100% in just 13 years.

Chapter Five: The Risk-Free Rate

Figure 5.1shows a plot of recent Bank of England spot rates:

Figure 5.1: Bank of England Spot Rates

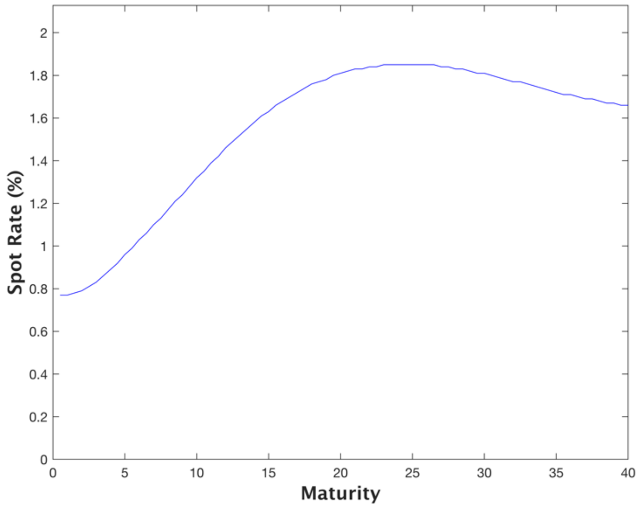

Source: Bank of England UK nominal spot rate curve, downloaded 6 March 2019.

If one wished to use a single 'representative' spot rate, then our earlier baseline choice of 1.5% is not unreasonable.

However, in a more sophisticated analysis there is no good reason not to use the whole spot rate curve, with spot rates for maturities greater than 40 years set, e.g., as equal to the 40-year spot rate.

Chapter Six: The Loan Rate

The loan rate or rollup rate is the rate that the lender charges on the ERM loan. This rate will be fixed for the life of the loan, but the rate applied to drawdown ERM loans will typically be floating, in the sense that the rate applied to a drawdown typically varies depending on when the drawdowns are made on a loan.

In his NNEG report, Tunaru (2019, p. 29) provides a nice plot of the evolution of ERM rates since 1999:

Figure 6.1: Evolution of ERM rates

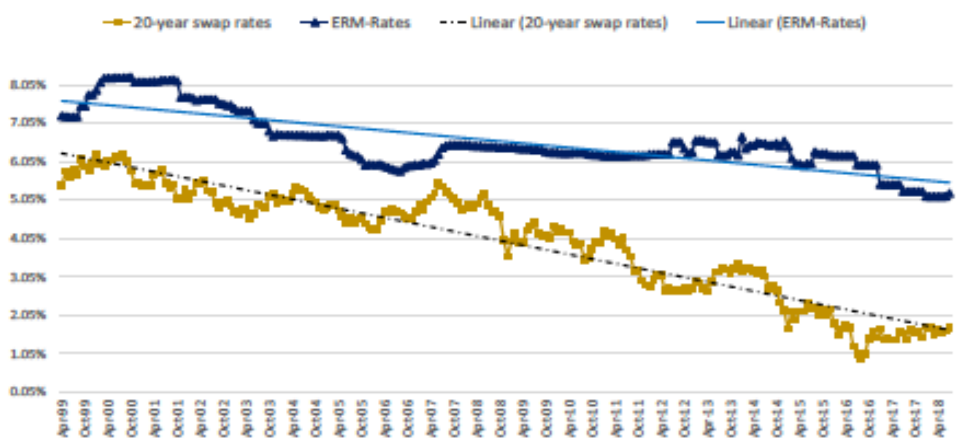

Figure 8: Average of top ERM customer rates against 20 year swap rates: Source Hosty et al. (2008) until 2006, Bloomberg and other combined by us using monthly interpolation.

Notes: Tunaru (2019, Figure 8).

This plot shows that loan rates were about 7% in 1999, then rose shortly thereafter to just over 8%, before trending downwards to about 5% in 2018.

In their Autumn 2018 market report, the Equity Release Council report (p. 7) that as of July 2018, the average loan rate across products was 5.22% and this value is consistent with the latest value in the Tunaru plot.

How much variation there is across products is difficult to say, but we have seen current rates that vary from 4.15% AER to 6.78% AER.

The latest available loan rates quoted by the ERC and shown in the Tunaru plot suggest that a baseline loan rate of 5.25% might be appropriate.

Chapter Seven: Net Rental Rate and Deferment Rates: Theory

Net Rental Rate vs Deferment Rate

Let's start by some clarifying definitions. The net rental rate is the rate that the landlord receives after deducting for void, management costs and maintenance costs. The deferment rate is the discount rate applied in the deferment price formula, reflecting the foregone income or use during the deferment period.

These two rates are frequently run together, and as we shall show later, they are in fact mathematically identical, but they are defined differently. In this chapter we shall show how to derive one from the other, based on the definitions given.

A First Principles Analysis

A more precise definition of the deferment rate is the discount rate that when applied to the freehold price of vacant possession results in the price of deferred possession. The deferment rate itself is not directly market observable, but it can be estimated as a function of market variables. The method proposed here uses net rental yields, as follows.

Let d be the *current* net nominal annual rental, the current time being the *beginning* of the year. (We use 'd' here because the approach we are using derives from the dividend discount model, where 'd' is used to refer to (nominal) dividends.) 'Net' means the gross or headline rental paid by tenants, less the costs incurred by the lessor such as management, maintenance and the expected costs of void or empty periods while the property is being re-let. Then we shall show that

(7.1) $$d/S = q$$

where d is the net rental as above, with the current time being the beginning of the rental year; S the estimated 'spot price', i.e., the freehold value of vacant possession estimated as the market value of an identical or similar property not encumbered by a leasehold; and q the deferment rate as defined above.

Then assuming that the value S of the property is the present value of all net rental receipts, which is what a market participant would reasonably assume, the following equation holds:

(7.2) $$S = d \times y \times (1 + y + y^2 + y^3 \dots) = d \times (y + y^2 + y^3 \dots) = d \times y/(1 - y)$$

where $y = (1 + g)/(1 + r + \pi)$, r is the risk free rate, π is the risk premium required by investors in residential property, and g is growth of net *income* (e.g., dividends or net rental, *not* property price).[30] So for any rental cashflow period we project the year-end net rental by the expected growth rate g, then discount back by the investor required cost of equity $r + \pi$.

Rearranging, it follows that

(7.3) $$d/S = (1 - y)/y$$

We now want the deferment value R_n, i.e., the value of the property S minus the lost net rental for n periods:[31]

(7.4) $$R_n = d \times y \times (y^n + y^{n+1} \dots)$$

Assume that there is no term structure to cost of equity or growth. (The effect of term structure will be discussed below.) The following equation is then true:

(7.5) $$R_n = d \times y \times (y^n + y^{n+1} \dots) = d \times y \times (1 + y + y^2 \dots) \times y^n = Sy^n$$

i.e.,

(7.6) $$R_n = Sy^n$$

Define q as the discretely compounded discount rate applied to spot S such that:

(7.7) $$R_n = Sy^n = S(1 + q)^{-n}$$

Hence

(7.8) $$y = 1/(1 + q).$$

From (7.3), substituting $1/(1 + q)$ for y:

(7.9) $$d/S = (1 - y)/y = (1 + q)(1 - 1/(1 + q)) = (1 + q) - (1 + q)/(1 + q) = q$$

and hence

(7.10) $$d/S = q$$

which was to be proved. Therefore we can estimate the deferment rate q using observable values for the value of vacant possession, S, and a value of the previous net rental amount $d0$.

Observe that (7.10) holds true whatever growth rate we choose, and whatever interest rate and risk premium are required by investors.

---

[30] (7.2) follows from applying the discount dividend model (e.g., Gordon, 1959) with property prices and rentals taking the place of stock prices and dividends. See, e.g., https://en.wikipedia.org/wiki/Dividend_discount_model.

[31] I.e. where lost rental $= d \times (y^0 + y^1 + \cdots + y^{n-1})$

Which is really strange, too. Tunaru (2019, p. 50) says "For risk-management calculation purposes then, it is very important to have an accurate measurement of q. Lack of data availability and long-term horizon makes this exercise extremely difficult, if not practically impossible." This claim seems plausible. Who could predict unobservables such as dividend growth, risk premia and so on? Yet we can calculate q by a simple formula using observable variables!

Term Structure of Growth

In the derivation of equation (7.1) it was assumed that there is no term structure to the cost of equity $(r + P)$ or to growth (g). Assuming a term structure would make a difference to the deferment rate. Decreasing the discount rate for early periods (say the first 20 years) or increasing the growth rate for the same period would have the effect of increasing q, because it would make the value of lost income higher as a proportion of the deferment value.

However, a pronounced term structure to either risk premium or growth seems unlikely. A forward rental rate is the rate one would pay to lease the property with a forward starting date for a certain period. But why would the market imagine that the forward rate between years 39-40 is significantly different from that between years 40-41, for example? It is difficult to see what information would justify such a jump, and there must be a presumption that we shouldn't introduce additional complicating factors without good reason.

Term Structure of Deferment Rate

Observations of deferment rates using leasehold prices show a term structure (of which more in the next chapter), with the deferment rate for short leases higher than for long leases. This effect arises because the value of a short-term lease approaches that of a short-term rental, and the short-term rental reflects the gross, rather than the net rental yield. As the leasehold term increases, its value will approach that implied by the net rental yield, which can be shown as follows.

Let q_0 be the short term (i.e., annual) rental rate *gross* of annual costs c such as void rate, maintenance, share of management etc. Assume the following (crude) model:

(7.11) $$q_0 = (1 - c)A = q_t(1 - c/t)A$$

where $A = (y + y^2 + y^3 \dots)$ and q_t is the gross effective rental rate over a let of t years. Then

(7.12) $$q_t = q_0(1 - c)/(1 - c/t)$$

Thus the gross rental converges to the net rental as the leasehold term increases.

Figure 7.1 gives an illustrative plot of the gross rental vs leasehold term.

Figure 7.1: Gross Rental vs Leasehold Term

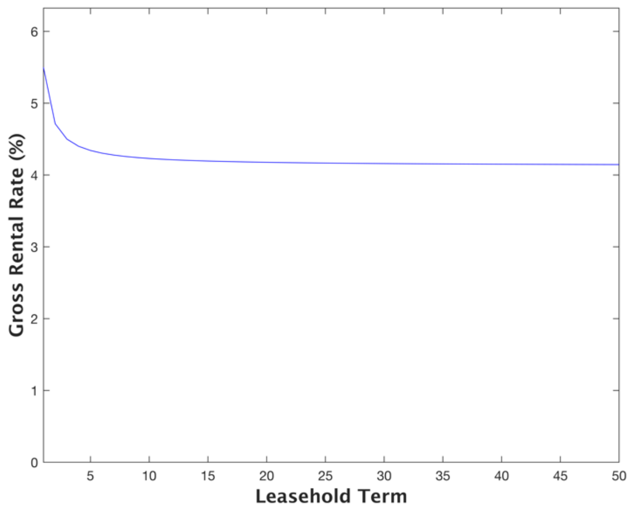

Notes: based on illustrate values of initial gross rental = 5.5% and c (costs) = 25%.

For the chosen calibration (q_0 = 5.5% and c = 25%), the gross rental goes from initial value, 5.5% towards a long-run value of 4.1%.

The Deferment Rate and the Real Risk-Free Rate

In Consultation Paper CP 7/19, the PRA proposed "to take account of movements in real risk-free rates when setting the deferment rate,"' in order to prevent variability in the real risk-free rate causing variability in the forward rate:

> The PRA would increase (reduce) the deferment rate if the review shows there has been a material increase (reduction) in long-term real risk-free interest rates since the last update.[32]

In our derivation of equation (7.10) above, however, we showed that the deferment rate is equal to the current net rental yield, i.e. the nominal net rental payment divided by the current nominal house price. The real risk-free rate does not even enter into it!

So what is going on here? Well it would appear that the PRA proposal implicitly depends on some assumed relationship or equivalence between the deferment rate and the real rate of interest, but no further details are given.

Digging deeper, it turns out that there *is* an equivalence, but only under conditions that do not hold. We start with the dividend discount equation:

(7.13) $$q + g = r + \pi$$

[32] CP 7/19 S2.4

where q is the deferment rate, g the growth of nominal net rental, r the risk free rate and π the risk premium. Now assume first that the property investment is risk free, i.e. that there is no risk premium π. Then

(A1) $$\pi = 0$$

Substituting into (7.13) we obtain

(7.14) $$q + g = r$$

Second, assume that g, the imputed growth in net rental, is equal to the general inflation rate i:

(A2) $$g = i$$

Hence

(7.15) $$q + i = r$$

Finally, assume (as stated in CP 7/19 para 2.5 that the nominal rate r is the sum of the expected general inflation rate and the real rate rr, a relationship known as the Fisher Effect:

(A3) $$r = i + rr$$

and where

(A4) $$E[i] = i$$

i.e., we assume that expected and actual inflation are the same. Substituting and subtracting both sides of (A3):

(7.16) $$q + i = rr + i$$

Hence

(7.17) $$q = rr$$

Given those four assumptions (A1-A4), it follows that the deferment rate and the real interest rate are identical.

We would question those assumptions, however. First, a property portfolio is clearly not risk free. An investment in a housing portfolio is commensurate to an investment in a *risky* index-linked bond, as opposed to an investment in an index-linked gilt, which is virtually risk free. This consideration suggests $\pi > 0$, so that the risky deferment rate will be higher than the real risk-free rate. Moreover π may vary through time.

Second, while rental inflation and general inflation are likely to be correlated, they are not the same. Nominal rentals will tend to go up line with inflation, indeed rental costs are part of the UK Consumer Price Index,[33] but as Figure 7.2 below shows, rental inflation and CPI are far from 100% correlated. There are long periods, such as 1991-2004, when rental inflation is consistently higher than CPI.

Figure 7.2: UK Consumer Price Inflation versus UK Rental Inflation, 1970-2017

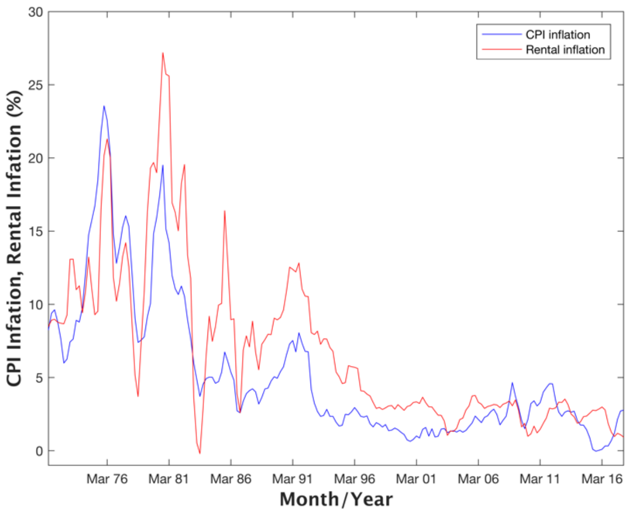

Source: OECD

Third, assumption (A3) above depends on the unobservable quantity, the *expected future* rate of inflation. This variable, as the Bank of England must know, is difficult to predict with any certainty, and hence is difficult to monitor.[34] By contrast, the net rental yield, which we proved above to be mathematically identical with the deferment rate, is relatively simple to observe.

Hence, if the PRA wants to monitor the deferment rate, it should monitor developments in the net rental yield.[35]

---

[33] Historical data from ONS 1988 - 2004 can be found here.

[34] It is possible that the PRA intends to monitor market expected real interest using the return on index-linked gilts. However, returns on index-linked gilts have been negative in the 2010s, whereas it is impossible for the deferment rate to be negative. The CP also notes (S2.4) that the deferment rate *will* always remain positive, in order to comply with Principle III of SS 3/17, but give no rationale of why this should be so.

[35] Note that the mathematical equivalence of the deferment rate and net rental yield also depends on the dividend discount model, but without additional assumptions like (A1)-(A4) above.

Chapter Eight: Net Rental Rate and Deferment Rates: Calibration

We can decompose the deferment rate q as follows:

(8.1) $$q = y - v - c - m$$

where y is the gross rental yield or the yield paid by the tenant, v is the void rate, c is management cost and m is the maintenance cost.

Define the maintenance cost m as the rate of expenditure (as a percentage of gross rental) required to keep the property in perfect condition (i.e. such as to achieve the best sales price for a property of that size in the same area), and define the tenant maintenance share (s) as the proportion of m that the tenant is likely to spend on maintenance. s will typically vary between 0 and 100%. For a short rental, s will be close to zero, and for a long let we would expect s to be close to 100% in the early years of tenancy, falling over time. In the final years it might fall to zero, even for a standard tenancy, given the lack of incentive to keep in full order for the landlord's benefit.[36] For an ERM, it would seem unlikely that the 'tenant' at end of life, perhaps in the situation where the NNEG had bitten, would have any incentive to keep the property in good condition, so we would expect s to fall towards zero in that case too.

We now use the following calibrations:

- Void: we use the standard '1 month in 12' rule of thumb, i.e., $v = (1/12) \times y$.[37]
- Management cost: following Tunaru (2019, p. 32), we assume management cost $c = 10\% \times y$.
- Maintenance cost: again following Tunaru (2019, p. 32), we assume maintenance costs $m = 15\% \times y$.

Thus, the maintenance cost borne by the landlord and to be subtracted from the gross rental yield is $m = 15\% \times 50\% \times y = 7.5\% \times y$.

We then have

(8.2) $$q = y \times (11/12 - 0.1 - 0.075) = y \times 0.7417.$$

Thus, the net is 74.17% of the gross.

Again following Tunaru (2019, p. 31), we take

[36] In fact, we can also imagine $s < 0$. So if $s = 0$ reflects no active effort to keep the property in condition, $s < 0$ reflects a determined effort by occupiers to strip the property (e.g., of light fittings, marble fireplaces, etc.) or trash the property! It happens.

[37] An alternative is to use empirical void data. Average void period for landlords in private rented sector in the United Kingdom (UK) have varied from 2.4 weeks to 2.9 weeks. (Source: https://www.statista.com/statistics/421102/rental-properties-void-periods-in-the-uk/. Accessed 19 March 2019.) If we take the mid-point, 2.65 weeks, then the average void rate by this measure would be 2.65/52 = 5.1%, as compared to the 'rule of thumb' void rate of 11/12 = 8.3%.

(8.4) $$y = 5.6\%.$$

Therefore

(8.5) $$q = 74.17\% \times 5.6\% = 4.15\% \ \approx 4.2\%.$$

So we use $q = 4.2\%$ as our 'best estimate'.[38]

[38] If we use the empirical void rate of 5.1%, then net is 77.4%% of gross and we would obtain $q = 4.3\%$.

Chapter Nine: Dilapidation

It is a law of nature that all properties depreciate over time, even the Great Pyramid of Giza. We stop dilapidation by maintaining the property. The cost of maintenance is thus a reflection of this tendency to dilapidate: if there were no dilapidation, there would be no cost of maintenance. But whereas dilapidation is a natural process, the decision to maintain is a choice. We can choose to maintain our property or we can let it dilapidate.

The ERMonomics of Dilapidation (I): Borrower Moral Hazard

It is usually the case that the owner has a long-term interest in the upkeep of their property, and it is then reasonable to presume that he or she would make the investments in the property that are necessary to counter the dilapidation that would otherwise occur.

Now consider equity release. Naturally, the lender would prefer that the borrower look after the property as much as they would if they still had a long-term vested interest in it, but the borrower no longer has any such interest. In any case, the borrower would be cash poor (because why else would he or she have taken out an ERM loan in the first place?) and as he/she ages, it would become more difficult to maintain the property even if the borrower wished to do so. There is a classic *moral hazard problem*.

Of course, the lender can impose conditions about maintenance, but has limited means of monitoring maintenance or enforcing them. So any such covenants would have limited effectiveness and there is still a moral hazard problem.

The likelihood then is that the borrower might make some investments in the property early on (e.g., a new kitchen or patio) but as they age, any such investments will be based on short-term considerations or become increasingly cosmetic (e.g., a paint job) or absolutely necessary to be able to continue living in the house (e.g., paying to fix a water leak). By the time the lender takes possession of the property, it will likely have depreciated in value, relative to the value one must assume it would otherwise have had, and certainly relative to the value it would have had had it been owned by someone younger, with the means, ability and incentive to maintain it properly. Consequently, an ERMed property will likely depreciate in value over the lifetime of the ERM, relative to the value of an otherwise similar property that had not been ERMed.

This dilapidation effect is reasonably well known, but there is no consensus on how to deal with it. To quote PRA CP 48/16:

> Opinion on the appropriate adjustment [needed for it] was divided, suggestions included adjustments to property value, property volatility, HPI, and *a margin for dilapidation*. Some felt that systematic underperformance risk due to adverse selection should be allowed for in the valuation. (pp. 25-26)

The ERMonomics of Dilapidation (II): Stochastic Dilapidation and Lender Moral Hazard

There are also other factors involved that are not so well known.

Consider the following figure.[39]

Figure 9.1: NNEG Claims: Aviva ERF4 Dataset

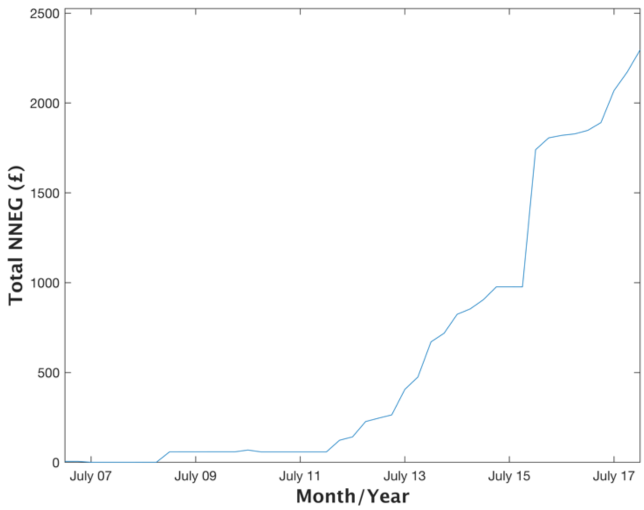

Source: Aviva: https://www.erfunding.co.uk/literature-library/erf4.html

The chart shows cumulative NNEG claims on ERF4, a set of Aviva ERMs which starts in 2004. Most ERMs start with a loan to value of lower than 50%, and property prices have gone up in most areas of the UK since 2004, so it takes a few years for the compound loan amount to reach current property value and reach NNEG territory. So low NNEG claims in the early years are to be expected.

But there is a surprise in store.

The exercise of the NNEG was in no case due to the compound loan value hitting the house price index. Aviva helpfully provide (i) and 'indexed house value' at exit, together with (ii) the original valuation and (iii) the realised sale value. The indexed value is the original value projected by the increase in the Halifax index since the valuation. 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 when the loans were repaid. Instead, 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 we estimate about 7% based on

[39] We thank T. Pocock for pointing out this Aviva ERM dataset: https://www.erfunding.co.uk/literature-library/erf4.html.

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%*). So it was also with 44% of the properties where the NNEG was exercised: 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.[40]

What is happening here?

Well, we know that as people get older, they are less inclined and less able to maintain their property to the standard they maintained earlier. We also know that they have less incentive to, because they have mortgaged it away. We discussed these issues already. Even so, the magnitude of the effect is surprising. It seems that the exercise of the put option appears to be due to the underperformance, often a dramatic underperformance, of the properties used as collateral. Even more surprising is that from the perspective of the lender, the underperformance seems to have a considerable random element, i.e., we are dealing with stochastic dilapidation![41]

The next Figure shows scatterplot of the 'achieved' values of ERMed properties, i.e., the amounts that the lender was able to realise after the borrower exited, expressed as a percentage of the indexed value, based on the Shared Appreciation Mortgage Securities (SAMS) originated by HBOS in the late 1990s:

[40] One of our correspondents who understands these data much better than we do suggests that this example is "quite clearly a fraud case." He backs up this claim with a persuasive analysis and concludes that the only non-fraud explanation he can think of is if the site had become contaminated by radioactivity. We defer to his judgment. Needless to say, firms should then either be modelling the possible impact of fraud or they should put suitable mitigating controls in place.

[41] See also CP 7/19 p. 8: "A key point to draw out is that lenders of equity release mortgages are exposed to the risk of individual property prices. This is because insurers provide a no-negative-equity guarantee to every borrower. Modelling approaches focused on house price indices do not capture all the risks – a portfolio of options is a very different thing to an option on a diversified index. Indeed UK insurers have experienced a number of these guarantees crystallising in recent years despite the rapid rise in UK house price indices over the past decades. One reason is that different localities of the United Kingdom have seen widely varying house price inflation – a national index masks the range of outcomes … . Another is that some properties may become dilapidated if elderly borrowers are unable to maintain the property. Willingness to maintain may be lower where borrowers have limited or no equity remaining in their properties. Equity release contracts generally require properties to be maintained. But, in practice, losses do occur and cannot necessarily be recovered."

Figure 9.2: Indexed vs. Achieved House Prices

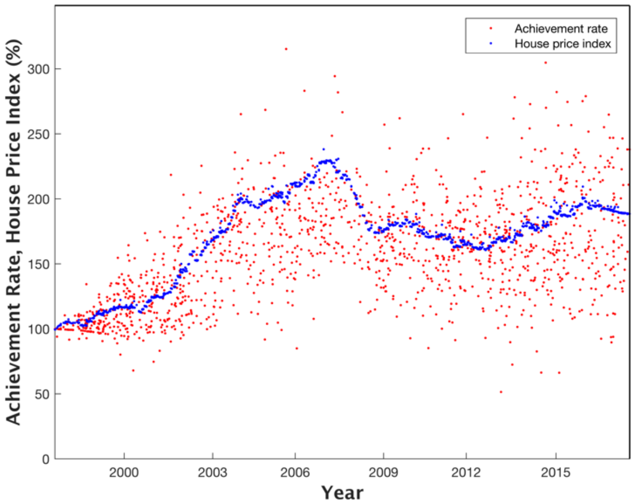

Source: SAMS. The achievement rates are real data; the indexed value is the original value projected by the increase in the Halifax index since the valuation.

The red dots are the scattershot of the individual achievement rates in the sample. The darker blue random-looking line is a simulated house price index.

What jumps out is that the achieved values are all over the place relative to the index. On average the achieved value is about 94% of the index but there is a huge dispersion around the index. The volatility of the index does not capture the volatility of the dispersion around the index!

An insightful explanation is provided by one of our EUMAEUS correspondents. To paraphrase:

> The achievement rates are like betas in stocks. Every house is individual. But the dispersion of results around the average and the volatility of that dispersion are both enormous. The impact of older people allowing their properties to subside in later life is the first reaction of most people, but is in fact only a small part of what is going on.
>
> Now practical experience suggests that the difference in price between the cheapest house on a street compared to the most expensive house on a street – assuming they're all the same size and design – should be no more than about 20%. That's from our practical experience of taking ownership of properties arising from securitisations and managing the sale process ourselves rather than relying on bank servicers. That's what it costs to fix a new bathroom and kitchen, take control of the gardens and generally paint the place before sale. So the maximum spread of a series of spruce houses and neglected houses shouldn't be more than about 20%, and this spread should put a floor under the impact of dilapidation.
>
> The trouble is that it doesn't.

> The problem is that none of this process occurs in bank managed sales – all of which happen "as is" once the property falls into the servicer's ownership. No one in the chain of sales set up by general servicers has any incentive to take control of the value at sale in the interests of the owner. Everyone just wants the quickest possible sale and the removal of "in-hand" properties from the balance sheet. General servicing is passive and entirely accepting of any price as long as the property is sold.
>
> The net outcome is that the achievement rates in securitisations that need actual house sales in the future to achieve their cashflows is generally disappointing. But not because of old people becoming disinterested in the property. The main reason is the dislocation of active interest [i.e., incentive to secure a good price] arising from securitisation. The agents, lawyers, servicers and cash managers do not have the active interest in managing the actual house sale process that an owner would have. The securitisation process itself detaches ownership from effective sales management. So the low achievement rates we see in the data are mostly attributable to the involvement of the servicers, most of whom have no economic interest in maximising sales proceeds.[42]

There is, thus, a *second* moral hazard that works to the advantage of the lender relative to the borrower. ERM lenders appear to work on the presumption that ERM loans will usually expire in positive equity, and this assumption is consistent with the performance of most ERM loans to date. Given this presumption, then lenders have little incentive to manage the properties they take possession of in order to find a good selling price. Instead, their concern is simply to turn the properties around to recover the amounts owed, i.e., to make a quick sale, and any proceeds left over are returned to the borrower or their estate. If the borrower or their estate fetch a poor price, then that is not the lender's concern and the borrower or their estate has no control over the disposal process.

This second moral hazard is an additional reason for the low for ERMed properties to fetch low sales prices. We can think of this second moral hazard as creating an economic dilapidation effect in addition to the low sales prices created by physical dilapidation.

There is also a *third* moral hazard. To quote another of our correspondents:

> the issue from my experience is that everybody had an incentive to overvalue the property from the start, it's not an issue of dilapidation over time … overvaluation is one of the biggest risks and insurers who have no lending experience are just clueless to this. I have seen the same issue in debt consolidation mortgages across Europe: if the property value does not come from an actual transaction, when you sell the property you always get a surprise! …In a normal purchase the value of the property is quite certain because it is determined between a willing buyer and a willing seller, but in ERM there is no sale, only an appraisal. The lender, the broker, the appraiser

[42] A second correspondent confirms this analysis and refers us to a Fitch UK RMBS criteria addendum (18-May-2018) that states that a quick-sale adjustment of 17% of property value is applied to houses and 25% to flats. For illiquid properties (i.e. in the top 5% and 1% by price) a further 10% or 15% discount may be applied.

and the borrower all have an incentive to close the deal. Why? To meet their business plan volume. In these circumstances, history teaches that the property value will always be overestimated, but this will not appear in any Excel spreadsheet and therefore any modelling will be based off unreliable LTV assumptions.

In short, the impact of economic dilapidation as such would appear to be drowned out by multiple moral hazards, whose net impact might easily be confused with dilapidation unless one digs deeper into the underlying economic causes of low achievement rates.

These moral hazards, in turn, reflect underlying agency problems whereby those with fiduciary duties of care towards shareholders instead put their own interests first.

Statistical Analysis of the SAMS Achievement Rates

Table 9.1 shows the main statistical features of the achievement rates in the SAMS dataset:

Table 9.1: Main Statistical Features of the SAMS Achievement Rates

Mean	94.3%
Standard deviation	19%
Skewness	0.30
Kurtosis	3.82
Sample size	1420
Range	[31% 170%]
5% lower bound	64%
5% upper bound	125%

Source: SAMS.

We see that the achievement rates are highly dispersed, somewhat positively skewed and somewhat heavy tailed.

Note that the results in Table 9.1 reflect the whole dataset, and interpretation of these results is made difficult by the fact that these loans have different durations.

To make interpretation easier, Table 9.2 shows the same results for those loans that start in 1997 and end in 2017. These loans have an approximate duration of 21 years.

Table 9.2: Main Statistical Features of the SAMS Achievement Rates for Loans Spanning 1997 to 2017

Mean	91.0%
Standard deviation	23.4%
Skewness	0.01
Kurtosis	2.13
Sample size	69
Range	[47% 138%]
5% lower bound	48%
5% upper bound	129%

Source: SAMS.

Calibrating the Dilapidation Rate

We can then use these results to estimate the dilapidation rate.

For the 1420 loans in the entire SAMS dataset, the average duration is 10.74 years and the average achievement rate is 94.3%. The average annual dilapidation rate d is then

(9.1) $$d = -(1/10.74) \times \ln(0.943) = 0.54\%.$$

Alternatively, for the 69 loans in the SAMS dataset that start in 1997 and end in 2017, the average duration is about 21 years and the average achievement rate is 91.0%. The average annual dilapidation rate d is then

(9.2) $$d = -(1/21) \times \ln(0.91) = 0.45\%.$$

So for ERMed, properties, there is an additional physical/economic/stochastic dilapidation effect that is not present when dealing with typical non-ERMed properties, and the associated average dilapidation rate is about 0.5% a year.[43]

Finally, Table 9.3 shows these mean dilapidation results along with the 5% lower bound and the 5% upper bound for the dilapidation rate for the second set of loans.

---

[43] To quote a further comment from the previous correspondent. "Is the difference in absolute £ explainable with the costs of repairs and inefficiencies of an auction sale? Well, you can say that the sale is inefficient, or that the property was never worth that much, especially since the original value was not an arm's length price. There is so much fraud into those numbers, IMHO. That is not a stochastic process, rather systematic over-valuation. If I was a mortgage broker, I would say to a customer that an overvaluation of 20%+ should go totally unnoticed, and maybe we can push it to 30-40%. Beyond that, the lender probably has automated checks in place (Zoopla) though not in the 90s when the SAMS were originated. That fits with your 0.5% pa, just from a different perspective."

Table 9.3: Main Statistical Features of the SAMS Achievement Rates for Loans Spanning 1997 to 2017

Whole SAMS dataset	
Mean d	0.54%
Sample size	1420
SAMS loans spanning 1997 to 2017	
Mean d	0.45%
5% upper bound	3.5%
5% lower bound	-1.21%
Sample size	69

Source: SAMS.

These bounds reflect the 90% confidence bounds for the dilapidation rate, and span 3.5% at the upper end to -1.21% at the lower end. The small sample size tells us not to rely too heavily on these results, but they are indicative.[44]

We are tempted to conclude that there is a lot more going on here than mere physical dilapidation, although doubtless there would be that too. The 'other things going on' would appear to be the economic consequences of the various moral hazards we have mentioned. If these other factors dominate, as they appear to do, then the term 'dilapidation' is only a small part of the picture and, one might say, is looking a little dilapidated.

[44] A fuller analysis of the SAMS database might, e.g., look at the achievement rates and durations for each loan, estimate d for each, and thence we produce a sample of d rates, from which we can obtain a statistically stronger sense of them.

Chapter Ten: Volatility

Suppose that all the standard conditions usually assumed in the derivation of Black' 76 actually hold: complete markets, continuous trading and so on. In such circumstances we would be able to observe implied vols from options market data, and we would base our vol calibration on these implied vols.

Now suppose that we don't have these implied vols, e.g., because there is no traded market in residential property options. Our starting point then is to obtain a volatility or set of volatilities based on historical property prices, e.g., those from an historical house price index.

Existing Volatility Estimates

CP 13/18 states (p. 9):

> 2.16 The PRA estimated a value for the property volatility parameter from analysis of residential property price index data. Nationwide, Halifax and Office for National Statistics index data were analysed and several time series models were fitted to the quarterly log-returns of data sets over a variety of historical time periods. The PRA selected a parsimonious model that fitted the data well, and extracted from the model the unconditional volatility for various holding periods, allowing for autocorrelation. Further adjustments were made to allow for concentration risk and basis risk between the changes in prices of individual properties and the index. The PRA's *central estimate is of a 13% volatility assumption* for typical holding periods for ERMs, although use of alternative data choices gives a range of 13%-16%, and making an allowance for parameter uncertainty gives a range of 11%-18%. Estimates for property volatility provided to the PRA by firms are generally in the *range 10%-15%*.[45] (Our emphasis)

In his report, Professor Tunaru states (pp. 1, 20) that the Nationwide historical index data suggest a range of volatility values between 3.85% to 6.5%. His Table 1 (p. 19) then suggests Maximum Likelihood and Method of Moments estimates of around 3.95% for both Nationwide and Halifax datasets, and based on this evidence, he opts for a baseline vol of 3.9%. It seems to us that these volatility estimates are on the low side, and his own results reports for different UK regions and sample periods based on the Nationwide dataset suggest a range of volatilities spanning 3.85% to 6.5%, so we would suggest it would have been more prudent to have picked a value somewhere in the middle of this range, say about 5%, but arguably higher. But let's go with 5% *as a starting point*.

[45] Note however that CP 7/19 issues in April 2019 proposes to abandon the 13% volatility calibration and replace it with a regularly revised volatility estimate the first of which is seemingly yet to be determined.

Vol Estimation: A First Pass

Geometric Brownian Motion

To derive historical volatility, we can use any time series of prices so long as prices are available at a consistent sampling period k. This sampling period can be daily, weekly, monthly or any other period.[46]

The volatility is calculated by computing the returns (preferably log-returns) of the price series, then computing the non-annualised standard deviation $\sigma(k)$ of the returns. This volatility is specific to the frequency chosen – returns on the same asset can produce low volatilities if sampled daily, but higher volatilities if sampled less frequently, e.g., annually. However, standard pricing formulas such Black-Scholes/Black '76 assume that the returns have been sampled annually. To obtain the annualized vol or annualized standard deviation when returns are sampled at periods not equal to one year, we 'annualise' the standard deviation of returns by multiplying by the square root of the sampling frequency. For example, if returns are sampled every working day, i.e. with a frequency of 250 working days a year, we multiply the standard deviation of daily returns by root 250 to obtain an annualised volatility. If returns are sampled every month, we multiply by root 12, if quarterly (as typical for housing indices) we multiply by the square root of 4, and so on. Thus, if $\sigma(1)$ is the annualised vol, and $\sigma(k)$ is the (non-annualised) vol derived from a k year sampling period, then

$$\sigma(1) = non - annualised\ vol \times (1/k)^{0.5} = \sigma(k) \times (1/k)^{0.5} \tag{10.1}$$

or

$$\sigma(k) = \sigma(1) \times (k)^{0.5} \tag{10.2}$$

Table 10.1 shows a set of illustrative σ for a range of k values:[47]

Table 10.1: Illustrative σ for a Range of Resampling Periods

k	$\sigma(k)$
1	5%
5	11.2%
10	15.8%
15	19.4%
20	22.4%
25	25%
30	27.4%

Note: k = resampling period. Obtained using equation (10.2).

[46] Exceptions can be allowed, such as weekends, bank holidays, periods of missing data etc, and there are techniques for infilling data or correcting timing errors, but this is a well-understood and separate subject which we shall pass over.

[47] As the calculations underlying some of the tables in this chapter can be quite involved, we make available on our website an Excel workbook with the the relevant calculations for all the tables in this chapter.

So σ can vary from 5% to 27.4% depending on the sampling period k, given this range of k.

Autocorrelated underlying

The square root rule depends on the assumption that the asset log-returns follow Brownian motion, the condition for which includes randomness, i.e. the return at any time is independent of whatever happens any time in the past or the future. However the evidence indicates that property prices are autocorrelated so the independence assumption does not hold. In an autocorrelated process, a positive return yesterday increases the chance of a positive return today, and a negative return yesterday increases the chances of a negative return today. Consequently, the square root rule approach just outlined will not do and we need to use an alternative approach that takes account of autocorrelation.[48]

A practical way to implement this adjustment is via the Hurst exponent approach outlined in Appendix 1 to this chapter. We estimate the Hurst exponent H over our property index data set and obtain $\sigma(t)$ using

(10.3) $$\sigma(t) = 5\% \times t^{H}$$

where H would typically lie in the range from 0.7 to over 1. For UK data (source: Dallas Fed) $H \approx 0.82$, so

(10.4) $$\sigma(t) \approx 5\% \times t^{0.82}$$

For pricing purposes it is not the maturity t that matters but the rehedging period.

However, as noted above, standard pricing formulas such Black-Scholes/Black '76 assume that the returns have been sampled annually. If we rehedge every k years, we would then use the formula

(10.5) $$\sigma = 5\% \times k^{0.82}/k^{0.5} = 5\% \times k^{0.32}$$

to obtain the annualised volatility σ that we would input into our put pricing equation.[49]

---

[48] The theoretical solution to this problem was set out in an important paper "Option pricing and hedging with temporal correlations" by Lorenzo Cornalba, Jean-Philippe Bouchaud and Marc Potters. (L. Cornalba, J.-P. Bouchaud and M. Potters (2002) "Option Pricing and Hedging with Temporal Correlations," *International Journal of Theoretical and Applied Finance* 5(3): 307-320). This paper provides a fairly general analysis of the impact of temporal (i.e., auto-) correlation on option pricing and their conclusions are clear. "In the Gaussian case [the one considered in Black-Scholes], we find that the effect of [auto-] correlations can be compensated by a change in the hedging strategy and therefore *options should be priced using the standard uncorrelated Black-Scholes model* (our italics)." The required change can be implemented by measuring volatility on the same time scale as the rehedging, but this qualification merely amounts to an adjustment to the volatility calibration, if even that.

49 By 'annualised volatility' we mean here the volatility which, when *de*-annualised in the standard way by multiplying by root k, will give the same result as the empirically derived volatility of the autocorrelated series.

A proof of equation (10.5) is given in Appendix 2 to this chapter.

Table 10.2 shows a set of illustrative σ for a range of k values, assuming $H = 0.82$:

Table 10.2: Illustrative σ for a Range of Resampling Periods for $H = 0.82$

k	$\sigma(k)$	Annualised σ
1	5.0%	5.0%
5	18.7%	8.4%
10	33.0%	10.4%
15	46.1%	11.9%
20	58.3%	13.0%
25	70.0%	14.0%
30	81.3%	14.8%

Note: k = rehedging period, H=Hurst exponent = 0.82. $\sigma(k)$ is obtained using equation (10.4) with t set to k. Annualised volatility is obtained using equation (10.5).

We are now looking at σ in the range from 5% to 14.8% for this range of k, depending on our choice of k, i.e., our choice of hedging period.

Volatility Around the Index

So far we have only discussed the volatility of the index, but there is also the volatility around the index. This additional volatility would include the impact of regional variation around the index, but there are further contributory factors as well. These include, e.g., the impact of changes in consumer relative demand for different types of property, expansions of nearby roads, the impact of new housing estates, yuppification, middle class flight, the opening or closing of a good nearby school, and the stochastic dilapidation effects discussed in the previous chapter. One will recall Figure 9.2 from the previous chapter:

Figure 9.2: Indexed vs. Achieved House Prices

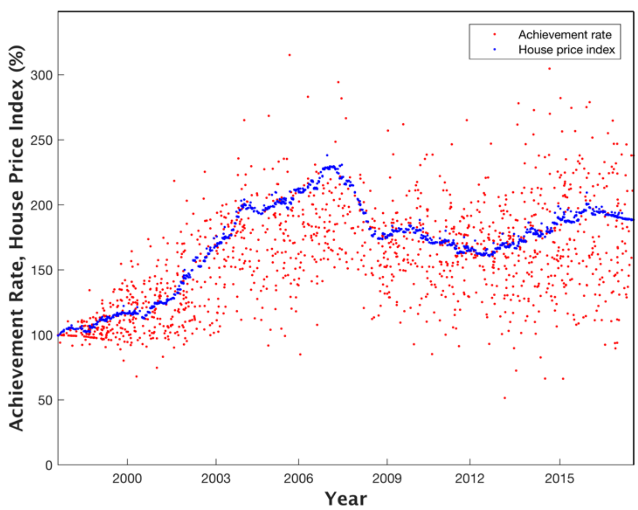

Source: SAMS

Recall that the darker blue random-looking line is a house price index and the red dots are a scattershot of the individual achievement rates in the sample.

We immediately see that the achieved values are much more volatile than the index. Above all, when seeking to calibrate the volatility, we need to keep in mind that it is that dispersion that matters, not the volatility of the index itself.

We can get an even better sense of the dispersion around the index from the next figure:

Figure 10.1: Indexed vs. Achieved House Prices (II)

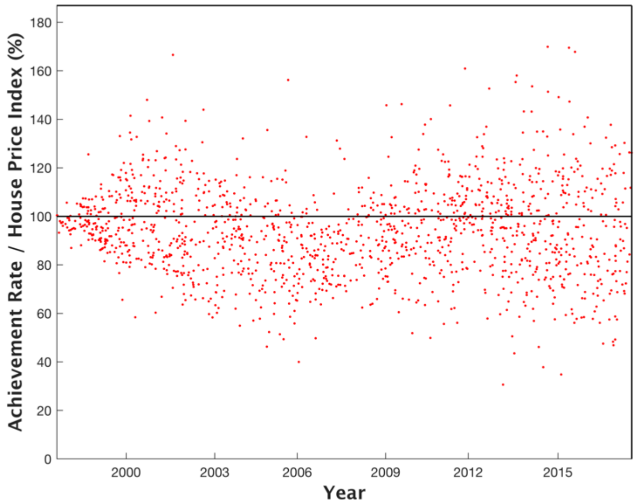

Source: SAMS

It would then behove us to revise our earlier index-based volatility estimates to take account of this additional volatility. We could start with our earlier volatilities reported in Table 10.2. Let us label this volatility as σ_{INDEX} whilst noting that it is dependent on the value of k. We now obtain σ_{AR}, the volatility of the achievement rate, as follows: take a rolling standard deviation of the achievement rate and divide by root time to get an annualised value. We assume here that there is no Hurst or autocorrelation effect for the achievement rate, i.e., that the square root law applies. We found that the annualised volatility values vary from 7% to 10%. Let's work with the middle value of 8.5%. We then have to assume a plausible correlation between the index vol and the achievement rate vol. Assuming zero correlation, which is not unreasonable, we then obtain the results reported in Table 10.3:

Table 10.3: Illustrative σ for a Range of Resampling Periods for $H = 0.82$

k	σ_{INDEX}	σ_{AR}	$\sigma_{INDEX\ and\ AR}$
1	5%	8.5%	9.9%
5	8.4%	8.5%	11.9%
10	10.4%	8.5%	13.5%
15	11.9%	8.5%	14.6%
20	13.0%	8.5%	15.6%
25	14.0%	8.5%	16.4%
30	14.8%	8.5%	17.1%

Note: k = rehedging period, H=Hurst exponent = 0.82, and the assumed correlation between the index and the achievement rate is zero. The terms in the rightmost column are obtained by Pythagoras.

The 'combined' or 'INDEX plus AR' volatility now varies from 9.9% for $k = 1$ to 17.1% for $k = 30$, but note that these numbers are little more than educated guesstimates and the 'true' numbers could be higher, e.g., if the correlation were positive.

The table and the figure show the strong effect of 'stochastic dominance', i.e. the tendency of a higher constituent volatility to dominate a lower one when correlation is negligible or, in this case, zero. In such cases the total volatility is given by Pythagoras. For example, if the first volatility is 10 and the second 1, the sum of squares is 101, the root is 10.05. So the lower volatility, while 10% of the higher one, has a marginal contribution of only 0.5%. Table 10.3 shows that at a maturity of one year, the achievement rate of 8.5% is driving the combined volatility of 9.9%. By contrast, at 30 years it is the rescaled index volatility of 14.8% that is the driving factor.

Interest Rate Risk as a Further Contributor to Volatility

We have hitherto assumed (as per Black-Scholes/Black '76) that the interest rate is constant, i.e., that there is no interest rate risk. In fact, interest rate risk not only exists, but arises from two sources. Consider the following components of the Black model, reproduced here in slightly simplified from Chapter 3:

(10.6) $$put = e^{-rt}[K \times N(-d_2) - F \times N(-d_1)]$$

(10.7) $$d_1 = [ln(F/K) + \sigma^2 t/2]/(\sigma\sqrt{t})$$

(10.8) $$d_2 = d_1 - \sigma\sqrt{t}$$

(10.9) $$F = S \times e^{(r-q)t}$$

where as usual: put is the value of the t decrement put, K is the strike price, F the forward price, σ the annualised input volatility, t the time to expiry in years, r the interest rate, S the price of 'spot' possession of the property, and $N(\ldots)$ is the cumulative normal distribution function.

The interest rate term r appears first (see (10.6)) as a *discount term* wrapped round the terms representing the future value of the put option, which brings the future value (i.e., $[K \times N(-d_2) - F \times N(-d_1)]$) back to present value. Here r plays the role of an outer discount factor.

r then appears again (see (10.9) as a *projection term* or inner discount factor taking us from the spot price S to the forward price F.

Each appearance gives rise to interest rate risk, but in different ways.

Discount rate risk

The first can be called discount interest rate risk. This risk can be hedged relatively easily, the gist of it being to swap floating into fixed.

A more detailed explanation goes as follows. When a trading desk sells an option, it places the premium in an account called the 'hedging account'. This account earns interest from the firm's central funding desk and the interest earned will typically be close to the firm's overall funding rate. To hedge the risk arising from changes in this outer discount factor, the desk should make an internal or external IR swap into a fixed rate with maturity at the option expiry date.

It can then be shown that this swap guarantees that, with no other change taking place in the market, the hedge account will earn the fixed rate r in the outer discount factor e^{-rt}. The demonstration goes as follows. Let

(10.10) $$P = FV \times e^{-rt}$$

where P is the option premium paid, r here is the long term rate earned on the option account, and FV is the future value of the option given by the *un*discounted Black formula.

Now suppose the long-term interest rate r changes, but there is no change in the forward price F. Such a circumstance would occur where the spot rate S changed by an amount ΔS in such a way that F remained constant under the formula connecting S with F, i.e.

(10.11) $$\Delta S = S(e^{-\Delta rt} - 1)$$

where Δr is the change in discount rate. In this situation F will remain the same and hence the future value FV of the put option will also remain the same. The change ΔP in the value of the option premium will then be a simple discount function:

(10.12) $$P + \Delta P = FV[e^{-(r+\Delta r)t} - e^{-rt}]$$

Assuming the amount P is currently held in the hedging account, we could replicate the change ΔP if P were invested in a long dated zero-coupon bond with maturity t. In practice the same effect can be achieved by investing P at the firm's short-term funding rate, but swapping the short-term floating payments into a zero-coupon swap.

Projection risk

In the previous example we assumed that the forward price remains constant while the spot price changes, i.e., a rise in long term interest rates will force the spot price lower, while a fall in the interest rate forces the spot price higher. This effect might be explained by a market expectation of unchanged future nominal rental cashflows, whose discounted present value would fall or rise according to the long-term interest rate operating as a discount factor.

The opposite case can also occur, i.e., we could have a situation where the spot remains steady, but the forward changes due to the way in which the interest rate operates as a projection factor.

Using the standard formula for the forward house price (i.e., (10.9)), and assuming constant q, the return on the forward is calculated as follows.

(10.13) $$\text{forward return} \approx \Delta HP + (\Delta IR - \Delta q) \times T$$

where $\Delta HP = \ln((S + \Delta S)/S)$ and T is the maturity of the forward at any point in the historical time series for the given combination of interest rate (IR), deferment rate (q) and house price index (HP).[50]

A proof of (10.13) is provided in Appendix 3 to this Chapter.

A number of points follow from (10.13).

Correlation between interest rate and house price index

First, given that there are four risk factors (Index, achievement rate, interest rate and deferment rate) impacting the forward price, we need to consider their correlations. We have already assumed that the correlation between the Index and the achievement rate is zero. The table below shows the correlation between the 10 year interest rate (which we take as a benchmark for the whole term structure) and the housing index, for 10 representative countries.

[50] Note that we cannot hedge away interest and deferment rate volatility on the assumption that the index price is less volatile than the forward. As we have pointed out in the Appendix to Chapter 3, the relevant volatility for the Black-Scholes option is not that of the spot price alone, but rather that of the d_1 term used to determine the option delta. The numerator in the d_1 term includes both interest rate and deferment rate – explicitly in Black Scholes, implicitly in Black 76.

Table 10.4: Correlation between 10Y Interest Rate and Index

Country	$\rho_{IR,INDEX}$
AUS 10Y	0.22
CAN 10Y	-0.05
GER 10Y	0.11
ESP 10Y	0.06
FRA 10Y	0.16
GB 10Y	**0.10**
IRL 10Y	-0.03
SWE 10Y	0.14
US 10Y	0.03
JP 10Y	0.27

Source: OECD (10Y interest rate) and Dallas Fed (House Price indices)

The takeaway points from this table are that the correlations between interest rates and house price indices are generally low and that a reasonable correlation for the UK would be zero.[51]

Correlation between q and house price index

Equation (10.13) also indicates that the deferment rate q is a further source of volatility. Take equation (7.1), which says that the deferment rate is equal to the rental yield divided by the house price, then replace the house price by the HP index and add a time t subscript. We then obtain:

(10.14) $$q_t = d_t/HP_t$$

where d_t is the aggregate nominal rental. We have no time series data on aggregate nominal rentals, but we can estimate their change using rental and house price indices. Data from OECD suggest that the deferment rate q is not constant (the annual volatility of q for the UK is of the order of 0.3%) and that changes in q are *negatively* correlated with changes in the index. These effects are shown in Figure 10.2:

[51] The low correlations reported in Table 10.4 are a bit of a surprise, considering that a lower interest rate immediately transforms into higher affordability and therefore - in the absence of new supply – into higher house prices. One reason could be that Table 10.4 looks at 10yr rates whereas the key driver might be short-term rates. Were the correlations between interest rates and house prices higher, the resulting 'combined' volatilities (of which more below) would be higher as well.

Figure 10.2: UK Nominal House and Rental Indices and Implied Deferment Rate

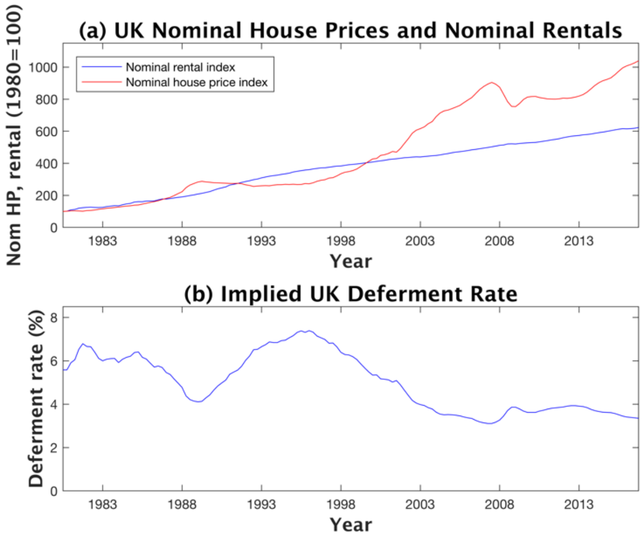

Sources: OECD

As Dean comments in a blog posting from earlier this year:[52]

> When I worked at the PRA on the paper that became CP 13/18, I had assumed that the deferment rate stays roughly constant. The rationale is that if rentals are expected to increase, this would increase the market value of properties, *all other things being equal*. But all other things aren't equal: there is strong evidence that *nominal* rentals track price inflation, and also strong evidence that interest rates anticipate inflation. So an increase in expected nominal rentals should correlate strongly with an increase in the interest rate used to discount the same rentals, and the rental yield, hence the deferment rate, should remain roughly constant.[53] I assumed this, and I imagine the PRA assumed this too.

A possible explanation for the volatile q rate and the negative correlation with house prices might then go as follows. Nominal house prices, which in theory should reflect the net present value of all future nominal (net) rental cashflows, tend quickly to anticipate – perhaps to over-anticipate – future upward or downward changes in rentals. Nominal rentals are sticky however and respond slowly.[54] Thus the large fall in house prices which occurred in the housing recession of the early 1990s was not reflected in rental prices, which continued to rise slowly, and so q rose in that period. Conversely, the significant

[52] "It Moves." The Eumaeus Project (14 January 2019): http://eumaeus.org/wordp/index.php/2019/01/14/it-moves/

[53] Reason, $q = r - g$ using a derivation of the Gordon model, where q is deferment rate, r is nominal interest rate, g is expected growth in nominal *rentals*. If g rises, and r is highly correlated with rental inflation, r will rise also, and the two effects will cancel out.

[54] This effect is well known in the literature, although there is no consensus on the explanation. For example, Campbell and Hercowitz (2009) find that "movements in U.S. house price-rent ratios cannot be fully explained by movements in subsequent rent growth" (J. R. Campbell and Z. Hercowitz (2009) "Welfare Implications of the Transition to High Household Debt." *Journal of Monetary Economics* 56, 1-16. For a review of the literature, see P. Gelain and K. Lansing "House Prices, Expectations, and Time-Varying Fundamentals," Federal Reserve Bank of San Francisco Working Paper 2013-03.

rise in house prices from the late 1990s to 2007 was notably higher than the rise in rentals, so q fell over this later period.

As an important aside, this combination of a volatile q that is negatively correlated with house prices has an interesting policy implication. If house prices go up, the loan-to-value of an existing equity release mortgage will fall, which will decrease the cost of the NNEG. At the same time, the graph above suggests the deferment rate will also fall, which will make the NNEG even cheaper, given that the deferment rate is the main driver of NNEG cost. Conversely, a fall in house prices will make the NNEG more expensive because of the fall itself, and will then make the NNEG even more expensive because of the implied rise in the q rate. The cost of the embedded guarantee is thus doubly geared to the state of the housing market. Ouch!

The PRA would appear to be still unaware of this double exposure, which has implications for how it should design its capital requirement regime for equity release. But we digress.

This negative correlation effect is not unique to the UK. Table 10.5 shows evidence for a strong and consistent negative correlation between the deferment rate and the house price index of our ten countries:

Table 10.5: Correlation between q and Index

Country	$\boldsymbol{\rho_{q,INDEX}}$
AUS	-0.90
CAN	-0.95
GER	-0.79
ESP	-0.88
FRA	-0.84
GB	**-0.82**
IRL	-0.43
SWE	-0.80
US	-0.86
JP	-0.92

Source: OECD (10Y interest rate) and Dallas Fed (House Price indices)

Volatility and term structure

Another corollary of equation (10.13) is that the impact of changes in interest rates on the forward return will decrease as time passes and the maturity shortens. For example, other things being equal, the impact of a 10bp change in interest rates on a 30 year forward will be to change the forward price by 10 x 30 = 300bp, which is significant. By contrast, the effect of the same change on a contract with 3 months to maturity will be 10 x 3/12 = 2.5bp, which is not significant. This changing sensitivity throughout the life of the contract means that the volatility caused by changes in interest (and deferment rates) is not constant. Instead, this volatility starts high and then falls towards zero as the contract approaches expiry.

We now wish to determine the average lifetime volatility of the contract. If X is the series of maturities and Y is the series of returns, it can then be shown that the volatility of the product $\sigma(XY)$ is the following simple function:

$$\sigma(XY) = \sigma(Y) \times T/\sqrt{3} \tag{10.14}$$

A proof is given in Appendix 4. The point to note is that the volatility of the product is now directly proportional to T.

Total Forward Volatility

Finally, we might consider the effect of all four risk factors (Index, IR, q and AR) in the forward rate (10.13) to give what we might call the total forward volatility. We can do so by estimating a correlation matrix between the four risk factors as shown in Table 10.6:

Table 10.6 Correlation Matrix for the Four Main Risk Factors

	Index	IR	q	AR
Index	1.00	0.00	-0.82	0.00
IR	0.00	1.00	0.00	0.00
q	-0.82	0.00	1.00	0.00
AR	0.00	0.00	0.00	1.00

Table 10.7 shows the volatilities for the 4 component risk factors:

Table 10.7: Volatilities of Component Risk Factors

Component Volatility	Value
σ_{INDEX}	5%
σ_{AR}	8.5%
σ_{IR}	0.58%
σ_q	0.17%

We next combine the correlations in Table 10.6 with the component volatilities in Table 10.7 to obtain the term structure for the total forward volatility shown in Table 10.8:

Table 10.8: Term Structure of Total Forward Volatility

t	Total Forward Volatility
1	10.0%
5	10.6%
10	12.2%
15	14.2%
20	16.5%
25	19.0%
30	21.6%

If we worked with these results, we would apply a 10% volatility to the put for decrement $t = 1$, a 10.6% volatility to the put for decrement 5, and so on, and an 21.6% volatility to the put for decrement 30.[55]

This term structure means that in principle there is no single volatility input for the series of puts that constitute the NNEG.[56]

Figure 10.3 shows a plot of the total forward volatility over a horizon of up to 50 years.

Figure 10.3: Volatility Term Structure

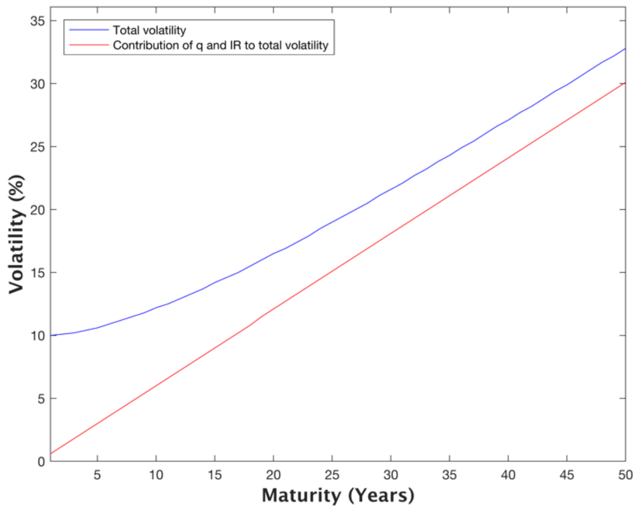

Notes: As per Table 10.6.

Note once again the effect of stochastic dominance. The interest rate volatility (0.58%) and deferment rate volatility (0.17%) make only a small marginal contribution at the shorter maturities. But recall (see equation 10.14) that the impacts of the interest rate and the deferment rate are proportional to maturity. After about 15 years, these risk factors become the 'strongest' driving factors and dominate almost entirely by about 30 years. This dominance explains why the blue total volatility curve increasingly looks like a straight line beyond those maturities. For comparison, the figure also includes a line (in red) representing the contribution of q and IR only, that is, the volatility of the forward in the events that (a) the index never changed and (b) there was no achievement rate volatility. The two lines move increasingly in parallel but do not in fact converge, the reason being the strong negative correlation between the index and q. If that correlation had been zero, the two curves would have eventually converged and we would have had the counterintuitive consequence that the volatility of the forward at these maturities would have not been driven by the housing index at all, but rather by the interest and deferment rates!

---

[55] A more comprehensive approach would take account of correlation under less frequent hedging and give a more detailed consideration of the correlation between the index and the achievement rate.

[56] However, for any given term structure and any given age, we shall show in Chapter 27 that it is still possible to impute a single volatility that gives the same NNEG valuation as one would get using a volatility term structure. However, this single volatility number must be consistent with the term structure and will be age-dependent.

Conclusions

The calibration of the volatility parameter is a more involved subject than is commonly suggested. We must take account of no less than four risk factors that affect volatility – the house price index, the achievement rates around that index, the interest rate and the deferment rate. We then end up with a volatility term structure in which different NNEG puts have different volatilities. We also find that the various volatility constituents have differing impacts depending on the maturity, with the house price constituents dominating at shorter maturities and the deferment rate and interest rate constituents dominating at the longer maturities.

Appendix One to Chapter Ten: A Hurst Exponent Approach to Autocorrelation in Property Prices

The empirical evidence suggests that property returns are autocorrelated and autocorrelation has implications for volatility extrapolation.

We can handle this autocorrelation using a Hurst analysis.

A Hurst Exponent Approach to Autocorrelation

By way of background, this problem appears to have first been observed by the hydrologist Harold Hurst (1880–1978) working on the Nile Basin in the 1930s.[57] Hurst was concerned to design an ideal reservoir that never overflows and never empties, based on observations of discharges from the lake. In any year t there will be an influx $\xi(t)$ of water into the reservoir, with a regulated discharge $\langle\xi\rangle_\tau$ from the reservoir, where τ represents a long period over which the reservoir operates. We need to estimate the storage required such that the average amount of water released over the period equals the average influx, without the reservoir emptying or overflowing at any intermediate time.

The average influx is

(10A1.1) $$\langle\xi\rangle_\tau = (1/\tau)\sum_{t=1}^{\tau}\xi(t)$$

which should be equal to the amount released per year. Let $X(t,\tau)$ be the accumulated departure of the influx $\xi(t)$ from the average $\langle\xi\rangle_\tau$:

(10A1.2) $$X(t,\tau) = \sum_{u=1}^{\tau}\xi(u) - \langle\xi\rangle_\tau.$$

The difference between the maximum and minimum accumulated influx is called the range R, which also represents the required storage capacity for an ideal reservoir that never empties nor overflows. The figure below (from Feder 1988, p. 151) sketches such a reservoir with influx $\xi(t)$, discharge $\langle\xi\rangle_\tau$ and range R. The height of the dam must be consistent with this storage capacity.

57 See H. E. Hurst, *Long-Term Storage: An Experimental Study* (London, Constable, 1965). The following account is adapted from J. Feder, *Fractals*, New York (Plenum 1988), pp. 149-151.

Figure 10A1.1: Reservoir Storage Capacity

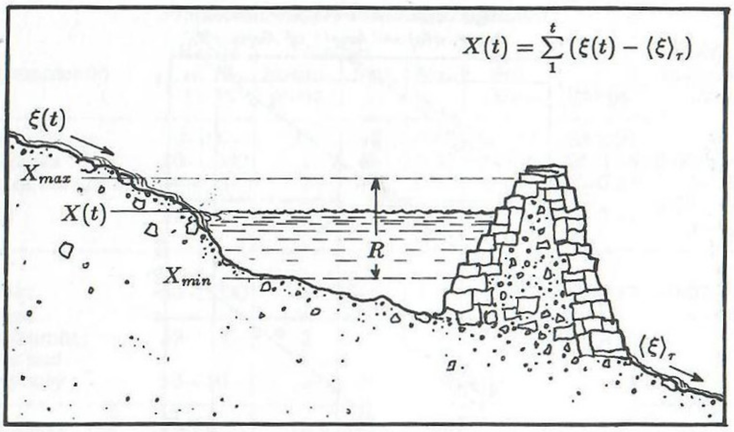

FIGURE 8.2: Sketch of a reservoir with an influx of $\xi(t)$, and an average discharge $\langle\xi\rangle_\tau$. The accumulated difference between influx and regulated outflow is $X(t)$. The range, R, is the difference between the maximum and minimum contents of the reservoir.

If influxes for successive annual periods are uncorrelated then it would be possible to work out the required range, hence work out an appropriate reservoir design, for any period τ simply by applying a square root law, using the following relation:

(10A1.3) $$R = S(\tau/2)^{0.5}$$

where S is the standard deviation of $X(t, \tau)$, on the assumption that successive annual inflows were random events. Empirical results, such as from the flow records of Lake Albert, which Hurst was engaged to work on in 1938, were somewhat different. He found that reservoir capacity based on the empirical results was larger than estimates based on the square root law. Extensive work by Hurst demonstrated that the flows were better explained by the following function

(10A14) $$R = S(\tau/2)^H$$

where H, the so-called Hurst exponent, typically varies empirically from 0.7 to 1.[58]

We can estimate H by fitting a power law to the data, i.e., by plotting $\ln(R/S)$ as a linear function of $\log\tau$. This fitting process is illustrated in Figure 10A.2:

[58] Originally called K by Hurst. Mandelbrot named it 'H' for Hurst.

Figure 10A1.2: Calibrating the Hurst Exponent Using a Linear Plot

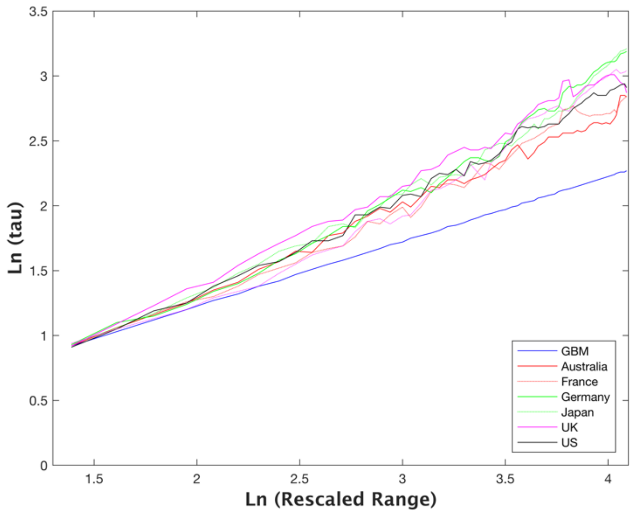

Notes: Based on Dallas Fed data.

The figure shows plots of the rescaled range $\ln(R/S)$ against $\ln \tau$ for a variety of national house price indices and for a simulated GBM series. The slope of each line gives the Hurst exponent, H. If the time series is generated by a random walk (or a Brownian motion process) then the slope or H exponent has a value of 0.5. The Figure shows that the various national house price series have H values between about 0.7 and 1, but the simulated GBM series has a slope of 0.5, as we would expect.

We can also estimate H using a method recently developed by Ceballos and Largo (2017).[59]

We analysed quarterly data for housing markets in 21 different countries and found the same phenomenon as Hurst found with the Nile, with results for H varying from 0.69 (Australia) to 1.01 (US). See Table 10A1.1 below:

Table 10A1.1: Hurst Exponent (H) by country

Country	H	Country	H	Country	H
Australia	0.69	Japan	0.82	Spain	0.89
New Zealand	0.78	S. Africa	0.84	Sweden	0.92
Denmark	0.81	Italy	0.86	Germany	0.95
Belgium	0.81	Finland	0.86	Luxembourg	0.95
Canada	0.81	S. Korea	0.86	Ireland	0.99
Norway	0.81	Switzerland	0.86	France	1.00
UK	0.82	Netherlands	0.88	US	1.01

Data source: Dallas Fed. The estimation uses the Ceballos and Largo adjustment for small values of τ.

So why does the Hurst exponent matter to us? The answer is because of its implications for volatility extrapolation. Under GBM, we would extrapolate volatility using the square root rule:

---

[59] See R. F. Ceballos and F. F. Largo (2017) "On the Estimation of the Hurst Exponent Using Adjusted Rescaled Range Analysis, Detrended Fluctuation Analysis and Variance Time Plot: A Case of Exponential Distribution," *Imperial Journal of Interdisciplinary Research* 3(8): 424:434.

$$\sigma(t) = 5\% \times t^{0.5} \quad (10A1.5)$$

where t is the holding or rehedging period used to determine σ, but under autocorrelation, we should extrapolate using

$$\sigma(t) = 5\% \times t^{H} \quad (10A1.6)$$

We assume the result of Parkinson (1980) showing that the volatility of returns over any holding period t could be estimated using the high and low in t.[60]

Smith and Jeffery (2019:24) also find, using empirical data from the UK housing market, that the volatility of returns depends on the chosen holding period, and rises with it.

Whether autocorrelation really matters to us depends however on the choice of hedging strategy underlying the option pricing. In particular, we would use (10A1.6) to obtain a projected t-period volatility if we are pricing the option using a rehedging strategy that calls for the synthetic option or underlying position to be rebalanced every t years. But in other cases autocorrelation is not an issue for us. One such case is where $t = 1$, in which case (10A1.5) and (10A1.6) coincide.

[60] M. Parkinson (1980) "The Extreme Value Method for Estimating the Variance of the Rate of Return." *Journal of Business* 53, 61–65.

Appendix Two to Chapter Ten: Proof of Equation (10.5)

Step #1: $\sigma(k)$ by definition is the volatility we empirically derive by sampling returns with period k, and taking the root mean square, without any kind of adjustment.

Step #2: To obtain the 'annualised' value σ we divide by root k. Thus

(10A2.1) $$\sigma = \sigma(k)/\sqrt{k}$$

where σ is the value we input into Black '76.

Step #3: From Step #1 above, it follows that $\sigma(1)$ is the volatility we empirically derive by sampling returns with period 1.

Step #4: Assuming GBM, then

(10A2.2) $$\sigma(1) = \sigma$$

Step #5: With the Hurst effect, (10A2.2) is false. Specifically,

(10A2.3) $$\sigma(k) = \sigma(1) \times k^H$$

where $H \neq 0.5$.

Step #6: From Step #2 above,

(10A2.4) $$\sigma = \sigma(k)/\sqrt{k} = \sigma(1) \times k^H/\sqrt{k}$$

which was to be proved.

Appendix Three to Chapter Ten: Proof of Approximation (10.13)

The forward rate F at any time t and for any maturity T is as follows:

(10A3.1) $$F_{t,T} = S_t e^{(r(t,T)-q(t,T))T}$$

where S_t is the spot price at time t, $r(t,T)$ is the interest rate of maturity T at time t and $q(t,T)$ is the deferment rate of maturity T at time t. With the passage of time Δt, the forward rate will change as a result of changes in S, r and q, and of course with the passage of time itself. Thus

(10A3.2) $$F_{t+\Delta t,T-\Delta t} = S_{t+\Delta t} e^{\left(r(t+\Delta t,T-\Delta t)-q(t+\Delta t,T-\Delta t)\right)(T-\Delta t)}$$

This expression is fairly complex, but we can make a number of simplifying assumptions as follows. First, we can assume that the term structure of both r and q is continuous. We have assumed throughout a flat term structure q, so it follows that $q(t+\Delta t, T-\Delta t)$ equals $q(t+\Delta t, T)$. We cannot assume that the term structure of interest rates is flat, because it will usually slope upwards or downwards at any time. However, we can reasonably assume that changes in the term structure will make no significant contribution to volatility. That is, a change over 1 month to the 10 year interest rate will not be significantly different from the change in the 9 year 11 month interest rate. Thus $r(t+\Delta t, T-\Delta t)$ will be approximately equal to $r(t+\Delta t, T)$, for small Δt. Hence

(10A3.3) $$F_{t+\Delta t,T-\Delta t} \approx S_{t+\Delta t} e^{\left(r(t+\Delta t,T)-q(t+\Delta t,T)\right)(T-\Delta t)}$$

The outer term $(T-\Delta t)$ can also be eliminated, as it represents a constant carry through time. As time passes, if r is greater than q, the forward price will gradually fall, or if r is less than q, the forward price will gradually rise. But volatility corresponds to the mean *difference* from the average, whereas the carry term will be close to the average itself. Hence

(10A3.4) $$F_{t+\Delta t,T-\Delta t} \approx S_{t+\Delta t} e^{\left(r(t+\Delta t,T)-q(t+\Delta t,T)\right)T}$$

We assume that the determinants of forward volatility are the changes in spot, interest rate and deferment rates alone, and that the passage of time is an insignificant contribution to volatility.

To determine the volatility, we must first determine the forward price return:

(10A3.5) Forward price return = ln $(F_{t+\Delta t,T-\Delta t}/F_{t,T})$

Substituting from the equation above:

(10A3.6) $$\ln(F_{t+\Delta t,T-\Delta t}/F_{t,T}) =$$

$$\ln\left[S_{t+\Delta t}\, e^{\left(r(t+\Delta t,T)-q(t+\Delta t,T)\right)T}\right] - ln\left[S_t e^{\left(r(t,T)-q(t,T)\right)T}\right] =$$

$$ln[S_{t+\Delta t}/S_t] + [r(t+\Delta t,T) - r(t,T) + q(t,T) - q(t+\Delta t,T)] \times T$$

Now make the simplifying assumptions that $r(t+\Delta t,T) - r(t,T) = \Delta r_t$ and $q(t+\Delta t,T) - q(t,T) = \Delta q_t$. We then obtain:

(10A3.7) forward return $\approx \mathit{\Delta} HP_t + (\mathit{\Delta} r_t - \mathit{\Delta} q_t) \times T$

which was to be proved, where $\mathit{\Delta} HP = \ln((S+\Delta S)/S)$.

Appendix Four to Chapter Ten: Proof of Equation (10.14)

We need to determine the volatility of a time series of prices for a forward contract, given that the maturity T of the contract is constantly decreasing. Assume the following standard result for two independent variables X and Y:[61]

(10A4.1) $$Var(XY) = Var(X)Var(Y) + Var(Y)E[X]^2 + Var(X)E[Y]^2$$

Let X be the series of maturities, and Y be the changes in interest rate Δr_t (or deferment rate Δq_t). Assume that the average interest rate or deferment rate change is zero, i.e. that $E[Y] = 0$.

(10A4.2) $$Var(XY) = Var(X)Var(Y) + Var(Y)E[X]^2 + Var(X)E[Y]^2 = Var(Y)[Var(X) + E[X]^2]$$

Then we can treat the series of maturities as a uniform distribution from the starting maturity T down to zero. The variance $Var(X)$ and the average $E[X]$ of a uniform distribution over the interval (x, y) are as follows.

(10A4.3) $$Var(X) = (y - x)^2/12 = T^2/12$$

(10A4.4) $$E[X] = T/2$$

Substituting:

(10A4.5) $$Var(XY) = Var(Y)[Var(X) + E[X]^2] = Var(Y)[T^2/12 + T^2/4] = Var(Y) \times T^2/3$$

(10A4.6) $$\sigma(XY) = \sqrt{Var(XY)} = \sigma(Y) \times T/\sqrt{3}$$

which was to be proven.

[61] L. A. Goodman (December 1960). "On the Exact Variance of Products," *Journal of the American Statistical Association* 55 (292): 708.

Chapter Eleven: Modelling Mortality

1. Introduction

This chapter provides an introduction to the mortality modelling issues that arise with equity release valuation.

2. Realised Mortality Rates

Consider the following Figure showing the realised mortality rates of England & Wales males for ages 70 and above.

Figure 11.1: Realised Mortality Rates for Males 70 and Above

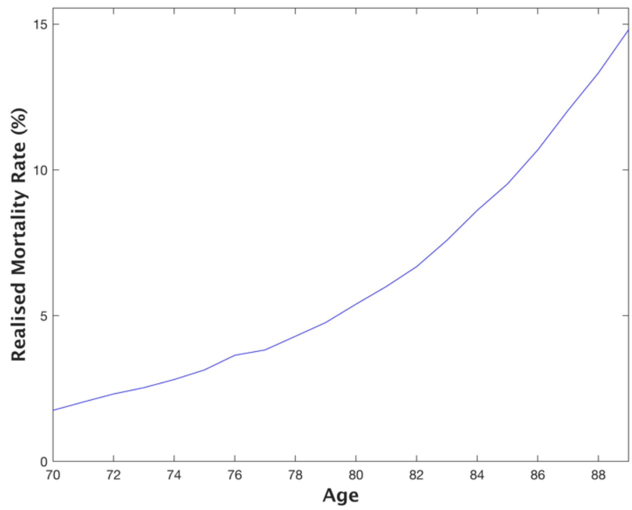

Notes: Based on CBD-M5 model (Cairns *et alia*, 2006, 2009) projections using England & Wales male deaths rate data estimated over ages 55:89 and years 1971:2017. Source: llma.org.

By realised we mean the mortality rates experienced for these males over the sample years, 1971:2017.[62]

We see that the mortality rate is a little under 2% for age 70 and then rises with age, as we would expect. However, there are two problems with the mortality rates in this Figure. The first is that they only go up to age 89, which is the maximum age in our sample age range, and for NNEG valuation we want mortality rates for higher ages as well. The

[62] Let D be the number of people of a certain age who die in a certain year and let E be the number of corresponding exposures or people of that age at risk of dying that year. The death rate $m = D/E$ and the mortality rate is $q = 1 - e^{-m}$ (see, e.g., Cairns *et alia*, 2009, p. 3). The mortality rate is a more convenient rate to use than the death rate because it is mathematically more tractable and because it is guaranteed to be less than 100%, whereas death rates can exceed 100% due, e.g., to errors in the exposures data, which might be due to misreported birth dates or the occasional unreported homicide.

second is that these realised mortality rates take no account of any anticipated longevity improvements.

3. Projected Mortality Rates

A traditional solution to this latter problem is to use some specially prepared mortality tables, which are in essence expert educated guesses about future mortality rates. Such tables can be obtained from the Continuous Mortality Investigation, for example.[63] This approach is easier to apply but is subjective and lacks transparency. A more scientifically grounded and more transparent and therefore better approach (although there is a bit of a learning curve) is to use mortality rate projections from a stochastic mortality model.[64] One such model is the CBD-M5 model mentioned earlier, which was designed specifically for old age mortality projections.

Figure 11.2 gives the realised and projected future mortality rates for males just turned 70:[65]

Figure 11.2: Prospective Mortality Rates for Males Aged 70

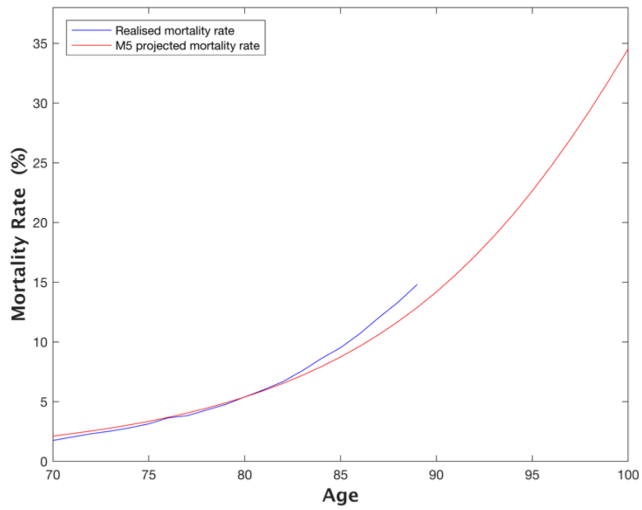

Notes: As per Figure 11.1.

We have shown the projected mortality rates out to age 100, but the model allows us to project them to any age we wish.[66] We see that the projected mortality rates grow at a

63 See https://www.actuaries.org.uk/learn-and-develop/continuous-mortality-investigation/cmi-mortality-and-morbidity-tables.

64 These models not only allow users to obtain likely projected mortality rates, but also allow them to obtain statistically grounded prediction intervals and scenario analyses, and to take account of refinements like parameter uncertainty, individual death risk and the impact of Bayesian prior beliefs.

65 These mortality rates are known as cohort *q* rates, because they follow the cohort currently aged 70 as they age over time and their mortality rates reflect their increasing age.

66 For NNEG valuation, we actually use projected mortality rates out to age 120, i.e., for modelling purposes we assume that any individuals who reach their 121st birthday are then automatically dispatched. It is important to take account of the extreme old age 'toxic tail' when dealing with lifetime financial products.

lower rate than the realised rates. This difference between the two plots illustrates the impact of projected improvements on future longevity.

4. Expected Survivor Rates

The next step in the analysis is to obtain the corresponding expected survivor rates, S_t, i.e., the probability that an individual alive now will survive to year t. If we let q_t be the mortality rate for year t, then the following holds for S_t:

(11.1) $$S_0 = 1$$
(11.2) $$S_1 = (1 - q_1)$$
(11.3) $$S_2 = S_1 \times (1 - q_1) = (1 - q_1) \times (1 - q_2) \text{ etc.}$$

Figure 11.3 shows the expected survivor rates for the people on their 70th birthday:

Figure 11.3: Expected Survivor Rates for Males Aged 70

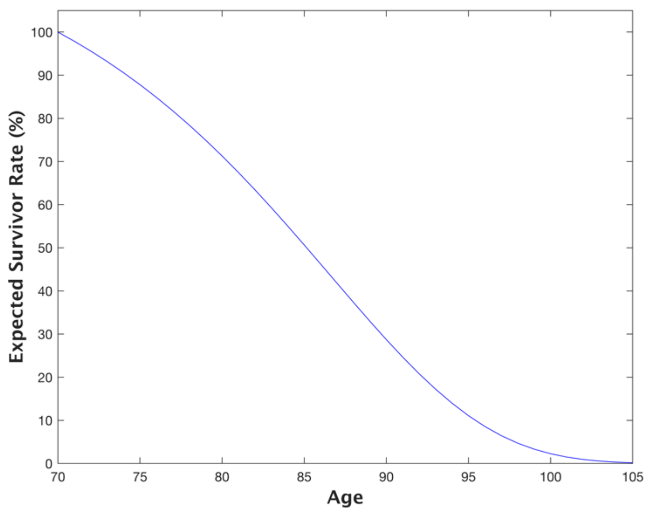

Notes: As per Figure 11.1.

The expected survivor rates fall from 100% on day 1 down eventually towards zero.

5. Expected Mortality Rates

We take the expected mortality rate (of a cohort of given current age) for future period t to be the product of the mortality rate for t and the expected survivor rate for t, i.e., the expected mortality rate equals $q_t \times S_t$. The expected mortality rates for a cohort of males currently aged 70 are shown in the next Figure:

Figure 11.4: Figure 3.1 Exit Probabilities for Males Currently Aged 70

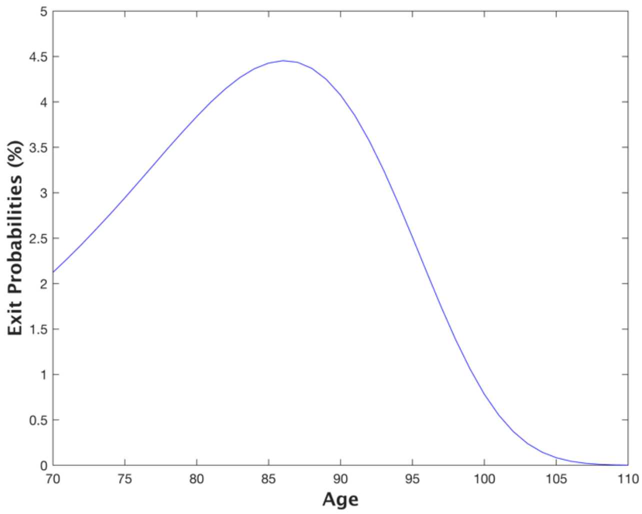

Notes: As per Figure 11.1.

which is essentially the same as Figure 3.1, if we ignore the possibility of house exit by going into a nursing home.

We see that the expected mortality rates rise to peak in the mid to late 80s, then fall and eventually go to zero.

Now remember that these expected mortality rates are the weights that apply to the NNEG put options. This weighting schedule tells us that the low maturity puts are not so important, because of the low probability of exercise, but the puts become more important as their maturity rises and the ones that are of most significance of those with maturities of maybe 10 to a little over 20 years. The puts with longer maturities have declining and eventually insignificant importance.

6. Model Risk in Mortality Projections

A problem with these projections is that they are dependent on an assumed mortality model, M5. They are therefore exposed to mortality model risk, i.e., the risk of error from the use of the M5 mortality model.

The best way to address this issue is to consider alternative models and one that would be suitable is M7. This model is an extension of M5. Model 5 has two period effects, but M7 adds a third period effect and a cohort or year of birth effect to M5. Further details can be found in, e.g., Cairns *et alia*, 2009.[67]

Figure 11.5 shows the expected mortality rates for 70 year old males based on projections from both mortality models:

---

[67] Those familiar with these models might be wondering why we don't also include model M6. The answer is that for reasons we don't quite understand, M6 gives potentially unreliable results for England & Wales females. We therefore drop it from consideration.

Figure 11.5: Expected Mortality Rates for Males Currently Aged 70: Models M5 and M7

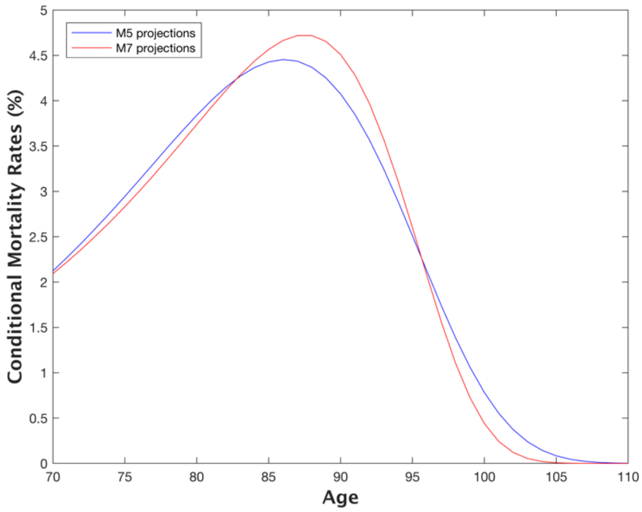

Notes: As per Figure 11.1.

The plots are distinct but have much the same shape. Consequently, we would expect these models to give fairly close NNEG valuations.

7. Model Risk in Mortality Projections

Table 11.1 shows the valuations corresponding to each of the mortality models:

Table 11.1: Baseline NNEG and ERM Valuations for Males Aged 70

Current House Price	Loan Amount	Model	L	$NNEG$	ERM
£100	£40	M5	£74.84	£32.19	£42.66
£100	£40	M7	£74.29	£31.46	£42.83

Notes: L is the present value of the loan component of the Equity Release Mortgage, $NNEG$ is the present value of the NNEG guarantee and ERM is the present value of the Equity Release Mortgage. Based on the baseline assumptions: male aged 70, LTV=40%, r=1.5%, l=5.25%, q=4.2% and σ=14.8%. Exit probabilities are based on projections of the M5 and M7 variations of the CBD model using England & Wales male deaths rate data spanning years 1971:2017 and ages 55:89.

We see a certain amount of variation in the loan values and NNEG valuations, but much less variation in the ERM valuations due to the offsetting impacts of the loan values and the NNEG valuations on the ERM valuations.

8. Females

Table 11.2 gives the corresponding valuation results for females:

Table 11.2: Baseline NNEG and ERM Valuations for Females Aged 70

Current House Price	Loan Amount	Model	L	$NNEG$	ERM
£100	£30	M5	£80.79	£40.28	£40.52
£100	£30	M7	£81.99	£41.84	£40.15

Notes: As per Table 11.1 but for females. The appropriate volatility for females aged 70 is 15.7%.

The female loan value and NNEG valuations are higher than for the males, as we would expect from higher female life expectancy. The valuations from the two models are also very close.

9. Joint Lives

Many ERM loans are to couples rather than individuals. In theory, one should model the exit probabilities associated with such loans in a way that takes account of the longevity prospects of both partners and the point that exit will occur when the longest surviving partner exits the house. Such an analysis is a little involved, however, and the standard approach is to treat such an ERM loan as if it were a loan to the youngest partner.

A typical case would be a couple in which the male is 70 and the female 66, i.e., so we have a younger female but bear in mind that females have longer life expectancy. Some results are given in Table 11.3:

Table 11.3: NNEG Valuations and Life Expectancies for a Typical Couple

Current House Price	Loan Amount	Partners	$NNEG$	Life expectancy
£100	£40	Male aged 70	£32.19	15.1
£100	£40	Female aged 66	£48.87	20.6

Notes: $NNEG$ is the present value of the NNEG guarantee. Based on the baseline assumptions: LTV=40% for the male and 36% for the female, r=1.5%, l=5.25%, q=4.2% and σ=14.8% for the male and 17.3% for the female. Exit probabilities and life expectancies are based on projections of the M5 variation of the CBD model using England & Wales male deaths rate data spanning years 1971:2017 and ages 55:89

The table shows NNEG valuations and associated life expectancies, where the NNEG is that of an ERM loan to each partner considered on their own. The NNEG valuations are £32.19 for the loan to the 70 year old male and £48.87 for the loan to the 66 year old female.

Bear in mind that the NNEG for a loan to a couple, with exit deemed to occur when the last remaining member exits, will always be larger than the NNEG from a loan to either individual member alone. The reason is that the latter NNEG valuations do not account for the risk (to the lender) of the other partner exiting later than the partner to whom to the individual loan was made. Thus, the true NNEG valuation would be bigger than either of the NNEG valuations shown in the Table and it is immediately apparent that one could get a major under-estimation of the $NNEG$ if one had treated the loan for NNEG valuation purposes as if it just a loan to a male aged 70.

Even the higher (female 66) NNEG of £48.87 will under-estimate the true NNEG, but if one looks at the life expectancies, it is also clear that the female is expected to live 5.5 years longer than her male partner. It is therefore unlikely that the male will outlive her, and one can make a plausible hand waving argument to the effect that the error in the NNEG valuation will not too important.[68]

However, any such argument relies on the specifics of the case at hand – and in particular, on the younger member having a notably longer life expectancy – and that consideration will not always apply. For example, if the male and female were reversed in age, then the 66 year old male would have a life expectancy of 17.1 years and his 70 year old female partner would have a life expectancy of 18.4 years. There is then a much bigger risk of the older member of the couple outliving the younger one, and the NNEG valuation error, associated with treating the loan as if it were a loan to the younger member, will be greater than in the previous case.

10. Impaired Lives

ERM companies will sometimes offer more attractive terms to borrowers with impaired lives, i.e., those with reduced longevity prospects. A first pass at modelling fair-value NNEG guarantees for impaired lives borrowers is to establish the impact of their health condition on their expected longevity, then offer them an LTV based on that expected longevity. To give a simple example, if a borrower aged 70 has the life expectancy of an 80 year old, then the lender might offer them the LTV i.e. loan terms of an 80 year old, e.g., so the borrower might get a loan based on an LTV of 40% instead of the standard 30%. For NNEG valuation purposes, then, we might treat the 70 year old as if he were an 80 year old.

As a further refinement, we might also take account of whether and if so how, their health condition might affect the prospects for the length of their end of life period in care.

[68] There are also two offsetting effects on any NNEG under-valuation. First, there is the possibility that the surviving member of the couple will move out after first one has died. Second, there is some evidence for a 'heartbreak syndrome' by which the death of the first member raises the mortality rate of the other, although anecdotal evidence suggests that the opposite can sometimes occur too.

Chapter Twelve: Long-Term Care

We have hitherto assumed that exit occurs with death. In reality, many older people leave their homes to move into some form of long-term care (LTC) and the time spent in care can be substantial.

Length of Time in Long-Term Care

We understand that a rule of thumb in this area is that people are expected to spend two years in long-term care. Or perhaps a little longer: "On average, older people stay in a residential care home for 30 months," states a recent Independent Age article citing an earlier (2009) report by LangBuisson.[69]

These numbers are comparable to those from the U.S. For example, Thomas Day (2010) states that, the "average stay for elderly patients who die in a nursing home is just shy of 2 years,"[70] whilst Brad Breeding cites U.S. reports of 2009 and 2010 that report average lengths of stay in assisted living facilities of about 28 months and 29 months respectively.[71]

But what everyone agrees on is that the length of time in care, and whether one will need to spend any time in care at all, are highly uncertain.

Taking Account of Expected Time in Care

The difference between the expected time to death (i.e., life expectancy) and the length of time to house exit could then make a material difference to equity release valuations.

The problem is how to take account of this difference.

The standard actuarial practice in the UK is to obtain exit probabilities from (conditional) mortality probabilities by imposing loading factors on the latter. A loading factor involves multiplying those mortality rates by a factor that reflects some LTC 'add on'. For example, for the purposes of obtaining exit probabilities from mortality rates, we might add an additional 2% to the latter. Our exit probabilities would then be 102% times the (conditional) mortality rates.

Hosty *et alia* (2007) offer a schedule of such loading factors:

---

[69] See "Cost of average length of stay in a residential home is equivalent to 26 years' worth of family holidays." Independent Age 26 Oct 2017; LaingBuisson *Care of Older People Market Report*, 27 edition, London: LaingBuisson.
[70] T. Day (2010) "About Nursing Homes."
https://www.longtermcarelink.net/eldercare/nursing_home.htm
[71] B. Breeding (2016) "So I'll Probably Need Long-Term Care, But for How Long? myLifeSite (6 July).
https://www.mylifesite.net/blog/post/so-ill-probably-need-long-term-care-but-for-how-long/

Table 12.1: Hosty *et alia* (2007) Mortality Loading Factors

Age	Male (%)	Female (%)
≤ 70	2	3
(70,80]	4	12
(80,90]	5	13
(90,100]	4	8

Source: Hosty *et alia* (2007, Table 13).

These loading factors rise with age and then fall again and reflect an intuition about the expected length of time spent in care outside the home, depending on the age when one goes into care.

There are problems with this approach, however. One was set out by Tunaru (2019, p. 67):

> For multi-state modelling considering the interaction between long-term care entry and mortality is paramount because there is significantly higher mortality experienced by long-term care residents compared to "at home" mortality means that to maintain the same aggregate assumption for mortality by age lighter than average mortality should be assumed for "at home" lives.

People in care will be less well than people of the same age still at home and so their mortality rates will be higher. Since the population mortality rates are averages of the mortality rates of people in care and people outside care, then any adjustment for the former requires an offsetting (but not necessarily equally offsetting) adjustment for the latter. The problem here is that without further data on the numbers in care vs the numbers outside care, then any such adjustments are difficult to carry out.

There is also a deeper issue. Even if we were confident in our projections, we have no reason to think that these loading factors are any good in the first place. To calculate them 'properly' from first principles, we would need (a) reliable projections for the mortality rates of people in care, (b) reliable projections for people outside care and (c) reliable projections of the relative sizes of these two populations. The problem is that we don't have any of these. Instead, we have are loading factors pulled out of thin air, i.e., guesstimates, admittedly by life actuaries with some intuition for such issues, and that is all we have.

So we shouldn't place much reliance on the Hosty *et alia* loading factors.

Fortunately, we don't have to, because there is a better way and it is simple too. Suppose we believe that the expected time in long-term care is 2 years. If we have a 70 year old male, then his life expectancy is about 15 years, and we expect him to spend his last 2 years in care, i.e., we expect him to exit his home in 13 years. So we can approximate his time to exit by giving him the life expectancy of a 72 year old, and we can do that by giving him the projected mortality rates of a 72 year old. Thus, we take our NNEG and ERM functions, and input ages of 72 instead of 70.

Table 12.2 gives some illustrate valuations.

Table 12.2: NNEG Valuations and Life Expectancies for a Typical Male

Current Age	$NNEG$	ERM
70	£31.19	£42.66
72	£26.95	£44.33

Notes: $NNEG$ is the present value of the NNEG guarantee. ERM is the present value of the ERM. Based on the baseline assumptions: LTV=40%, r=1.5%, l=5.25%, q=4.2% and σ=14.8% for age 70 and 14.1% for age 72 . Exit probabilities and life expectancies are based on projections of the M5 variation of the CBD model using England & Wales male deaths rate data spanning years 1971:2017 and ages 55:89.

We see that this age adjustment makes a substantial difference to the NNEG valuation but a much smaller one to the ERM valuation.

If we believe that the expected time in care is 2.5 years, then we can obtain our approximations by taking our NNEG to be average of the NNEGs for ages 72 and 73, and so forth.

Ideally, we might want to build some fancy model, but in the absence of data, we can't. In the meantime the simplistic approach just suggested could be the best we can do.

Chapter Thirteen: Delayed Possession

We have hitherto implicitly assumed that the loan is immediately repaid at the time of house exit. In practice, however, there is likely to be some delay. For example, in their NNEG study, Li *et alia* (2010) define δ as the average delay in time from the point of exit to the sale of the house. They use a baseline calibration of δ = half a year (p. 16) and provide some sensitivity results, but they do not provide any empirical justification for their δ calibrations.

In theory, we take account of the impact of delays on our valuations if we have good data that allows us to estimate the expected delay. Let's suppose, for example, that we expect there to be a half year delay between house exit and house sale. We can then handle this delay by treating the borrower as if he or she were half a year younger. So if the borrower is actually 70 when the loan is taken out, we treat the borrower as if he or she were 69 and then obtain the NNEG or ERM as equal to the average of the NNEGs or ERMs for 69 year hold and 70 year old.

The actuary or modeler can presumably use their firm's account database to obtain some sense of the delayed experienced on past ERM loans.

However, it may not always be so straightforward, and for two reasons. First there will be dispersion in the 'time to sale' statistics, and some of observed times to sale can be long. In one case we have seen, the time to sale was three years. A second issue is whether the loan fixes at exit or death and has delayed settlement or keeps rolling until settlement. In the latter case, there is the potential for some serious abuse. One of our correspondents, an expert in ERM property management practices, informs us that

> a bank or building society can also use long term assets (like ERMs) to hide its errors of valuation at the beginning of a long dated book ... I have seen some naughty institutions filter the harvest of ERM assets as they roll off by posting the cash results of those that expire at LTV < 100% and just allowing the > 100% ones to roll on and on, without addressing or crystallising them. So from a reports and accounts perspective (where little of this level of analysis is exposed to the public) it looks as though the cash is rolling in nicely as planned at inception, but the horrors (the NNEGs that should eat up the accumulated cash accounting value of the assets) are also piling up unnoticed, because uncrystallised – deal with them later as it were. So the book looks like it's the best cash generating asset on the balance sheet – cashflow soaring and other loans still rolling up. Because there is no explicit "credit" effect on an ERM, there is no obvious "default" if no-one decides to recognise it.
>
> Don't forget that even if someone dies, the recognition of that fact and the movement to sale of the underlying collateral can be delayed for ages too as you can claim to be negotiating with the heirs etc etc, and all the while the asset is still rolling up at the compound rate (even though economically the asset should have been crystallised much closer to actual death/LTC) when in fact it is getting ever more above 100% LTV. But you can argue that it is still a legal claim on the estate, even if its unlikely ever to be fully realised – you can

> say that you were enmeshed in negotiations with the estate and hoping for a sudden doubling of the housing market. I have seen actual cases in the UK where the time from both mortgagors final death date to my portfolio appraisal date as putative buyer has been greater than 7 years – and STILL the asset shows as current, nicely rolling up etc etc, when it passed 100% BEFORE they died, and when presumably its now full of squatters and falling down!
>
> Lift the corner of this mat and all the little bugs come rushing out……...

Even so, he suggests, “a bit of lagging is small beer compared to the wide [achievement rate] dispersion no matter what.”

Chapter Fourteen: Credit Spreads

We have so far assumed that we are discounting at risk-free, i.e., with no credit spread. Not everyone agrees with this approach, however.

There are two reasons why someone might want to add a credit spread to the discount rate. The first is because they want to get lower NNEG valuations, and the second is because they believe that there should be a credit spread. Both of these are wrong and the first self-evidently so: there is never a good reason for tweaking calibrations to get desired outcomes. But even if well-intended, it is still a mistake to add in the credit spread.

A good example of the credit spread argument is provided by David Land at the IFoA sessional event on ERMs on 11 December 2018. He works for Rothesay Life, who are a major provider of annuities and a major user of the Matching Adjustment.[72] They also just bought a whacking great portfolio of Equity Release Mortgages from the taxpayer.[73] It makes sense that Rothesay would be interested in an approach that lowers NNEG valuations.

We then wonder if Rothesay are using a higher-than-LIBOR funding rate to cheapen the cost of their guarantees. If they are, they shouldn't be.

In the discussion at the December 11 event, Mr. Land asked a pointed question. If the working party hasn't yet fixed the right method of calculating the forward, isn't that a pretty major source of possible error? No coherent answer emerged, but Land raised an interesting point. So what is going on? Well, if we can't lower the value of the no negative equity guarantee by putting in an optimistic growth forecast, perhaps we can tweak the funding rate instead. He drops a hint when he suggests that there's a large range of possible funding rates that we might consider and goes on to state that "The PRA thinks that you could possibly fund a house at LIBOR flat, which seems remarkably difficult." The implication is that we should be adding in a spread.

Land's argument is a seemingly plausible one and provides a much better case for lowering NNEG valuations than the so-called 'real world' approach that the industry are so enamoured of. It is still wrong, but it is wrong in a more interesting way.

Now we agree with him that it would be an unusual lender who made a risky loan at 'LIBOR flat'. If they are to lend at all, they would charge a spread to compensate for the risk of loss from default. After all, the borrower might default and leave the lender with a loss because insufficient collateral has been posted.

---

[72] See their *2017 Solvency and Financial Conditions Report* at https://www.rothesaylife.com/media/1183/rothesay-life-sfcr-2017.pdf

[73] See http://eumaeus.org/wordp/index.php/2018/09/28/doing-gods-work/ and http://eumaeus.org/wordp/index.php/2018/10/05/casting-magic-upon-daylight. The substance is that in September 2018 Rothesay Life bought an £860m portfolio of equity release loans from UK Asset Resolution Limited (UKAR), a government agency established on 1 October 2010 to 'facilitate the orderly management' of the closed mortgage books of both Bradford & Bingley and NRAM.(Source: Oliver Ralph, FT, 27 September 2018, UKAR website). The precise amount paid by Rothesay was not disclosed

But consider the following. Suppose we did add a spread and suppose (because it is true and easily verified) that adding the spread produced a lower NNEG. You then have a situation where the worse the credit quality of the borrower, the higher the spread; and the higher the spread, the *cheaper* the NNEG, because the NNEG is based on a set of put options priced off the forward. So the spread argument doesn't pass the sniff test.

Let's try to draw out what Land may have had in mind. Start with:

(14.1) $$R = S \times e^{-qt}$$

(14.2) $$F = S \times e^{-(r+s-q)t}$$

where R is the price of a reversion aka deferment contract, S is the spot price of the income producing asset e.g., a property, q is the discount rate aka deferment rate representing the present value of lost income over the term of t years. F is the price of a forward contract, r the risk free and s the funding spread over LIBOR.

If we use Black 76 then the underlying is the forward price F rather than the deferment price R. Since the NNEG for any decrement t is a put option, the higher the forward price F, the more the option is out of the money, and the lower the cost of the NNEG guarantee. So we can make the NNEG at least appear cheaper by assuming a non-zero spread s.

Suppose then we are a firm holding a deferment contract for possession at t and we want to hedge our risk by selling the corresponding forward contract. Let's also take the case most favourable to the credit spread argument, which is where our counterparty has no collateral at all. If we price the contract per equation (14.2) above, it is certainly true we may want to charge a spread over the risk-free rate r to compensate for the risk that the purchaser of the forward contract will default at t. So how do we represent the value of our hedged position in our books? Assuming no risk of default, the present value of the position is

(14.3) $$R' = R \times e^{st}$$

But there clearly is a risk of default, because we just assumed one when we applied the spread s! We must therefore reserve against that risk, but by how much? Well if s is a spread representing the cost of default, i.e., if the difference in present value attributable to s is precisely the cost of default, then that same difference is the amount we would reserve. Consequently

(14.4) $$\text{Reserve } = R \times e^{st} - R$$

We must then subtract the reserve from our portfolio value:

(14.5) $$\text{Portfolio value } = R' - \text{reserve} = R \times e^{st} - (R \times e^{st} - R) = R$$

We are then back where we started and the correct spread is zero! You cannot cheapen the value of a deferment contract by hedging it in the derivatives market. Whatever increase in *value is attributable to applying a spread over LIBOR represents a compensation for risk, which must be subtracted again as a reserve.*

Remember too that this case is the one most favourable to the credit spread argument, that in which the counterparty provides no collateral. If the forward contract is fully collateralised, on the other hand, then there would be no credit spread because there would be no credit risk. So either the spread goes in and then comes out again, or it doesn't go in to start with. And in the intermediate case where the contract is partially collateralised, we have a weighted average of these two cases and again no spread.

We have only considered the case of a forward, but the same applies to an option. We have just shown that we cannot raise the forward rate by tweaking the funding rate. The funding rate should be the risk free rate, not the higher rate that might be offered by a bank to, say, a buy-to-let investor. But if we cannot change the forward rate, then it follows that the NNEG valuation must remain unchanged as well.

Chapter Fifteen: Drawdown

Drawdown ERMs differ from the 'conventional' ERMs (often called Lifetime Mortgages, LTMs) that we have considered so far in that the borrower contracts to receive a drawdown facility rather than a loan taken out at any one time. This facility would give the borrower the right but not the obligation to draw down on the facility at his/her discretion, up to the maximum amount of the drawdown facility. Typically, the borrower would draw down an initial amount when the facility is set up with a view to making subsequent drawdowns later.

The loan contract would specify a current drawdown lending rate, which we understand would typically be a little higher (maybe 5-6 basis points higher) than the lending rate on an LTM mortgage. Once a drawdown is made, the lending rate on that loan tranche is fixed. The drawdown lending rate, like the LTM loan rate, will change over time as the lender periodically adjusts the rate in response to changes in market conditions and its own lending policy.

The existence of the drawdown facility makes Drawdown ERMs more involved than LTM ERMs, but there has been little discussion of how to handle the drawdown facility from the ERM valuation perspective.[74]

One way to model valuations on drawdown ERMs and NNEGs would be to use past historical data – which the firm would have – on its drawdown experience. Its valuation specialists could then use these data to project the timing of future drawdowns and the amounts drawn down each time. For example, they might anticipate that a female aged 70 would draw down 50% of the remaining facility at age 74 and draw down the rest of the facility at age 78. They could also forecast the drawdown lending rates at these future times. For ERM and NNEG valuation purposes they could then treat the drawdown ERM as equivalent to a portfolio of three separate LTM loans and we already know how to value the ERMs and NNEGs of LTM loans: there would be an initial LTM loan taken out now when the borrower is 70, a second LTM loan that is expected to be taken out in 4 years' time by the same female when she is 74 and a third LTM loan that is expected to be taken out in 8 years' time by the same female when she is 78.

An alternative way to model valuations is just to assume that the whole of the drawdown facility is drawn down in one go when the contract is made. In that case, the drawdown ERM loan would be treated as equivalent to a single LTM loan.

[74] In its CP 7/19 (section 2.15 to 2.18), the PRA has an extensive discussion of drawdown issues that is, none the less, quite solution-lite. See PRA Consultation Paper 7/19 "Solvency II: Equity Release Mortgages – Part 2," April 2019.

Chapter Sixteen: Early Repayment

Many ERM contracts allow the borrower to repay earlier under certain conditions. Now early repayment of mortgages is a notoriously fiendish problem.[75] It is well known, for example, that prepayment rates on mortgage-backed securities are sensitive to interest rates and MBS modellers have had great difficulties getting this modelling 'right'. One might however suppose (or at least hope) that the prospects of future repayments induced by lower interest rates would not be such a problem at current low interest rates.

The usual approach suggested in the NNEG literature is to work with assumed repayment rates, which might be either pulled out of thin air or based on historical repayment behaviour. Tunaru (2019, section 10.9) has a good discussion of this subject, but not much is known.[76] Historical prepayment rates tend to be high in the early years but then tail off. Such rates make sense intuitively: in the early year, borrowers may regret the loan and still be able to pay it back; in the later years, it becomes increasingly difficult to pay off early as the loan amount rolls up and their financial circumstances deteriorate.

The best we can do is probably to use historical repayment rates as a proxy for future repayment probabilities.

How to Handle to Early Repayment

One way to make use of these for NNEG and ERM valuation goes as follows. Imagine for the moment that an ERM contract had no facility for early repayment. Then the value of the NNEG, $NNEG$, and the value of the ERM, ERM, would be what they were in previous chapters, and we can move on. Put it this way: if $NNEG^{no\ repay}$ is the value of the NNEG without the possibility of early repayment and so forth, then

(16.1) $$NNEG = NNEG^{no\ repay}$$

(16.2) $$ERM = ERM^{no\ repay}$$

We now introduce the possibility of early repayment and ask ourselves: how does the possibility of early repayment affect NNEG and ERM valuations?

Let's start with a pre-existing ERM loan, and let its value be ERM^{old}. If the old/pre-existing ERM loan is repaid, then a certain amount – X, say – is paid to the lender and the ERM is extinguished. X would include any early repayment charges stipulated by the ERM contract. The important point for the modeller is that we can easily work out what X might be.

---

[75] For more on these issues, see, e.g., JP Morgan *MBS Primer*, June 2006, or A. Davidson and A. Levin *Mortgage Valuation Models: Embedded Options, Risk, and Uncertainty*, Oxford University Press, 2014.

[76] Another problem is that the ratings agencies take a very conservative approach to repayment. They typically model mortgage securities using a yield to maturity approach, and then account for prepayment as reinvestment at swaps minus 50 bps. This approach makes prepayment appear very expensive for lenders.

We need now to make an assumption about the lender does with the money it receives. It could invest in T-bill, hold it in cash and so on, but let's suppose that the lender uses the money that is repaid to make a new ERM loan, ERM^{new}.[77]

Let's suppose also that process of converting the loan repayment to the point where a new ERM loan is made costs the firm Y, where Y would be the costs of marketing, distribution and so on, net of any such charges to be applied to the new customer. Again, the modeller can work out what Y might be.

Consequently, the amount left over to be loaned to the new customer is $X - Y$. To make the loan amount explicit in the valuation of the new ERM, call the latter $ERM^{new}|(X - Y)$. Thus, $ERM^{new}|(X - Y)$ is the value of the new ERM loan, where $X - Y$ is the amount loaned at inception. Note that the input calibrations of the new ERM loan might and typically would be different from those of the old ERM loan, e.g., the loan rate might have changed. However, the modeller would have the information needed to obtain the value of the new ERM loan.

So when an ERM loan is paid off early, the ERM lender loses an asset worth ERM^{old}, but acquires an asset worth $ERM^{new}|(X - Y)$, and by hypothesis we know how to obtain $ERM^{new}|(X - Y)$.

The 'transaction' may or may not be beneficial to the lender, but the lender has no choice but to accept the loan repayment if the terms of original loan allow it, and we are assuming that the best the lender can do with the money repaid is to invest it into a new ERM loan.

Ex ante, we don't know whether the borrower will repay early or not, but we can suppose that the modeller has access to data on past loan repayment frequencies, e.g., if the modeller is an ERM actuary working for a large firm, then he or she will have access to the firm's loan histories, and can then estimate past loan repayment frequencies. From these, the modeller can estimate ϕ, the probability that a given ERM loan will be repaid over some future period. The modeller can then obtain the value of the ERM taking into account the possibility of early repayment, e.g.:

(16.3) $$ERM = \phi \times ERM^{new}|(X - Y) + (1 - \phi) \times ERM^{no\ repay}$$

[77] This assumption will not always be appropriate, however. As one of our readers points out: "valuing one ERM asset on the assumption that you can originate another, at will, to order, when needed, out of thin air, is going to struggle to pass IASB or SII or MAP constraints. Looks like falsely capitalising the value of reinvestment risk to me." His points are well-taken. If re-investing the proceeds from an early repayment into a new ERM is an issue, then one needs take a view on how else those proceeds should be invested. For example, would they be invested in gilts, and so on. But the basic approach we set out here can easily be tweaked to handle such alternatives.

Chapter Seventeen: Fees, Charges and Expenses

In principle, we might want to take account of fees and charges. We say "might" for good reason: if fees, charges and expenses are fairly set, i.e., they are neither excessive nor too low, then if those fees etc. are paid by the borrower, then they have no impact on the lender and hence would be irrelevant for NNEG valuation. They would also be irrelevant for ERM valuation, at least from the lender's perspective.

Whether fees and so forth are fairly set is another matter and is beyond the scope of this report. Accordingly, in this chapter we set out what information we have on fees and charges, and leave it to practitioners to decide if, and if so, how, these might affect NNEG valuations. For what it is worth, our view is that these are second-order issues relative to the big issues like the calibration of the deferment rate or volatilities. On the other hand, fees and charges would be of considerable interest to borrowers and would certainly relevant when doing value-for-money analyses. Those however are beyond our remit.

Hosty *et alia* on fees and charges

Hosty *et alia* (2007, p, 29) report the following policy expenses summary:

Table 6. Policy expenses summary

Item	Range	Assumption
Distribution and sales	1%-6% of the customer advance for a broker distribution model, similar for direct telesales, 5% or more for a direct sales model offering home visits	2.5% of the customer advance
Marketing	From 0.5% for limited support to a broker distribution model to 5% for cold advertising to generate leads for a direct sales force	1.0% of the customer advance
Administration		
Initial	£250-£750	£500
In force	£40 p.a.-£120 p.a.	£60
Termination	£200-£500	£350
Property Management		
Initial valuation	Charged to customer	£nil
Re-inspections and maintenance issues	Included in "in force" admin fee	
Property sale	Charged to customer except for NNEG situation	
Funding	Assumed to be covered in "cost of funds"	
Expense Inflation	Upwards of assumed CPI at 2%-3% p.a.	3% p.a.

Notes:

1. It is assumed the valuation costs, provider's legal fees and disbursements are all charged to the customer. In addition a customer accessing the product through a broker channel may also have to pay an additional advice fee.
2. The cost of reinspections will depend very much on the providers' practices. We have assumed costs for a provider who reinspects properties only as needed to cover very specific situations e.g. complaints handling. For providers who operate a programme of sampling or systematic periodic revaluations then costs may be significantly higher.

Our understanding is that the 'distribution and sales' and 'marketing' charges both apply. Thus, the total charge for distribution, sales and marketing would be 3.5% of the loan.

Hosty *et al.* (2007, p. 31) also report a 'specimen product specification' that includes similar items but in addition includes:

- provider's legal fee (£300) and property valuation fee (£1 per mille) to added to the loan;
- property sales expenses equal to 2% of final property value: these would be charged to the borrower unless the NNEG is activated;
- additional costs in event of negative equity claim: £500 to cover costs of additional valuation and administration; and
- early repayment charges: "Mark to market with 25% cap", whatever that means.

These charges and fees were in 2007 £s. According to the Bank of England's inflation calculator the cumulative inflation since then has been 13.6%, so we would have to multiply these numbers by 1.136 to get their equivalents at today's price level.[78]

Hosty *et alia* (2008) provide the additional clarification that 'other costs' are passed to the customer, presumably with the exception of property sales expenses if the NNEG is activated, because to do otherwise would violate the NNEG.

Contemporary Fees and Charges

Loan rates: we have seen contemporary loan rates varying from 4.15% AER to 6.78% AER. It pays to shop around.

Repayment charges: we have seen cases where the charges to repay loans were about 9.6% of the amount owed.

For one large firm we have seen charges and fees that include:

- Loan rate = 5.35% AER
- Start fee = £600
- Advice fee = £995
- Legal fee = £650
- End fee = £125.

Our understanding is that for this firm, all other costs and expenses would be passed to the borrower, unless doing so would violate the NNEG.

There is also the issue of repayment charges, and we are aware of one contemporary case in which a borrower faced a repayment charge of about 16% of the rolled-up loan amount and it struck us that this charge was high.

There could also be other fees and charges, but we have been unable to do a systematic search.

A surprising example recently came to light, however. A recent *Private Eye* article revealed a nice little scheme involving Age UK.[79] Potential borrowers looking at the Age Concern website who are interested in equity release are encouraged to "dip [their] toe in the water with help from the Age Co UK Equity Release Advice Service, provided by Hub Financial Solutions Ltd." As the *Eye* article continues. "Hub, it turns out, is owned by, er, Just Group. But its not just good business for Just. Age UK itself takes a handy commission of "up to 0.75 percent of the amount advanced under each equity release plan sold, together with a contribution towards marking support.""

[78] https://www.bankofengland.co.uk/monetary-policy/inflation/inflation-calculator
[79] "Just Reward." *Private Eye* No. 1490 22 February -7 March 2019, p. 39.

Chapter Eighteen: Scenario Analysis and Stress Testing

A scenario analysis is a hypothetical 'what if' exercise in which we examine what might happen to some variable of interest (e.g., a NNEG or ERM valuation) if some future scenario were to play out. For example, we might examine what would happen according to our model if future house prices were to behave in a particular way.

A stress test is a scenario analysis in which the posited scenario is an adverse one (e.g., a large drop in house prices).

Type 1 House Price Scenario Analysis or Stress Test

One type of scenario analysis/stress test is to model the impact of an immediate one-off house price fall. We assume that house prices fall five minutes after the ERM loan has been made. This exercise is easy to carry out. We first value the NNEG or ERM at the initial house price value. We then re-value them immediately after the house price fall using the new LTV. For example, if the initial LTV = 40% and house prices fall by 50%, then the new LTV will be $0.4 \times 1/(1-0.5) = 80\%$. For this type of stress test we do not make any projections of future variables, e.g., future house prices, other than that house prices fall shortly after the ERM loan is made.

Table 18.1 gives the impact of an immediate one-off house price fall of 50%.

Table 18.1: Impact of an Immediate 50% Fall in House Prices

	$\boldsymbol{L}$	$\boldsymbol{NNEG}$	$\boldsymbol{ERM}$
Pre-stress	£74.8	£32.2	£42.7
Post-stress	£74.8	£48.4	£26.4
	$\boldsymbol{\Delta L}$	$\boldsymbol{\Delta NNEG}$	$\boldsymbol{\Delta ERM}$
Impact of stress	£0	£16.2	-£16.2
% Impact	0	50.4	-38.1

Notes: $NNEG$ is the present value of the NNEG guarantee and ERM is the present value of the Equity Release Mortgage. Based on the baseline assumptions: male aged 70, r=1.5%, l=5.25%, q=4.2% and σ=14.8%. For the pre-stress scenario, we assume house price = £100 and LTV=40%; for the stress scenario, we assume house price = £50 and LTV=80%. Exit probabilities are based on M5-CBD model projections using England & Wales male deaths rate data spanning years 1971:2017 and ages 55:89.

For the given set of calibrations, $NNEG$ rises by £16.2 and ERM falls by the same amount. $NNEG$ increases by 50.4% and ERM falls by 38.1%.

Figure 18.1 shows the impact on NNEG and ERM valuations of a range of immediate one-off house price falls.

Figure 18.1: Impact of a Range of Immediate House Price Falls on NNEG and ERM

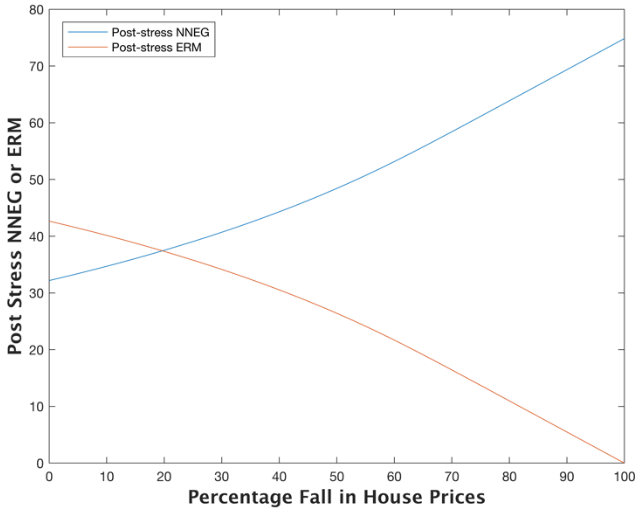

Notes: As per Table 18.1.

We see that the greater the fall in house prices, the greater the fall in ERM valuation.

Type 2 House Price Scenario Analysis or Stress Test

In a second type of exercise, we assume a hypothetical rate of growth of hpi and show how, e.g., the ERM valuation would behave over time under the posited scenario. This type of exercise works as follows. We start by using our ERM valuation model to obtain the current value of the ERM. Then we obtain the expected value of the ERM after one year, which would be 1 year exit probability times the expected payoff if exit occurs in year 1, plus the 1 year probability of no exit times the ERM value after one year. This latter ERM value is obtained using our ERM valuation model, but taking account of the new age after 1 year (i.e., 1 year older), the new house price after one year (which is equal to the old house price times e^{hpi}), and the new LTV after one year (which is equal to the old LTV times e^{l-hpi}). We carry on in like manner for the expected value of the ERM after two years and so forth.

In practice, we would often want to compare two different scenarios: a base scenario, which might be the scenario we expect, and a stress or adverse scenario. For example, we might follow Just Group and assume $hpi = 4.25\%$ for our expected or base scenario (see Chapter 24 below). We then posit some stress scenario, e.g., $hpi = -1.7\%$, which was the average hpi in Japan over 1990:2017. Figure 18.2 shows a plot of these two scenarios:

Figure 18.2: Expected ERM Valuations Under 4.25% vs. Japan House Price Growth Scenarios

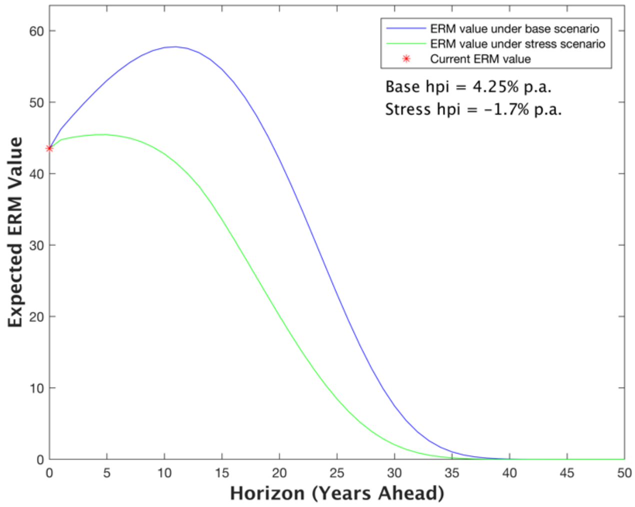

Notes: As per Table 18.1.

We see a much lower expected ERM valuation projection under the stress scenario. The lesson here is that if we were relying on $hpi = 4.25\%$ but actual hpi turns out to -1.7%, then the large ERM increases we were expecting will not come pass, and future ERM valuations will decline to zero considerably more quickly than we had expected.

Expected Cashflows under Type 2 House Price Stress Test

Another type of stress test is project cashflows under an assumed house price scenario. Figure 18. 3 show the projected cashflows under the same two scenarios.

Figure 18.3: Expected ERM Cashflows Under 4.25% vs Japan House Price Growth Scenarios

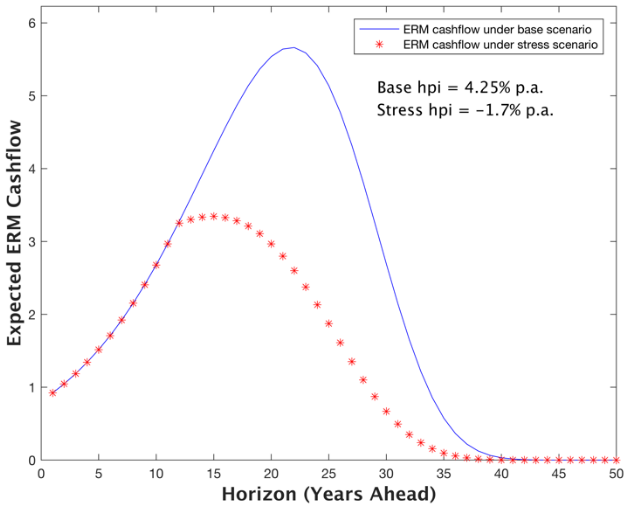

Notes: As per Table 18.1.

If the 4.25% scenario were to transpire then subsequently realised cashflows would be much larger than if the Japan scenario were to transpire, and the Japan-style cashflows would be even lower if we also experienced, say, 95% achievement rate instead of the 100% achievement rate assumed so far:

Figure 18.4: Expected ERM Cashflows Under 4.25% vs Japan House Price Growth Scenarios (II)

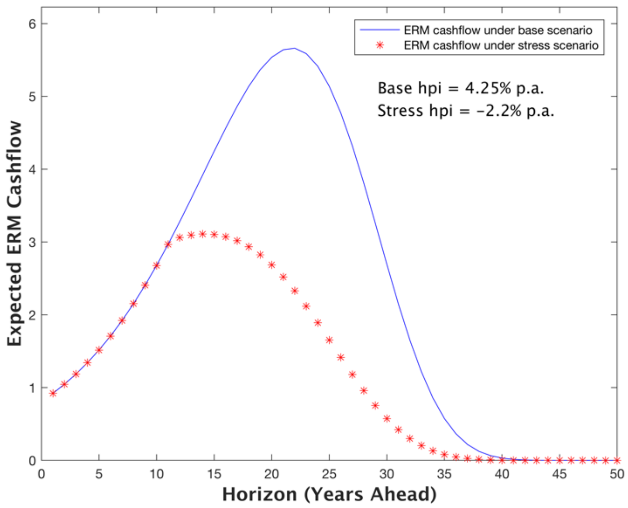

Notes: As per Table 18.1.

Longevity Scenario Analysis or Stress Test

We can also carry out scenario analyses or stress tests based on other projections. An important one in the equity release context would be a longevity scenario analysis, in which we posit some change to expected longevity and consider the impact of that change on NNEG or ERM valuations.

An easy way to carry out such an exercise is to assume a particular change in longevity, say, we assume that longevity suddenly increases by 3 years. We can then approximate the impact of this scenario by reducing by 3 years the age of the individual inputted into our valuation model, whilst keeping other parameters (and especially the LTV) the same as they were.

Table 18.2 gives the results of such an exercise.

Table 18.2: Impact of a 3 Year Increase in Longevity

	L	$NNEG$	ERM
Current	£74.8	£32.2	£42.6
Under scenario	£82.8	£42.4	£40.4
	$\Delta NNEG$	$\Delta NNEG$	ΔERM
Impact of scenario	£7.9	£10.2	-£2.3

Notes: L is the present value of the risk-free loan component of the ERM, $NNEG$ is the present value of the NNEG guarantee and ERM is the present value of the Equity Release Mortgage. Based on the baseline assumptions: house price=£100, male aged 70, LTV=40%, r=1.5%, l=6%, q=4.2%, and σ=14.8%. For the scenario, we input an age of 67. Exit probabilities are based on M5-CBD model projections using England & Wales male deaths rate data spanning years 1971:2017 and ages 55:89. Source of data: llma.org.

We see that both L and $NNEG$ rise by comparable amounts and largely offset in their net impact on the value of the ERM.

It is also interesting to note that the sign of the impact on ERM is *negative*. An increase in longevity decreases the value of the ERM. This result means that the impact of increased longevity on the value of the ERM has the wrong sign to function as hedge to an annuity book. The liabilities of an annuity book increase when longevity rises, so to function as a hedge, an asset must also increase in value, but in this case the ERM decreases in value instead.

Chapter Nineteen: The PRA's Good Practice ERM Valuation Principles

In its Supervisory Statement SS 3/17 published in July 2017, the UK Prudential Regulation Authority set out certain good practice principles relating to ERM portfolios. These principles include two that impose upper bounds on ERM valuations.

Principle II

Principle II states:

> The economic value of ERM cash flows cannot be greater than either the value of an equivalent loan without an NNEG or the present value of deferred possession of the property providing collateral.

i.e.,

(19.1) $$ERM \leq L \text{ and } ERM \leq PV(F)$$

where PV(.) is the present value of the term in (.).

Proof that $ERM \leq L$

Start with

(3.1) $$ERM = L - NNEG$$

We know that

(19.2) $$NNEG \geq 0.$$

If $NNEG > 0$ then

(19.3) $$ERM = L - NNEG < L.$$

If $NNEG = 0$ then

(19.4) $$ERM = L - NNEG = L.$$

Hence

(19.5) $$ERM \leq L$$

which was to be proved.

Proof that $ERM \leq PV(F)$

The Present Value (PV) period t payoff to ERM_t is $\min[L_t, PV(F_t)]$, where L_t is the present value of the loan assuming it matures in t years and assuming that there is no NNEG involved, and $PV(F_t)$ is the present value of the period t forward contract, for all t. But

(19.6) $$\min[L_t, PV(F_t)] \leq PV(F_t) \text{ for all } t.$$

Therefore

(19.7) $$ERM_t \leq PV(F_t) \text{ for all } t.$$

Hence

(19.8) $$ERM \leq PV(F)$$

which was to be proved.

Principle III

Principle III states:

> The present value of deferred possession of a property should be less than the value of immediate possession

i.e.,

(19.9) Deferment house value < spot house value.

On the Validity of PRA Principle III, i.e., Why Deferment Property Values are Lower than Current Property Values

At the risk of belabouring the obvious (because we must!), we provide a number of alternative demonstrations of the validity of Principle III.

Demonstration #1

Compare the value of two contracts, one giving immediate possession of the property, the other giving deferred possession when exit occurs. The only difference between these contracts is the value of foregone rights (e.g., to rental income or to use of the property) during the deferment period, and the value of these foregone rights should be positive for the residential properties used as collateral for ERMs. It then follows that the present value of deferred possession should be less than the value of immediate possession, i.e., we obtain Principle III.

Principle III thus follows from elementary economics. Why would we not pay less to get less?

Demonstration #2

As an alternative demonstration, recall

(3.10) $$R_t = current\ house\ price\ \times e^{-qt}$$

where q is the *deferment rate* and R_t is the deferment price, and note that the spot house value and the current house price will be equal. Assuming

(19.10) $$q > 0$$

and it is reasonable to assume that (19.10) does hold, then

(19.11) $$R_t = current\ house\ price\ \times e^{-qt} = spot\ house\ value\ \times e^{-qt}$$

which implies

(19.12) $$R_t < spot\ house\ value$$

It is then reasonable to suppose that the deferment house value will be equal to R_t and Principle III follows.

Demonstration #3

A longer and more rigorous demonstration goes as follows:

Let q_0, q_1, q_2, ... be the set of net rental rates for a property from now, period 0, to forever. These net rental services are the use-benefits we get from living in a property (e.g., the benefits of having a roof over our heads) or the rental incomes we could obtain by renting the property out.

Let us assume that these are all positive. After all, zero or negative rental rates do not make much sense.

Let A be the set of those net rental rates q_0, ... for periods 0 to forever.

Let B be the set of net rental rates q_t, q_{t+1} .... from periods t to forever, where $t \geq 1$.

Let C be the set of net rental rates q_0, ... q_{t-1}, for periods 0 to t-1.

Assume for the moment that the prices of A, B and C all exist.

Since the sets of rentals are positive and hence valuable, then the prices of A, B and C should each be positive. By the law of zero arbitrage, the price of A should also be equal to the sum of the prices of B and C. But since the price of C is positive, it must follow that the price of B < the price of A, i.e., the deferment price must be less than the current price and Principle III is established.

To challenge this conclusion, it is necessary to argue that some of these prices do not exist. Since the price of A is the spot price, then the price of A clearly does exist, so one would have to argue that the prices of B and/or C do not exist.

Let's note to begin with that the empirical basis of any such claim is arguable. Whilst it is manifestly obvious that the prices of B and/or C will rarely exist for some specific property, it is often possible to infer proxy prices for different types of property from comparisons of freehold and leasehold prices and it is these proxy prices that one would use for valuation purposes.[80]

For example, consider a leasehold on a London flat with 99 years to run. The price of this leasehold would typically trade at about 95% of the price of a vacant freehold, and the corresponding freehold, i.e. the right to exclusive possession after 99 years, would trade at about 5% of the vacant value, and gives us the price of possession deferred by 99 years.[81]

The following chart shows implied deferment prices – the deferment prices implied by leasehold prices, expressed as a percentage of the freehold vacant possession value[82] – for RICS prime central London 2009 and Leasehold Valuers LLP 2017:

[80] Typically, one would obtain deferment prices for a particular property by applying rules of thumb to the prices for different property types. These would take account of particular features of a property such as location, parking availability, the size of the garden and so on.

[81] See the 'relativity graphs': http://www.graphsofrelativity.co.uk/.

[82] We assume that there is no 'marriage value', i.e. that the sum of the market value of the leasehold and freehold equal the value of vacant possession. For a discussion of this issue, see the next footnote.

Figure 19.1: Empirical Implied Deferment Prices

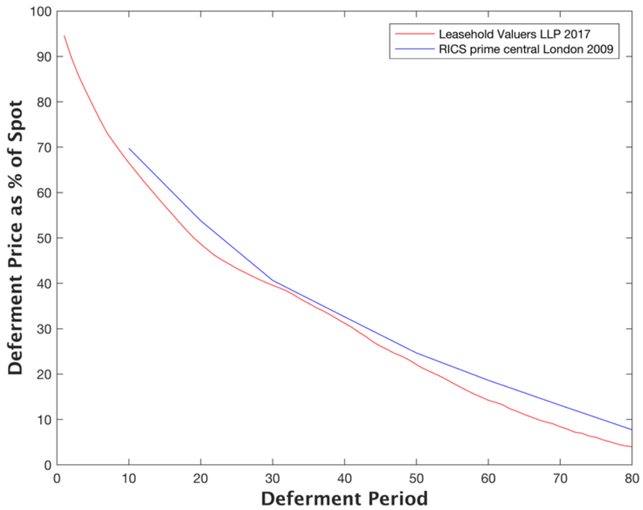

Notes: RICS = Royal Institute of Chartered Surveyors. Sources: http://www.graphsofrelativity.co.uk and Leasehold Valuers LLP.

These fall as the deferment horizon lengthens and are always less than the spot price.[83]

But suppose for the sake of argument and contrary to the evidence just presented that some of these prices do not exist and do not have near approximations or proxies in terms of other market prices. In this situation, we simply switch the metric from prices to values and we can establish the validity of Principle III in much the same way as before. For example, if we assume that each net rental rate has a positive value, then it immediately follows that each of A, B and C has positive value, so the value of B must be less than that of A and Principle III follows. Indeed, even if we assume that the current net rental rate is positive and the others are merely non-negative, then Principle III still follows.

To challenge Principle III, one is then left having to argue that net rental rates or the values of net rental rates are negative.

Let's consider possible examples.

One is where the property and the land which it stands are polluted beyond any feasible repair. Chernobyl comes to mind: even if the land could be restored to a usable state, the costs of doing so would be prohibitive. In this case, all q_0, q_1, ... are negative and will remain so. The property and the land itself would then be abandoned. This type of situation is rare, however.

[83] These results are illustrations only, and other factors come into play. For example: (a) Current leasehold values reflect the right to extend at a market value, whereas ERM borrowers have no such right (the 'lease' ends when they exit into long term care or die, and the estate has no right of extension. (b) We have assumed no marriage value – marriage value is the additional value an interest in land gains when the landlord's and the leaseholder's separate interests are "married" into single ownership (see Law Commission, 2018, p. 23, n. 62) - but adding marriage value would increase the implied q rates further. and (c) ordinary leaseholds tend not to terminate with the property in ruin, whereas there is evidence that very old ERM borrowers tend to neglect their property.

A comparable case would be where the land on which a house is built erodes off a cliff into the sea. But the counter argument is that we would not expect many lenders to be ERMing properties that could fall into the sea after a couple of storms.

A less rare case is where the property is uninhabitable and repair would be uneconomic, but the land itself is valuable. Parts of Detroit come to mind. One might then say that the (current or near current) net rental proceeds were negative, but this situation would not last because the land itself is valuable. The property would be demolished, perhaps after being sold off, and the site redeveloped to restore a positive net rental stream.

A third and more common case is where the property needs repair and repair is economically feasible. The property might not generate any current net rental, but it would be repaired and a positive rental stream restored. This situation is not uncommon, but is still relatively infrequent, in that it does not apply to most properties most of the time.

The general case is that most properties most of the time generate a positive net rental stream. Therefore, when looking for a general rule to assess deferment value, the only sensible rule is to assume a positive net rental stream – and a positive net rental stream implies that the deferment value will be less than the current property value.

In short, if the prices of A, B and C all exist and are positive, then the validity of Principle III follows from zero arbitrage. If any of the prices of A, B and/or C do not exist, however, then we can still obtain Principle III by switching over to a rational valuation argument, in which it suffices to argue that the values of A, B and C are all positive because the underlying rentals have positive value.

Demonstration #4: The fiduciary principle

There is also a normative argument that one can call the 'fiduciary principle'. Even where market prices do not exist, accounting principles say that the accountant *should* value economically similar assets in the same way and imply that valuation should reflect rational investor preferences. The word 'should' or 'ought' appears, e.g., in IFRS 13 B14a: "Cash flows and discount rates *should* reflect the assumptions that market participants would use when pricing the asset or liability." The fiduciary principle says that an accountant or auditor or some other person, who has an obligation of trust towards a less knowledgeable investor, must value an asset or liability as a *rational knowledgeable investor* (or market participant, or knowledgeable, willing independent person) would. This principle provides a safeguard against interested parties coming back along the lines of "no arbitrage doesn't apply here, so we can make up any price that benefits management, other non-fiduciaries or anyone else we choose." Applying this principle, the accountant, actuary etc. must acknowledge that rental services have positive value and this acknowledgement suffices to establish Principle III.

Misconceptions About Principle III

Guy Thomas

Guy Thomas takes issue with Principle III in a recent posting (Thomas, 2018). In his piece, he acknowledges that the loss of foregone rights (e.g., to income or use of the property) during the deferment period [i.e., the argument underlying Principle III] "appears a reasonable argument" but even so, adds that "there are also reasonable counter-arguments." As he put it:

> Housing today is owned mainly by owner-occupiers. They have a preference for a current interest to a deferred interest, because they need a roof over their heads, they like long-term security of occupation, they like being able to make their own choices on extensions and repairs, etc. In other words, they like the practical and sentimental benefits of home ownership. A minority of owners are buy-to-let landlords: they like understandable form of the investment, the unusual ability to finance it largely with borrowed money, and perhaps the disengagement it facilitates from the distrusted pensions and savings industry.

We would put it a little differently. Anyone who lives in a property gets the 'net rental services' of that property – the use-value benefits of a roof over their heads and so forth. Some people choose to obtain those benefits by buying their property and others by renting the property they live in. In the latter case, the property owner gets the benefit of the rent tenants pay, and in most plausible situations, the owner who rents out their property will receive a rent that more than covers the costs of maintaining their property. There are exceptions as we have explained, but these are unusual.

> For an insurer, on the other hand, these practical and sentimental benefits of a current interest in a house have no relevance. The main potential benefit of a current (as opposed to deferred) interest is the potential income from letting.

True, and this point applies to any owner who rents out their property.

> But a current interest also has several disbenefits [sic]: tenants need to be managed, houses need to be maintained, from time to time there are costs (Including possibly PR costs) of evicting tenants in arrears, and there is a possibility (through existing or new legislation) that tenants might acquire new rights.

Yes, there are costs and risks to having tenants.

> If on the other hand houses are kept vacant, this gives another set of problems: council tax, security and maintenance costs, and possibly very considerable PR costs of owning substantial amounts of empty housing.

Yes, there are also costs from keeping properties in vacant possession.

> These disbenefits are not fanciful; their materiality can be inferred from the observable fact that despite the excellent long-term performance of housing as an investment, neither insurers nor any other financial institutions have shown any enthusiasm over the past several decades for housing as an asset class.

These passages are a roundabout way of saying that there are benefits and costs of owning property but if an owner regards the costs as outweighing the benefits, then the sensible choice for the owner is to sell. The property will then end up in the hands of an owner who *does* value the benefits as more than the costs – otherwise they wouldn't have bought the property and *someone else* would.

The lack of enthusiasm (or otherwise) of financial institutions for housing as an asset class is another question. He continues:

> So current interests in houses are evidently not attractive to insurers and other institutional investors. Deferred interest might well be more attractive, particularly if in the form of cash-settled financial contracts, so that all the problems of current interests are permanently avoided. Even if a deferred interest is not strictly preferred, *the relative valuation of a deferred interest compared to a current interest seems very likely to be much higher for an insurer than a typical individual owner*. (Our emphasis)
>
> Now if there were a substantial market for deferred interests, the money weight of individuals' preference for current interests versus insurers' preference for deferred interests would determine the relative *market prices* for the two types of interest (i.e. what the PRA calls the 'deferment rate'). But we have the same problem as with the hedging arguments: *the market for deferred interests does not exist on any meaningful scale*. (Our emphasis)

Leaving aside that a market for deferred interests does exist (see above), Thomas is comparing one hypothetical non-market valuation (i.e., insurers' valuations of current possession) against another (i.e., their valuations of deferred possession). A comparison of the relative valuations of spot and deferred possession made by a party that is *ex hypothesi* not a major player in the market does not establish (a) anything about the market prices or plausible values for current possession or the market prices or plausible values for deferred possession or any relationship between them. In any case, no such comparison establishes (b) that deferred, forward or future 'interests' have the negative value necessary to undermine the validity of Principle III.

To make point (a) in a different context, suppose we value a typical stately home as being worth 2 times the value of a typical castle, but the market values a typical stately home as being worth 3 times the value of a typical castle. Our views might be sincerely held, but they are of no relevance if we don't have any portfolios of castles or stately homes and are not in the market trading them. Because we are not in the market trading these things, our views about their relative valuations have no relevance to anyone but ourselves. The only valuations that matter are those of the market.

Or to give another comparison. We might believe that Hollywood movies are rubbish and Bollywood movies are great, or the other way round. Makes no difference. The only inferences we could reasonably draw from that expression of relative preferences is that we would never go to see the rubbish movies. But a lot of people do go to see the rubbish movies because they don't share our good taste, so the rubbish movies have economic value. Therefore we might rationally invest in them, even if we would never watch them. The point is not to confuse subjective taste with economic value, and what is the best estimate of the economic value? It is what the market will pay.

In any case, no such comparisons establish that deferred, forward or future 'interests' have the negative value necessary to undermine the validity of Principle III.

Thomas's argument also runs afoul of the fiduciary principle: even if an actuary or auditor has a private view of the relative valuations of deferred and current possession, they are still required to report the market valuations, and only those valuations.

In short, the validity of Principle III can be buttressed by sound economic theory, solid accounting principles and abundant empirical evidence, but the counter arguments cannot.

Radu Tunaru on Principle III

Tunaru also challenges Principle III. In his report he argues (p. 50) that the "prepaid forward price", i.e. the deferment price, will always be lower than the current house price (i.e. the freehold price of vacant possession):

> the idea that the prepaid forward price should always be lower than the current house price can be challenged. This condition will work obviously in normal market conditions and for shorter maturities. However, in the aftermath of financial and economic crises, conditions may be reversed. For example, in the aftermath of the subprime crisis house prices dropped significantly. The usual question of "how much should you pay to get a house in five or ten years time?" should be replaced with the question "what price can you get on the market to sign now for possession of a house in five or ten years time?" Even if the house price market was depressed in the aftermath of the subprime crisis, the expectation of the market would naturally be that the market will recover after some time and the forward curve will be in contango. Thus, it is possible that the market will require a prepaid forward that is higher than the current house price.

This argument is not theoretically defensible and we are not aware of any evidence for it. From a theoretical point of view, why would a rational investor who expected house prices to recover pay more for a contract for possession in (e.g.) 5 years time than they would pay for a contract for vacant possession? The prices of both contracts will be identical after 5 years, so in both cases the investor will benefit from the desired recovery. However with vacant possession, the investor will be able to use the property or rent it out, whereas with deferred possession there is no such benefit. The rational preference would be for vacant possession, and as we noted above, an accountant who was valuing the contracts for the benefit of shareholders is professionally obligated to reflect that rational preference.

From an empirical point of view, we know of no evidence suggesting that long dated freeholds in the London market were trading above the vacant possession price in the aftermath of the 2008 collapse. A 99 year lease on a London apartment would typically trade at 95-99% of the vacant freehold value. We have no evidence that leasehold values collapsed to zero, or went negative, during the crisis. Perhaps such evidence exists, but given the theoretical unlikelihood, the onus should be on Tunaru to produce it: strong claims require strong evidence. He points (p.51) to the fact that short-dated EUREX futures contracts traded with negative deferment rates for short a period following the crisis, but these refer to futures not forward, and do not refer to contracts for long-dated deferred possession such as in the freehold and leasehold markets, which are the forwards relevant to the case of equity release.

Consider also evidence such as the following from a Daily Mail article published on 24 July 2009.[84] The article mentions ‘the recent large price falls in the housing market’ and then points out:

> “Recent price reductions of up to 25 per cent in most areas mean an equivalent *drop* in the cost of extending your lease or buying the freehold,” says Angus Fanshawe, of London estate agency Douglas & Gordon.

But if Tunaru were right, the freehold (or the extension) should be getting more expensive, not less, given market expectations of price rises in the future, no?

> At the very posh end of the market, an apartment in Chelsea's exclusive Cadogan Gardens has a 43-year lease and is on sale for £2.95million (020 7581 3022, friend and falcke.co.uk), but may be worth up to £4.5 million if the lease were extended to 99 or 125 years.

So the 43 year leasehold has a positive value, not the negative one implied by a negative deferment rate.

> … a flat in Central London with an 88-year lease would sell for 97 per cent of its full market price but the same flat with 22 years on the lease would sell for only 56 per cent and with a ten-year lease for 32 per cent.

So if the deferment rate were negative, we could buy a flat with vacant possession for 100%, carve out an 88 year lease and sell for 97%, then sell the remaining freehold for possession after 88 years at over 100%. That would be a nice trade!

All the relevant evidence indicates that deferment rates are positive as Principle III maintains.

[84] https://www.dailymail.co.uk/property/article-1201766/Time-new-lease-life-Freehold-isnt-way-secure-future-profits.html

IFoA Misconceptions on Principle III

In June 2016, the Institute and Faculty of Actuaries issued "DP 1/16: Equity Release Mortgages: IFoA Response to the Prudential Regulation Authority," its official response to the PRA's earlier Discussion Paper DP 1/16, which had had asked for industry views on ERMs. To quote from this response:

> 33. For the second relationship in paragraph 4.9 [i.e., Principle III] to hold, in theory, there needs to be a deep and liquid market. Otherwise the implication is that the average value of the HPI [House Price Inflation] assumption is less than or equal to the discount rate assumed in the valuation of the NNEG. In practice, the approach to setting the HPI assumption varies significantly from firm to firm.

There are several howlers here:

- Mistake #1 is that for Principle III "to hold, in theory, there needs to be a deep and liquid market." The validity of Principle III has nothing to do with a deep and liquid market and we have just shown that its validity holds under general conditions.
- Mistake #2 is to suggest that the "average value of the HPI assumption is less than or equal to the discount rate assumed in the valuation of the NNEG." This statement is just plain wrong. The correct statement is that we can assume any HPI we want to, but the assumed value of the HPI is always irrelevant to the valuation of the NNEG.

Para 35 then gives some illustrations of circumstances in which Principle III allegedly might not hold:

- One is the claim that Principle III "is a statement of 'value' and applies to any individual. However this is not necessarily true in terms of the exchange value." This strange statement is an imaginative addition to the economic theory of value but is unfortunately also wrong. The claim that the Principle III "is a statement of value and applies to any individual" is true, but the corollary is that it also applies to all individuals *including* (and not excluding!) when they engage in trade at market or exchange values.
- Another is the claim that "in a negative yield curve scenario, the relationship (Principle III) would fail as the premise that deferral could lead to a lower present value no longer holds." This statement is a head scratcher but one can see that it must be wrong because the deferment price (or value, makes no difference here) is equal to $S_0\, e^{-qt}$ and this expression does not include any interest rate or yield, negative or not. To repeat, Principle III depends only on the q rates being positive (or mostly positive) and it is difficult to imagine plausible situations where that would not be the case.

So how come the distinguished actuaries of the IFoA could make such mistakes? A possible clue is that the covering letter opens with the following statement:

> The IFoA's *Equity Release Members Interest Group* (ER MIG) and Life Board have been involved in the drafting of this response. The contributors to this response *include members who are actively engaged with use of equity release assets by life insurers*. (My italics)

The IFoA had allowed itself to be used as a mouthpiece for ERM industry practitioners to broadcast their misunderstanding of their products.

But the authors of the IFoA official response to DP 1/16 are not alone in misunderstanding these principles. Consider these passages from a recent Deloitte communiqué on ERMs:

> In our view, the third principle (that future possession of a property cannot be more valuable than current possession) is likely to attract the most future debate.

But Principle III is just elementary economics!

> Very importantly, this principle implies that assumed future house price growth cannot exceed the discount rate applied in the valuation. …

No it does not.

> The PRA expects there to be a positive value associated with possession of a property.

Yes, obviously.

> The practical implication of this is that the assumed house price growth within the NNEG option pricing calculation cannot exceed the discount rate, as this would imply that future possession is more valuable.
>
> This principle therefore effectively sets a cap on firms' house price growth assumptions.

These statements are nonsense. Principle III has *no* implications about assumed future house price growth. You can make any assumptions about future house price growth that you like and Principle III would be still be valid.

> We would expect firms investing in ERMs and other direct investments to see an increased level of scrutiny and questioning from the PRA, with the *bar set very high for management's understanding of the valuation of such investments*. (Bulley *et alia*, 2017, our italics)

They are clearly off to a flying start on that one.

The lead author, Andrew Bulley, is a partner in Deloitte's Centre for Regulatory Strategy. Prior to joining Deloitte, Mr. Bulley was Director of Insurance Supervision at the Bank of England.

To challenge Principle III is thus to make an egregious intellectual error and it is remarkable that the IFoA has not only failed to condemn any such challenge but has explicitly given it its imprimatur. This situation is analogous to the UK's top mathematical institute, the Institute of Mathematics and its Applications, taking the official view that the validity of 2+2=4 is an opinion. You see, some mathematicians are of the opinion that 2+2=4 but others have a different view, including some who speak for the Institute.

We appear to have here another case of 'actuarial judgment' gone awry.

We are reminded of some comments made on this subject by Tim Gordon almost two decades ago (Gordon, 1999). He wrote (p. 4) about the actuarial conviction that "actuarial judgment is the only technique for valuing long-term liabilities" but 'actuarial judgement' produces an answer that "varies enormously depending on which actuary carries out the calculation." He continued:

> actuaries assume that judgmental methods are the only methods available which give sensible answers. What is more, the judgement involved is something which apparently only comes with years of experience. In other words, we claim to know the answer but cannot tell anyone else how to derive it in advance.

The experienced actuary knows it when they see it. Roman augurs had the same skill reading chicken entrails. As he continued further:

> The problem is that the difference that actuarial judgement can make to valuations using the traditional approach is enormous. It means that:
>
> • we are exposed to pressure from clients seeking to move answers in the direction which favours them, and
> • we lose credibility because we are unable to explain precisely how we arrive at an answer.

Actuarial judgment can also lose credibility when it produces answers that are demonstrably wrong.

Bounds on ERM and NNEG Valuations

To return to the main storyline, the impact of these two Principles is illustrated in Figure 19.2:

Figure 19.2: Illustration of Principles II and III

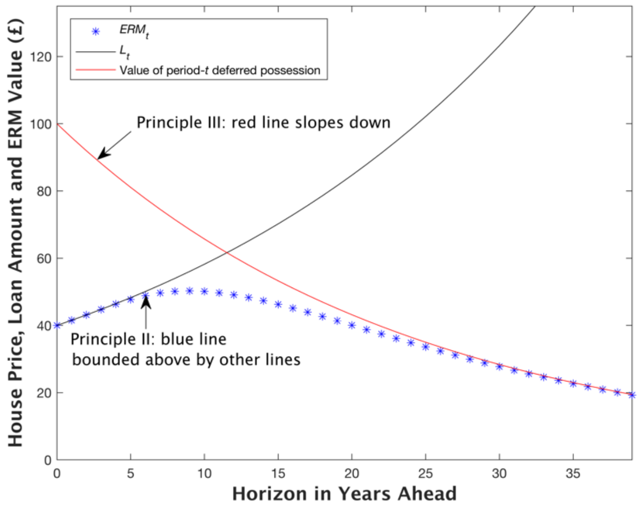

Notes: Based on the baseline assumptions: male aged 70, LTV=40%, r=1.5%, l=5.25%, q=4.2% and σ=14.8%. Exit probabilities are based on M5-CBD model projections using England & Wales male deaths rate data spanning years 1971:2017 and ages 55:89.

Principle II implies that the blue (ERM_t) line must be below both the green (L_t) line and the red (deferred possession) line, and Principle III implies that the red (deferred possession) line should slope downwards.

There is some interesting intuition underlying the Figure:

- For very low horizons, $NNEG_t$ is very out of the money and probability of exercise is very low. Hence the value of the option will be negligible and ERM_t will be indistinguishably close to the value of the loan L_t.
- For long horizons or high t, the option is well into the money and the probability of exercise is high and approaching 1. Therefore, the ERM_t line converges to the deferred house value line for period t.

Underlying these graphs are some elegant mathematics. ERM_t is given by

(19.13)
$$\begin{aligned} ERM_t &= e^{-rt}L_t - e^{-rt}[L_t N(-d_2) - F_t N(-d_1)] \\ &= e^{-rt}[1 - N(-d_2)]L_t + N(-d_1)e^{-rt}F_t \end{aligned}$$

where we have set the deferment price $D_t = e^{-rt}F_t$. From the standard equivalence $N(-x) = 1 - N(x)$, we then get

(19.14)
$$ERM_t = N(d_2)e^{-rt}L_t + N(-d_1)e^{-rt}F_t$$

This expression is simpler and reflects the shapes of the curves clearly. As d_2 gets positive, $-d_1$ gets negative, so $N(d_2)$ goes to $1, N(-d_1)$ goes to zero and ERM_t approaches the present value of the loan. As d_2 goes negative, it's the other way round, so the term on the left disappears and the term on the right approaches the deferment value $e^{-rt}F_t$. One sees these bounds at play in Figure 19.2.

Besides their mathematical elegance, these bounds implied by Principles II and III have a helpful practical use: they are an easily calculated cross-check on any proposed ERM or NNEG valuation. Consider Figure 19.3, which shows the upper bound for ERM_t made explicit and highlighted in blue.

Figure 19.3: ERM Upper Bound

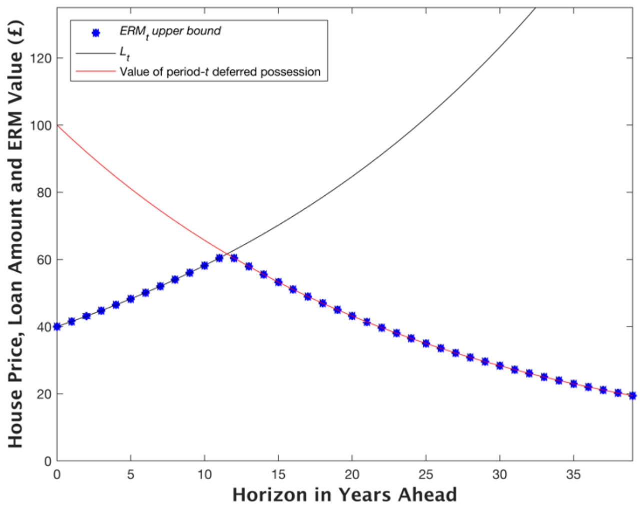

Notes: As per Figure 19.2

As an aside, if we start with a figure like Figure 19.2 and let the volatility get small, then it is easy to show that Figure turns into Figure 19.3 and the message is that the ERM valuation approaches the Principle II upper bound. But if the ERM valuation approaches its upper bound, then the corresponding Black 76' option valuation must approach the Principle II NNEG valuation lower bound, i.e., as $\sigma \to 0$, the Black '76 NNEG valuation approaches the Principle II NNEG lower bound.

We can obtain the ERM_t upper bound as the minimum of $e^{-rt}L_t$ and $e^{-rt}F_t$. Note that this upper bound can be estimated using only information about the current house price and LTV (which together give us the current amount loaned), the risk-free rate r, the net rental q, the loan rate l and the exit probabilities. For example, in the baseline case, we estimate the ERM upper bound to be £46.7, which compares to our earlier baseline estimate of ERM as £43.5. So even without estimating ERM or its $NNEG$ *or estimating any underlying option model or calibrating any additional parameters, such as the volatility*, we immediately know that any proposed value of ERM that exceeds £46.7 must be wrong.

But if we can estimate an upper bound for ERM without requiring an option-pricing model or relying on any volatility parameters, then by

(3.2) $$ERM = L - NNEG$$

we can also estimate a lower bound for $NNEG$ on the same basis. Given that L = £74.9 in our baseline case, the upper bound ERM estimate of £46.7 implies a $NNEG$ lower bound equal to £28.3. This lower bound compares to our earlier $NNEG$ estimate of £31.4. So *even without estimating the $NNEG$ or relying on any $NNEG$ valuation model or any* *volatility estimate that might go into any such model*, we know that any proposed NNEG value below £28.3 must be wrong.

To cut to the chase, given these various inputs – the assumed age and gender, the assumed house price and LTV, the assumed r, q, and l rates, and the inputted exit probabilities – it is *impossible* to get a NNEG value any lower than £28.3 *whatever option pricing model one might use and regardless of how it might otherwise be calibrated*.

At the risk of repeating ourselves, we would stress that this lower bound NNEG value is *not* dependent on Black '76. The recent Institute reply to CP 13/18 released on 28 Sep 2018 made a great deal of noise about how autocorrelation, mean reversion, lack of Geometric Brownian Motion and so forth undermined the validity of Black '76, and a number of participants at the LSE seminar on 1 October 2018 made similar points. We would dispute the validity of these claims – not least because they often confuse sufficient with necessary conditions for Black-Scholes type valuations to be valid, and we have more to say on these issues in Chapter 20 – but even if these claims were all valid, they do not apply to the bounds-based valuation offered here, because that argument is not dependent on any option pricing at all, Black '76 or otherwise.

We also have here a handy cross-check of any proposed NNEG valuation: if any proposed model based on the same input calibrations gives a lower NNEG valuation than either of the NNEG lower bounds, then it must be wrong.

Black' 76 vs. Principles-based bounds results

It is interesting to compare our baseline NNEG valuation results with the results we would have obtained had we dispensed with the option pricing model and used the Principles-based bounds instead:

Table 19.1: Baseline ERM/NNEG vs. Principle II Bounds Valuations

Current House Price	**Loan Amount**	$\boldsymbol{L}$	$\boldsymbol{NNEG}$	$\boldsymbol{ERM}$
£100	£40	£74.8	£32.2	£42.7
Current House Price	**Loan Amount**	$\boldsymbol{L}$	$\boldsymbol{NNEG}$ **lower bound**	$\boldsymbol{ERM}$ **upper bound**
£100	£40	£74.8	£28.1	£46.8

Notes: As per Figure 19.2.

These results indicate that the basic NNEG under-valuation story obtained earlier using Black' 76 *still holds true if we use the Principles-based bounds instead of any option pricing model*.

Thus, one cannot dismiss the NNEG under-valuation story based on arguments – right or wrong makes no difference – about the validity of Black '76 applied to NNEG valuation.

Table 19.2 gives the corresponding results for the Principle III bounds:

Table 19.2: Baseline ERM/NNEG vs. Principle III Bounds Valuations

Current House Price	Loan Amount	L	$NNEG$	ERM
£100	£40	£74.8	£32.2	£42.7
Current House Price	**Loan Amount**	L	$NNEG$ **lower bound**	ERM **upper bound**
£100	£40	£74.8	£13.2	£61.7

Notes: As per Figure 19.2.

Figure 19.4 gives plots of the Principle II NNEG lower bound against deferment rate and age, where the latter assumes that LTV follows the 'age minus 30' rule:

Figure 19.4: *NNEG* and NNEG Principle II Lower Bounds vs Deferment Rate

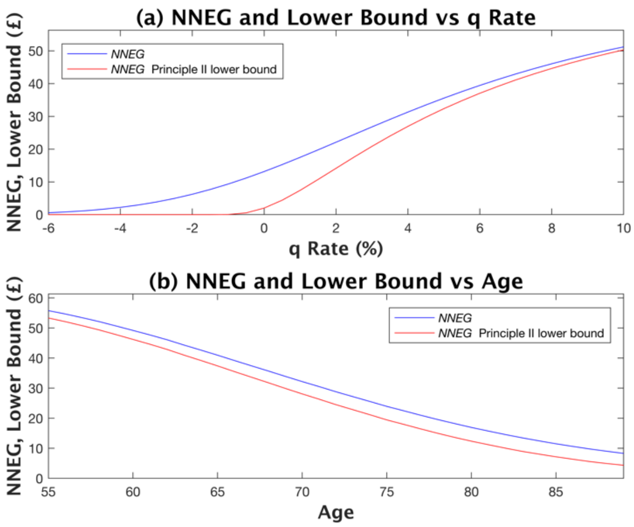

Notes: Otherwise as per Figure 19.2.

Figure 19.5 shows the corresponding ERM upper bounds:

Figure 19.5: *ERM* and ERM Principle II Upper Bounds vs Borrower Age

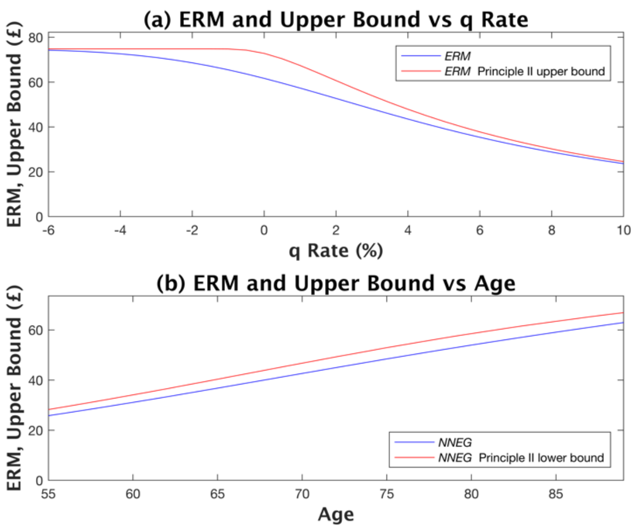

Notes: LTV follows the 'age minus 30' rule. Otherwise as per Figure 19.2.

Figures 19.6 and 19.7 show the corresponding plots for the Principle III bounds:

Figure 19.6: *NNEG* and NNEG Principle III Lower Bounds vs Deferment Rate

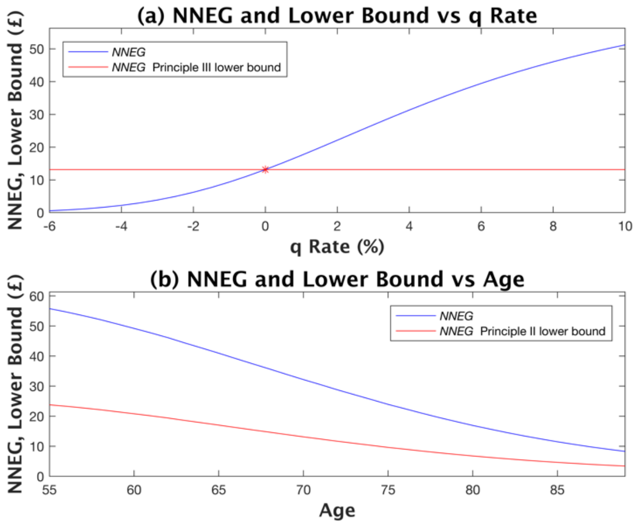

Notes: Otherwise as per Figure 19.2.

Figure 19.7: *ERM* and ERM Principle III Upper Bounds vs Borrower Age

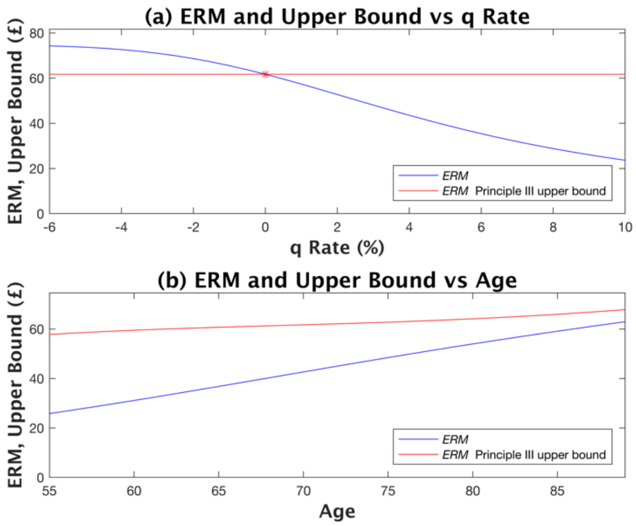

Notes: LTV follows the 'age minus 30' rule. Otherwise as per Figure 19.2.

We see that for Principle II, the bounds are quite close to the Black '76 valuations, especially in the plausible range of deferment rates. The Principle III bounds are less close.

Chapter Twenty: The Market Consistent Approach

Introduction

A 'market consistent' (MC) approach to valuation is one which gives 'market consistent' valuations, and a 'market consistent' valuation can be defined as a 'fair value' valuation based on the IFRS definition of a fair value price, namely:

> The price that would be received to sell an asset or paid to transfer a liability in an orderly transaction between market participants at the measurement date.[85]

Although it is often identified with Black' 76, the MC approach does not mandate the use of any particular option pricing model. For example, instead of Black' 76, the MC approach might be applied with an option pricing model that allows for, e.g., transactions costs or with a Monte Carlo (or stochastic) simulation model, provided (a big if) that the model is fit for purpose and that the calibrations are reasonable. The essence of the approach is market consistency, not adherence to any particular model. Any model is allowable, provided it meets the standards for market consistency.

The MC approach is sometimes called a 'risk neutral' approach in the actuarial literature. This alternative label reflects the underlying derivative pricing methodology based on risk-neutral probabilities. However, this latter term is unfortunate, in that it suggests to those who do not understand this methodology that it must be assuming that *investors* are actually neutral about risk and who would want to use a method that made a questionable assumption like that? Indeed, the argument that the 'risk-neutral' approach is flawed because investors are not risk-neutral is probably the most common misconception one encounters against the approach, as if to suggest we should want some alternative approach that does not make this assumption. But we cannot stress often enough that the market consistent or 'risk neutral' approach DOES NOT assume that investors are risk-neutral. The truth of the matter is that the methodology works regardless of whether investors are risk-neutral or not; their attitude towards risk is an entirely separate matter. The term 'risk neutral' has an unwarranted negative connotation when used in broader circles and is best avoided.

The MC approach is also subject to other misconceptions that have much confused actuarial discussions. These arguments claim that you can't use Black '76 or Black-Scholes because: there are no forward contracts or you can't do forward contracts, you can't short the underlying, house prices are autocorrelated rather than GBM, markets are illiquid or incomplete, and, more generally, the assumptions underlying the derivation of Black-Scholes family models do not hold empirically. If these assumptions don't hold, so it is claimed, then these models will be unreliable. But this argument is logically flawed. Yes, it is it true that in its usual classical/textbook derivations, Black-Scholes or Black '76 make a whole bunch of assumptions, some of which are empirically false. However, these assumptions are sufficient rather than necessary for the model's results to be valid, so

[85] See, e.g., https://www.iasplus.com/en/standards/ifrs/ifrs13.

you can't just kick aside some of these assumptions and conclude that Black' 76 results must be wrong. To prove that those results are wrong, it is necessary to kick aside assumptions that are *necessary* for the results to hold, not conclusions that are *sufficient* for them to hold, and that is a much more difficult task.

Furthermore, to demonstrate that these models give the 'wrong' answers, one must also demonstrate that some necessary condition is not only 'wrong', but also introduces a material error to the resulting valuations. Even then, a model can still be useful. A case in point is Black '76. This model gives us put valuations are biased, but still useful, because it can put a lower bound under the put valuation, i.e., so if the model tells us that the put value is £X, then we know that the true put value must be *at least* £X.

We often encounter criticisms of Black '76 made by actuaries who think that if they can, for want of a better way of putting it, knock the model off its perch then they can thereby justify lower NNEG valuations as if by default to an alternative that produces lower valuations. This view is seriously mistaken. If Black '76 does not hold, then the valid alternative is not the discounted projection or Tunaru approaches (and more on these in later chapters) but some other MC-consistent model that will deliver *higher* NNEG valuations than Black '76. We suspect that many actuarial critics of Black '76 would be a lot less critical if they understood this point.

This point is a recurring theme in the literature that addresses option pricing in the presence of transactions costs. The general result from that literature is that transactions costs increase option valuations above BS valuations or, in some cases, increase the volatility that needs to be inputted into the BS model, which amounts to the same thing. That literature also teaches us that in such circumstances we must take account of utility or risk preferences which, in turn, make for a range of correct valuations, i.e., there is no uniquely correct solution.[86]

In truth, we shouldn't put any model on a pedestal, whether that model be Black '76 or some superior alternative, let alone an inferior one. What *should* go on the pedestal is the pricing methodology. We should not be thinking, "We use Black '76 because it is the valid model." Instead, we should be thinking, "How do we apply the correct (that is, zero arbitrage rehedging) pricing methodology in the ERM context?"

We then come to the fundamental methodological issue that should be at the forefront of our thinking: we should spell out how the rehedging-based pricing methodology would work when the underlying is property.

Put Pricing from First Principles

Put pricing under continuity

---

[86] See, e.g., H. E. Leland (1983) "Option Pricing and Replication with Transactions Costs," *Journal of Finance* 40: 1283-1301; S. D. Hodges and A. Neuberger (1993) "Optimal Replication of Contingent Claims under Transactions Costs," *Review of Futures Markets* 8: 222-239; or G. Barles and H. M. Soner (1998) "Option Pricing with Transactions Costs and a Nonlinear Black-Scholes Equation," *Finance and Stochastics* 2: 369-397.

Instead of assuming that Black '76 holds, let's first assume that we are working in the idealised theoretical world assumed by Black-Scholes, which is the same idealised theoretical world as that assumed by Black '76. We wish to value a European put option on a forward contract on some asset and let's suppose that we haven't even heard of the Black '76 put option pricing formula. So what do we do?

The answer is that we construct a synthetic put, a series of dynamically rebalanced positions in the underlying contract which replicates the payoff of the option contract at expiry. The underlying contract can be a forward contract in which we pay at maturity to take possession of the asset at contract maturity or a deferment contract in which we pay now to receive the same asset at maturity.

As the price of the underlying moves around, we continuously rebalance our synthetic put, to ensure that the synthetic put always has the same expected payoff as the 'real' put, and we continue rebalancing until both puts expire.[87] The value of the synthetic option is then the payoff of the contract (for a put, the difference between the strike price and the underlying, if the latter is lower), minus the cost of hedging. The value of the synthetic will be the same as the value given by Black '76 if the hedging is continuous (which is impossible to carry out in practice). So whatever the payoff, and whatever path the underlying asset follows until expiry, the payoff minus the hedging cost will yield approximately the same amount, namely the theoretical Black '76 option premium.[88] We know that the synthetic put value must equal the Black '76 put value because this pricing process is the same one that Fischer Black used to obtain his Black '76 put pricing equation, including the impossible condition that the hedging is continuous.

Put pricing under discontinuity

The assumption of continuous trading (like many other textbook assumptions) is of course purely theoretical, because it applies only at the limit as $n \to \infty$. In practice, only discrete hedging is possible. It then turns out that the value of the synthetic option under discrete hedging does depend on the path of the asset to expiry. To illustrate this dependency, we simulated the value of a discretely hedged at-the-money put option for 20,000 different paths, expressing the difference (or hedging error, HE) between the discrete (H76) and the continuous (Black 76) price as a percentage of the continuous price, i.e.,

(20.1) $$HE(n) = (H76 - B76)/B76$$

[87] The standard textbook account is that we construct our position from a combination of risk-free bonds and the underlying. An institutional desk would take a slightly different approach. If the option is sold, the desk will receive a premium which it invests with the firm's own treasury department. It will receive interest based on the firm's own funding rate which reflects its external rating, less a spread taken as profit by the treasury. If the option is long dated, the desk may also hedge the interest rate risk by a swap with its own swap desk. If the underlying is a futures contract, then the cash position would need to be cash settled, i.e., rebalanced to reflect interim profit and loss on the underlying. If the underlying is a forward, no such rebalancing is necessary.
[88] We implicitly assume that the volatility is approximately constant.

where n is the number of times that we hedge (so we rehedge $n-1$ times). We assume a time to expiry of 1 year and make other assumptions, including an underlying continuous volatility of 13%. So if $n = 52$, we recalibrate every week, if $n = 250$, we recalibrate every day, assuming 250 trading days to a year, and so forth.

The results are shown in Figure 20.1:

Figure 20.1: Rehedging Error and Hedging Frequency

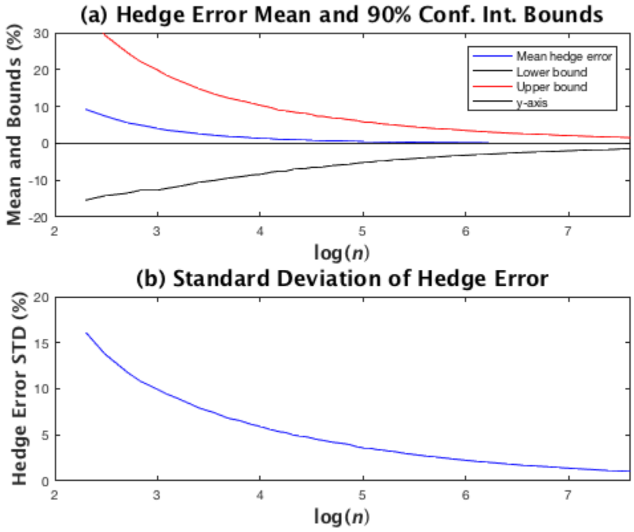

Notes: Based on 20,000 simulation trials, $F = 100$, $K = 100$, $r = 0$, $t = 1$ and $\sigma = 13\%$.

As we can see, the average hedging error, the standard deviation of the hedging error and the confidence bounds around the hedging error converge to zero as n gets large, as option theory suggests. Some specific results are shown in Table 20.1:

Table 20.1: Rehedging Error and Hedging Frequency

n	50	100	250	500	1,000
Average HE	1.54%	0.69%	0.32%	0.12%	0.04%
Std HE	6.15%	4.41%	2.85%	2.05%	1.45%
Lower bound	-8.62%	-6.29%	-4.06%	-3.01%	-2.15%
Upper bound	10.92%	7.32%	7.32%	3.09%	2.16%

Notes: As per Figure 21.1. 'HE' = hedge error. 'Lower bound' = lower bound of 90% confidence interval. 'Upper bound' = upper bound of 90% confidence interval.

The choice of n is a subjective judgment. With $n = 50$, i.e. rebalancing about weekly, the standard deviation is 6.15% of the at-the-money price. Given that an option trader would typically quote a spread on implied volatility, he could adjust the spread to improve the chances of making a profit on the trade. For example, if the trader estimates volatility at 13%, he will quote a 2% bid offer spread in the market, meaning he buys at 12%, sells at 14%. With an ATM option, one year to expiry, that equates to a bid-offer spread of 4.78 – 5.58. The difference from the mid price of 5.18 expressed as a percentage, is thus about 7%. Given that the standard deviation of error is around that number, it follows that the trader has an 84% probability of making a profit, just for that one trade. Alternatively with daily hedging, n would be around 250, implying a hedging error of about 2.84%, in which case the trader would quote a correspondingly lower bid-offer spread.

These results can easily be adjusted for different times to expiry. The two relevant variables are (i) the number of hedging periods and (ii) the volatility of the asset *at the hedging period*. For example, suppose we are pricing a 20 year option, with a chosen hedging period of 1 month, and suppose that the volatility of 1 month returns is 1%. (Note that we are *not* annualising the volatility here. We mean that the standard deviation of monthly returns is 1%.) Then we can pretend we have a 1 year option divided into $20 \times 12 = 240$ periods, and that the annualised volatility of the asset is $1\% \times 240^{0.5} \approx 15.5\%$. This $n = 240$ case is close to the $n = 250$ example above, implying a standard deviation of error of about 2.5%.

The hedging period is crucial not just because of the pricing error, but also because of the Hurst (or autocorrelation) effect that we noted in Chapter 10. If the underlying is autocorrelated – another departure from Black-Scholes – then the longer the period between rebalancing, the greater the impact of the Hurst effect, i.e. the greater the realised volatility compared with the theoretical volatility implied by the square root rule. For example, suppose our 20 year option is on an underlying asset which is highly autocorrelated, e.g. with $H = 0.9$, and that we are rebalancing every 5 years, rather than every month. The standard (GBM) volatility correction formula assumes

(20.2) $$\text{adjusted } \sigma = 1\% \times 60^{0.5} = 7.75\%.$$

However with the Hurst adjustment, the result is very different:

(20.3) $$H\text{-adjusted } \sigma = 1\% \times 60^{0.9} = 39.8\%.$$

Taking account of the Hurst effect (i.e., taking account of autocorrelation in the underlying) will then increase both the cost of the option and the standard deviation of the hedging error. Thus, this particular departure from the Black-Scholes world will lead to a higher option price than that given by Black '76.[89]

Dynamic Replication

It is frequently claimed that without a complete liquid hedging asset, we cannot dynamically replicate an option and so cannot derive the Black-Scholes formula. To quote the PRA agreeing with unnamed respondents who raise these objections, "The PRA agrees with the respondents that the attributes of the residential property market do not permit the derivation of the Black-Scholes formula via dynamic replication arguments" (PS 31/18 2.28, p.9).

Taken strictly and literally, this claim is false. The Black-Scholes formula is a derivation from first principles, just as the Pythagorean theorem is a derivation from first principles, and it is always possible to *derive* it.

[89] We would emphasise however that not all departures from a Black-Scholes world will produce put values that differ from the Black '76 put value. For examples, see D. Buckner "Weird distributions #1." *The Eumeaus Project* 17 October 2018; and "Weird distributions #2." *The Eumeaus Project* 17 October 2018.

Clearly something else is meant, but what? Is the claim that the trader cannot successfully make a market in an option that is either (i) impossible or (ii) expensive to hedge? In response to the first case, it is *always* possible to hedge an option. Just because there does not exist a listed market, does not mean that we can't value NNEGs and ERMs using some market hedging approach. Forgive the triple negative. Our point is that hedging possibilities will always exist. One possibility is an external hedge, i.e., to dynamically hedge the ERM in the OTC (over the counter) market. For each maturity and every so often, the ERM provider could estimate the number of properties expected to exit at that maturity, the average price of such properties, then agree to deliver properties of an equivalent standard to an external hedge provider. Most investment banks make a business of such arrangements for a price. There are non-trivial implementation issues, but they are manageable ones, and would be less difficult than those involved in, say, pension buyouts or buyins, in which there has been a flourishing market for over a decade.

As for the second, i.e. expensive to hedge case, of course the price of the hedge might be high, but a high price argues for an equivalently high valuation for the NNEG, relative to the valuation we might obtain using a volatility based on observed market mid-prices.

Or is the objection being made that we cannot use Black or Black-Scholes as an *accounting* basis for the option? There are two replies to this objection. The first, which we shall address in this section, is that we can use first principles to establish an upper bound for the value of the ERM, or a corresponding lower bound for the NNEG. The second, which we shall address in a subsequent section, is that we can use the standard option model to provide a P/L explanation for the payoff of the option, *even if the option is unhedged*.

Recall how we established in Chapter 19 that we do not require dynamic hedging arguments or complete liquid markets or an option pricing model or even a volatility estimate to establish upper bounds value for the ERM and corresponding lower bound valuations for the NNEG. Simple arguments suffice as per the PRA's Principles II and III. To recap, if a freehold, i.e. a contract for deferred possession at a certain maturity, is on the market for a certain price, then a rational investor would not pay more than that price for an ERM at that same maturity, and accounting principles require the firm to reflect the considerations that a rational investor would make (see Chapter 26 for more on the accounting issues). Similarly, a rational investor would not pay more for deferred possession than for spot. So we have lower bounds on the put value and these lower bounds correspond to the values we would get under Black '76 for limiting cases of the inputs, i.e., where $\sigma \longrightarrow 0$ and $q \longrightarrow 0$. The former limiting case is the lower bound under Principle II and the latter limiting case is the lower bound under Principle III.[90]

The impact of these bounds was shown in Figure 10.3. Figure 20.2 gives the same figure but with the Black '76 ERM_t based on $\sigma = 3\%$ superimposed on it.

[90] We would also point out that the PRA's principles (Principles II and III) are consistent with the standard option pricing model, and, indeed they underlie it. Principles II is about the minimum-of(A,B) condition that underlies all option models, whereas Principles III is about how the forward price is the backbone of the option.

Figure 20.2: Black '76 ERM Valuation and Principle II ERM Upper Bound

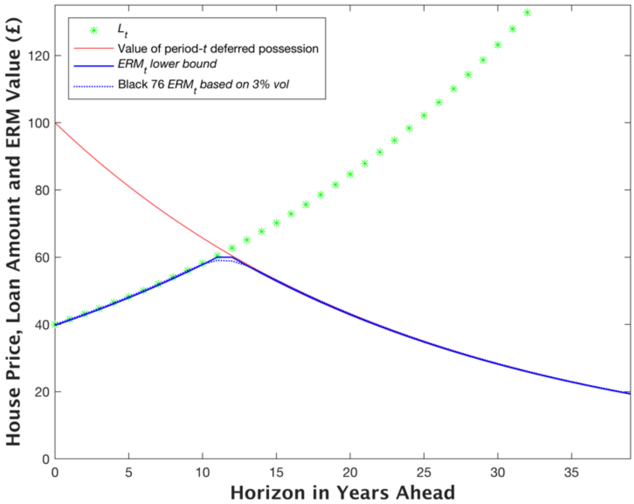

Notes: L_t and ERM_t are the loan value and ERM values for decrement t. Based on the baseline assumptions: male aged 70, LTV=40%, r=1.5%, l=5.25%, q=4.2% and σ=3%. Exit probabilities are based on M5-CBD model projections using England & Wales male deaths rate data spanning years 1971:2017 and ages 55:89.

As the volatility gets small, the ERM_t curve approaches the ERM upper bound and $NNEG_t$ approaches its lower bound.

Table 20.2 shows the Black '76 NNEG and ERM valuations compared to the Principle II bounds.

Table 20.2: Black '76 NNEG and ERM Valuations Compared to Valuations Based on the Principle II Bounds

Volatility	***NNEG/NNEG LB***	***ERM/ERM UB***
0.00001%	100%	100%
1%	100.03%	99.98%
2%	100.22%	99.87%
3%	100.54%	99.68%
4%	100.99%	99.41%
5%	101.57%	99.06%
6%	102.29%	98.62%
7%	103.16%	98.10%
8%	104.18%	97.50%
9%	105.34%	96.79%
10%	106.34%	96.01%
11%	108.07%	95.15%
12%	109.63%	94.21%
13%	111.30%	93.21%
14%	113.08%	92.14%
15%	114.95%	91.01%

Notes: $NNEG$ is the value of the NNEG guarantee, $NNEG\ LB$ is the value of the NNEG Principle II lower bound, ERM is value of the Equity Release Mortgage and $ERM\ UB$ is the value of the ERM Principle II upper bound. Based on the baseline assumptions: male aged 70, LTV=40%, r=1.5%, l=5.25%, q=4.2% and σ as indicated. Exit probabilities are based on M5-CBD model projections using England & Wales male deaths rate data spanning years 1971:2017 and ages 55:89.

Each line shows the ratio of the Black '76 NNEG value to the NNEG lower bound, and the ratio of the Black '76 ERM value to the ERM upper bound, for a given volatility. The volatilities range from 1% to 15%.

We see that the ratios are remarkably close to 100%, especially for the lower volatility values. For example:

- for volatilities of up to 5%, the discrepancy between the Black '76 ERM valuation and the ERM upper bound < 1%;
- for volatilities of up to 11%, the discrepancy between the Black '76 ERM valuation and the ERM upper bound < 5%; and
- for volatilities of up to 15%, the discrepancy between the Black '76 ERM valuation and the ERM upper bound < 10%.[91]

So if you (a) don't have an option pricing model that you like or (b) don't like the dynamic hedging assumptions of Black '76, the solution is simple: don't worry about the dynamic hedging stuff and just use the bounds instead! Remember that *the bounds are Black 76 as well, but with volatility set to an arbitrarily low value*.

[91] One interpretation of these results is in terms of tolerance levels. If we apply a tolerance level, and that tolerance level is 10%, say, then we would be comfortable with the ERM bound valuations provided we felt that the 'true' volatility was no more than 15%, and so on.

You can then rest assured that by using the bound valuation instead of pure Black '76 the error in your valuations will be small if Black' 76 is anywhere close to being the 'correct' model.

The only circumstance in which there might be a larger error worth nothing is that in which the 'correct' model is one that gives much higher NNEG valuations than Black '76. So for those critics of Black '76 who feel that Black '76 NNEG valuations are too high, we bear news of great joy: throw the model in the bin and use the bounds instead, but let's hear no more complaints.

Even if dynamic hedging via liquid markets is impossible, and so it is impossible for Black '76 to give the exact cost of hedging, we can always apply an MC approach using the PRA bounds and these bounds do not depend on any impossible conditions.

Indeed, we could also use Black '76 too, but *not* on the grounds that we are making assumptions about liquid complete markets that we know are false but necessary for the validity of the model. Instead, we could use Black '76 *merely* on the grounds that Black '76 gives us results close to the bounds. Compare with the practice of surveying. It is impossible in practice, perhaps even in theory, to measure the exact length of the radius of a circle, or its circumference. Our measuring instrument will never be a constant length due to the effect of heat, and the exact length of the instrument at the sub-atomic level is impossible to quantify. Nor for the same reason is it possible to compare its length exactly with the object measured. Nonetheless we are comfortable applying the formula $c = 2\pi r$, even though the formula is true only in the 'continuous' world of pure geometry. We are comfortable because the mathematical ratio is a good approximation to the ratio determined by practical measurement, to whatever degree of precision we want.[92]

Pricing Beyond the Boundary

In the previous section, we showed how a simple boundary approach could be used as a basis for accounting valuations of the NNEG and ERM. This approach required no assumptions about the possibility of dynamic hedging at all.

But as we shall now argue, it is even possible to account for the costs of dynamic hedging using a purely notional approach to hedging. We can simply assume that the provider has create an internal hedge by creating two separate desks – an option desk and a trading desk – which trade *internally* with each other.

The option desk (like any option desk) would create two separate accounts, an option account and a hedging account. The option account would face the external market and would be valued using the standard Black 76 approach. The hedging account would be opened with the premium for the option. The option desk would calculate the Black '76 deltas of the option (i.e. the sensitivities of the option to changes in the underlying forward contract) at each decrement, and the delta-adjusted number of properties

[92] Indeed, the method of Archimedes was precisely to start with a discrete approach to determining π, by approximating a circle to a polygon then adding sides to the polygon, on the assumption that as the number of sides approached infinity, the perimeter of the polygon would approach that of a circle.

forecast to exit, at each decrement, then agree to deliver the indexed value of that number of properties to the trading desk, at each exit date. It would thus sell corresponding amounts of property forward (i.e. for delivery and settlement at the exit dates) to the trading desk. The prices agreed would be by agreement between the desks, perhaps based on nation-wide or regional price indices.

The figure below shows illustrative (spot) values of the hedge vs decrement year.

Figure 20.3: Spot Value of Delta Hedge Against Decrement

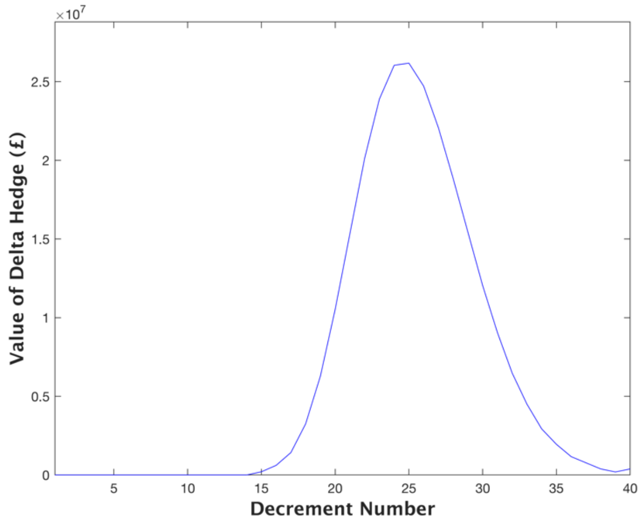

Notes: Our calculations based on Nationwide Index.

The figure plots the value of the delta hedge for each decrement number (or the time to maturity) of each put in the NNEG. For low decrement numbers the delta is zero but then rises according to a bell shape pattern to peak at decrement 25, and then falls thereafter towards zero.

For example, the option desk could agree to deliver the value of 22 average properties in the West Midlands region in year 15. The spot price could be determined by reference to the Nationwide index. Assume we are at the end of Q1 2019, when the average West Midlands property is worth £189,263, so 22 such properties would *now* be worth £4,163,782.[93] Assuming a net rental yield of 4% in that area, and that the firm's funding rate is close to risk free, say 1.5%. Then the agreed forward price F would be:

(20.4) $$F = £4{,}163{,}782 \times e^{((1.5\%-4\%)\times 15)} = £2{,}861{,}723$$

In 15 years' time the trading desk would pay the agreed price of £2,861,723, but at the same time would be paid the value of 22 West Midlands properties, according to the Index. Thus, supposing the Index had doubled in value over that period, the value paid

[93] In principle, properties could be broken down by type in a large number of different ways, e.g., by type of property (house or flat, price range, property size, garden size, if there is a garden, parking facilities, location, etc.). We gloss over those issues here, but one could fine tune the hedge to take account of them if one wished to.

would be £4,163,782 x 2 = £8,327,564, hence the total transfer would be £8,327,564 - £2,861,723 = £5,465,841.[94]

Assuming that the expected number of properties are in fact delivered by the borrowers' estates through the ERM agreement, and so long as the appreciation in the delivered stock matches the appreciation in the Index, the option desk will have broken even.

Now the option desk still bears the risk that the value of the properties delivered at exit will not match the value of the hedge position. Under the arrangement discussed, the desk promises to deliver the *indexed* value of 22 properties, not the realised or 'achieved' value. But there could be a number of reasons why the value promised could differ from the value realised. There is the risk that the predicted number of properties received on exit does not match the actual number – for example if borrowers lived longer than expected or if they exited at some earlier or later date. There is the risk, given that the provider can estimate only the number of properties that exit, but not the identity of the properties, that their average quality does not match the average quality of the index. Even so, much of the latter risk would be picked up and offset at other decrements. For example, if 22 extra properties exited in year 16, rather than in year 15, then the only loss would be in the cost of carry. Likewise if the quality of properties were below average in year 15, and if the overall book did reflect average quality, we would expect above average properties to exit in year 16, or at some point in the future. As we noted in chapter 10, there would be basis risk from the 'slippage' between the values of individual properties and the value of the index, but the dynamic hedging strategy would take care of (the bulk of) the period-to-period basis risk over the lifetime of the NNEG. In fact, that is its purpose.

The Appendix to this chapter provides a more detailed working of how such a hedge might work.

A possible objection to this two-desks argument is that such internal trading is artificial and does not represent reality.[95] Our response is that firms do sometimes use such arrangements *but they are best understood as accounting devices*. The aggregate p/l of the two desks will always add up to the external reported profits of the firm, because whatever profit is made internally by one desk will match exactly a loss made by the other. There doesn't even have to be a manned trading desk – the accounts could simply be made up on the basis of the arrangement above. The accounting arrangement is simply a form of p/l explanation, assigning to one book the profits or losses accruing from index and basis volatility, and assigning to the other book whatever accrues from purely proprietary trading positions.

The same reply could be made to the objection that no proprietary trading desk would be likely to accept a mandate from senior management to take on trades at cost from the option desk. Perhaps not, but given that the arrangement is merely an accounting exercise to explain the different sources of profit and loss, there wouldn't have to be any traders. Middle office staff could simply perform the work, for a basic salary. To repeat, the two desks arrangement is an accounting device, but from our perspective, the key

---

[94] The difference of £2,861,723 representing principal plus interest on the payment of £2,285,132 that the option desk would have received if the trading desk had paid them up front for the value of the deferment.

[95] We have encountered this objection in correspondence.

point is that it is a useful device to explain how a firm can implement a dynamic hedge against an option position.

Figure 20.4 plots the Black '76 vs. rehedge put values based on a historical simulation over Nationwide house price data for the period 1971:Q1 to 2018Q1. The time units here are quarters, so $k = 1$ here means rehedging every quarter.

Figure 20.4: Black '76 vs Rehedge Put Values: Historical Simulation for $k = 1$

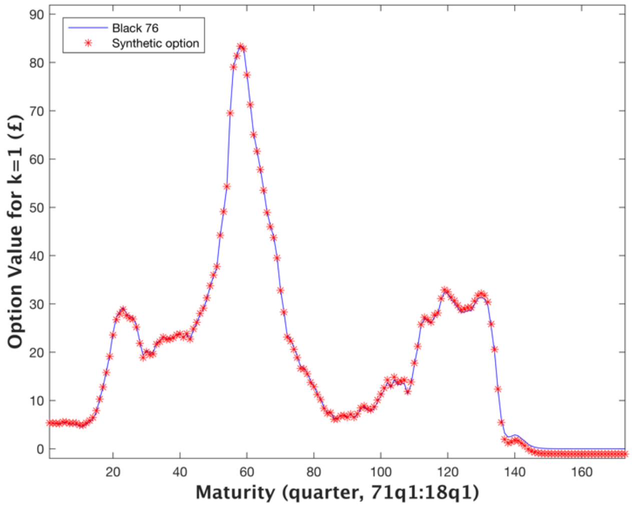

Notes: Illustrative put value simulation based on Nationwide UK property data for 1971Q1: 2018Q1.

We see that the synthetic put values closely track the Black '76 put values, implying that the rehedging approach gives much the same results as Black '76. Or, put the other way round, Black '76 would give much the same results as the rehedging approach.

Black '76 as a NNEG Lower Bound

To pull the discussion together, there are a number of reasons to think that Black '76 can under-estimate the NNEG put value. These include: (a) The presence of the Hurst effect for cases where $k > 1$. (b) The fact that any rehedging approach would use a discrete rather than continuous n (see, e.g., the positive mean hedge error in Figure 20.1(a). (c) The existence of transactions costs (see footnote 2). (d) The likelihood of basis risk, which might create the need for a reserve fund to absorb possible losses or, if we wish to book on the liability side of the balance sheet instead of the assets side, the basis risk might create the need for a capital requirement to cover the risk of loss. The price of the option is then the cost of the synthetic plus the cost of the reserve or capital required. The cost of the required reserve or required capital is a deep subject, however, and we do not examine it further in this report.

Conclusion

The MC approach can be implemented using (a) the PRA principles bounds, (b) Black ’76 or (c) some delta rehedging method. The principles bounds will produce valuations that are lower than those of Black ’76 and Black ’76 will produce valuations that are lower than those of the rehedging approach, but there are tradeoffs between ease of implementation and accuracy. The first is easiest to implement and makes the least calibration demands, but is the least accurate; the third is the most accurate but makes the most calibration demands and is the least easy to implement; and Black ’76 lies between the two. However, the differences between the three approaches are small and any of the three would be reasonable.

Again, we emphasise that the MC approach does not require that we use Black ’76 or depend on any assumptions that are impossible to hold.

Appendix to Chapter Twenty: A Worked Hedge Example

This Appendix goes through an illustrative exercise to show the practical steps one might go through to delta hedge an ERM exposure to a property price index.

To follow the calculations, we refer to the cell entries in the Excel spreadsheet "Chapter 20 Appendix ERM Example" under the tab "Hedge example". Note that the precise calibrations are unimportant. What matters is the hedge calculation process.

Assume the provider issues or acquires a book of ERMs on 8,100 properties scattered throughout the UK (see **B5**).[96] The average age is around 70 (see **F7**) and the average loan to value is around 40% (see **B11**). In practice such uniformity is unlikely, but the management of a more dispersed book is straightforward. With the average property valued at £213,000 (see **B6**) at the beginning of 2019,[97] there is an implied a total collateral value of around £1,723m (see **B8**), and (with average LTV of 40%) a total loan value of £689m (see **B7**).

We assume risk free rate of 1.5% (see **B9**), deferment rate 2% (see **B10**)**,** loan rate of 5% (see **B12**) and input volatility of 13% (see **B13**). The implied internal rate of return is 3.2% (see **B14**), comprising the risk-free rate, a NNEG spread of 1.8% (see **B15**) and an expected excess return of 1.7% (see **B16**).

In reality the provider will mark the ERMs to origination value and use the excess return as a spread over risk-free to discount the pension liabilities via a Matching Adjustment. For simplicity, we assume that the liabilities are valued correctly by discounting at risk free, but that the assets are booked at a day one gain.

These calibrations give a loan value, i.e., L, equal to £1,316m (see **E12**). L is the value of projecting the amount lent (weighted by longevity) at each decrement at the loan rate, then discounting at risk-free, on the assumption that there is no NNEG cost. The present value of the NNEG on these assumptions is £388m (see **E13)**, giving a net ERM value of £928m (see **E14**). Hence there is a day one profit of £928m - £689m = £239m (see **E15**).

The risk will be greatest around years 14 to 18 (see **M36:M40**). Below that year, the ERM portfolio will consist increasingly of non-defaultable loans. Above that year, there will be fewer borrowers. So consider decrement year 16.

The exit rate will be about 4.55% at year 16 (see **V38**), meaning that about 369 (~4.55% x 8,100, see **AG38**) of the 8,100 properties will be delivered. We don't know which of the 8,100 properties these will be, only that 369 will be drawn from the general population of properties. The delta of the ERM is 29.2% at that decrement (see **AH38**), which corresponds to 107 properties (see **AI38**). Assuming the average current value of properties of £213k (see **B6**), and a 20Y forward price of 90% of spot (see **E16**), the forward price of each property will be £192,453 (see **E17**). Assuming the forward price falls by £1, the value of the year 20 ERM falls by £107, i.e. 107 properties times £1 (see

---

[96] One of the Aviva ERF books (#2) started with exactly this number.

[97] Source, Nationwide HP index

E18). It follows that the forward price risk can be eliminated by selling 107 properties forward to the trading desk. The fall in value of the ERM will be approximately offset by the rise in value of the position held with the trading desk. Of course with the trading desk holding an equal and opposite position, i.e. being long the forward, the impact of the hedge on the firm's overall p/l will be zero, but as argued above, the internal trade device is simply an exercise in explaining p/l, to apportion the p/l in an appropriate way between those taking option risk, and those taking purely linear hedge positions.

One would apply similar analysis to the other decrements.

Though high simplified, this spreadsheet shows that the process of delta hedging an ERM exposure to a property price index is actually quite straightforward.

Chapter Twenty-One: Misconceptions About the Market Consistent Approach

The Market Consistent approach has come in for considerable criticism amongst the actuarial profession. Critics sometimes claim that you can't use Black '76 or Black-Scholes because these models are derived on the basis of a bunch of assumptions that are empirically invalid. If these assumptions don't hold, it is claimed, then we shouldn't use these models. Actuarial critics often write as if options practitioners were unaware of these issues, but in fact, the 'holes in Black-Scholes' are well-known to options practitioners who are expert in working their way round them ('traders' lore'). Fischer Black, the same Black as in Black '76, once wrote an article entitled "How to use the holes in Black-Scholes" (1988) in which he set out no less than ten of these unrealistic assumptions:

- The stock's volatility is known and doesn't change over the life of the option.
- The stock price changes smoothly and never jumps up or down.
- The short-term interest rate never changes.
- Anyone can borrow or lend as much as they want at a single rate.
- An investor who sells the stock or the option short will have the use of all the proceeds of the sale and receive any returns from investing these proceeds.
- There are no trading costs.
- An investor's trades do not affect the taxes paid.
- The stock pays no dividends.
- An investor can exercise the option only at expiration.
- There are no takeovers or other events that can end the option's life early.

He then showed how one might work around these limitations. "Since these assumptions are mostly false," he concluded, "we know the formula must be wrong. But we may not be able to find any other formula that gives better results in a wide range of circumstances" (1988, p. 67).[98]

Black concedes too much. He is right to acknowledge that these assumptions are mostly false, but it does not follow that the falsity of these assumptions makes the model invalid.

If proposition *A* implies proposition *B*, then establishing that *A* is false does not establish that *B* is false. The argument that *A* implies *B* so not-*A* implies not-*B* is the logical fallacy of denying the antecedent.

Let's say that we claim to have proved that the Moon is made of green cheese. You then disprove our proof. However, the Moon might *still* be made of green cheese even if our proof is invalid.

The point is that we need to be careful to distinguish sufficient and necessary conditions for Black-Scholes to hold. Yes, it is it true that in its usual classical/textbook derivations, Black-Scholes makes various assumptions, some of which are empirically false. However,

[98] See also T. Gordon "The Price of Actuarial Values," Staple Inn 1999.

most of these assumptions are sufficient rather than necessary for the model to be 'valid' in the sense we put forward in Chapter 20, namely that it provides a robust estimate of payoff minus cost of *discrete* hedging. So you can't just kick aside some of these assumptions and conclude that the Black-Scholes results must be wrong. For a start, there are multiple ways to obtain Black-Scholes – amongst others, there are the martingale approach, the binomial approach, Capital Asset Pricing Model (CAPM) and utility-based approaches[99] – and these will be based on different sets of sufficient conditions. To show that the model is 'wrong,' one would then have to establish *which particular* assumptions were necessary, and then demonstrate that one or more necessary conditions were not only wrong, but also introduced material errors into the resulting valuations. Any such an exercise is more involved than merely asserting that some particular assumption is empirically false.

Black-Scholes is Robust to Non-Randomness

Black-Scholes *works* when the realised profit or loss from *replicating* the option matches the assumptions used when *pricing* the option. As explained in the previous chapter, we start by computing the 'delta', or sensitivity of the option price to changes in the underlying price, then take a position equal and opposite to the delta, so that changes in the price of the option are offset by changes in the hedge position. Since the delta itself changes somewhat with the underlying price (an effect known as gamma), we might also occasionally re-hedge, but that is a separate issue.

It turns out that the accuracy of BS is surprisingly robust to a number of assumptions commonly made to obtain the BS formula.

One such assumption is GBM, so let's take a simulated Gaussian distribution, but sort it in order of magnitude, so that while it is still Gaussian, it is no longer random, so the 'Brownian' bit of 'Geometric Brownian Motion' no longer applies.

We then get the plots shown in Figure 21.1:

[99] For more on these, see, e.g., J. Andreasen, B. Jensen and R. Poulson (1998) "Eight Valuation Methods in Financial Mathematics: The Black-Scholes Formula as an Example," *Mathematical Scientist* 23: 18-40.

Figure 21.1: Synthetic vs Black-Scholes Put Option: Geometric Non-Brownian Motion

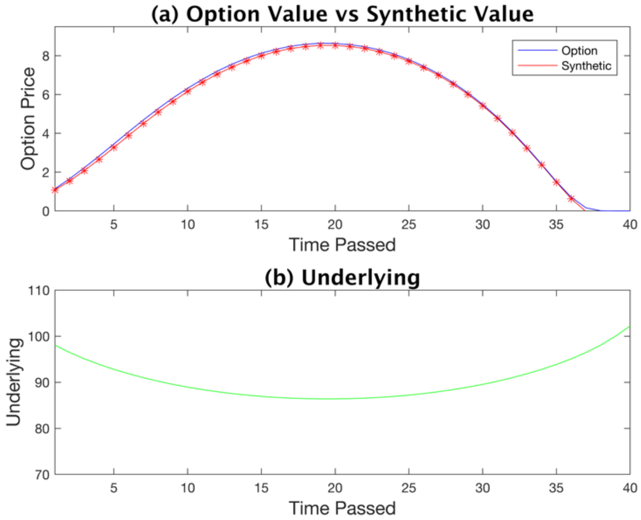

As we can see from the green line in the Figure, the price falls throughout the first half of the series and then rises.

The blue line shows the price of a put option price struck at 90, computed using the standard Black 76 option formula. As you expect, the price rises as the underlying falls below 95, then falls back as the underlying increases, ending at zero as the option expires out of the money.

The red line shows the value of the *synthetic* or *replicating* option constructed using the delta of the put. The (short) delta starts at 18%, since the underlying price begins at 95, rises to nearly 100% as the put is increasingly in the money, but then falls back to zero at the end.

The hedge is not perfect but it is close. The implication is that, while the assumption of randomness may be a sufficient condition for Black-Scholes to hold, it is by no means necessary.

Black-Scholes Robust to Mean Reversion

In its reply to CP 13/18, the IFoA states (p. 10) that

> Using the Black-Scholes formula in pricing NNEG will affect the cost of the guarantee, since allowance is not made for the features of mean reversion, momentum and jumps described above. Under geometric Brownian motion the volatility increases with the square root of time while for other models it does not; the value for long term derivatives such as NNEG *could materially differ from that assumed under the Black-Scholes model*.[100] (Our emphasis)

[100] "CP 13/18: Solvency II – Equity Release Mortgages: IFoA Response to Prudential Regulation Authority." 28 September 2018.

Let's have a look at this mean reversion claim. Figure 21.2 plots Black-Scholes against a put value under a mean-reverting process:

Figure 21.2: Synthetic vs Black-Scholes Put Option: GBM vs Mean-Reversion Process

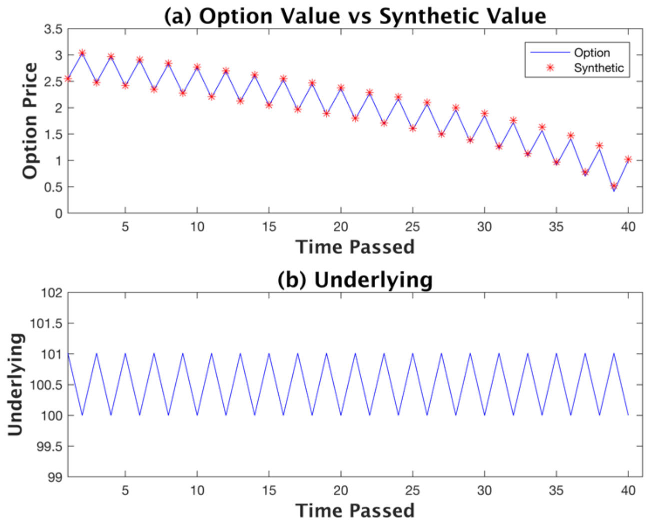

The chart above shows (green line) an underlying asset that follows the price series 95, 96, 95, 96 … i.e. the series has a mean of 95.5 to which it continually reverts. The blue line shows the price of a put option struck at 90, modelled by the standard BS option formula. The red line shows the value of a hedge, constructed from a derivative of the same formula ('delta') and they end up in approximately the same place.

We see that BS works well, even though the series is mean reverting and is not drawn from anything that even resembles a normal distribution.

So it is true that under a GBM process the volatility increases with the square root of time whilst for a mean-reverting process it does not, but this difference doesn't matter here. The longer the sampling period for the strange distribution above, the lower the sampled volatility: the prices series is going exactly nowhere. Yet the standard option pricing model still works well. Do a different simulation of the underlying path and you would get similar results.

Black ’76 is Undermined by Autocorrelation in House Prices

Another example is autocorrelation in house prices. On p. 14 of his report, Tunaru writes:

> GBM as a data generating process for house prices is totally inappropriate because it ignores serial correlation and stickiness of prices, as well as clustered volatility and downward jumps.

We agree with him that GBM does not provide an empirically accurate description of the empirical house price process, but we would still argue that Black ’76 *which assumes GBM* still provides a good framework to value NNEGs and ERMs.

Let’s leave aside issues of price stickiness, clustered vol or downward jumps, which really point to the calibration of the volatility parameter fed into Black ’76, and focus on the serial correlation or autocorrelation of house prices.

The issue is whether autocorrelation undermines the validity of Black ’76.

It is true that many random variables are autocorrelated and that that the usual derivations of Black ‘76 assume that the underlying random variable is not autocorrelated. However, recall that the question of how autocorrelation might affect option pricing was addressed in the Cornalba- Bouchaud-Potters article that we discussed in Chapter 10 and their conclusions were clear. “In the Gaussian case [the one considered in Black-Scholes], we find that the effect of [auto-] correlations can be compensated by a change in the hedging strategy and therefore *options should be priced using the standard uncorrelated Black-Scholes model*” (our italics). The key is to ensure that the volatility is measured on the same time scale as the rehedging but this qualification merely amounts to an adjustment, if any, to the volatility calibration, and BS still holds. The same argument can then made for Black ’76.

We discussed the connection between rehedging and volatility in Chapters 10 and 21. The essential point is that the volatility should correspond to that of the rehedging period. Suppose then that an option is being rehedged every 5 years. If the monthly volatility is 1%, then under GBM we would input

$$\text{adjusted } \sigma = 1\% \times 60^{0.5} = 7.75\% \tag{21.2}$$

into our option pricing equation, but if the underlying is autocorrelated with $H = 0.9$, say, we would input the following H-adjusted volatility

$$H\text{-adjusted } \sigma = 1\% \times 60^{0.9} = 39.8\%. \tag{21.3}$$

Black ’76 is still valid in the presence of autocorrelation provided we use the appropriate volatility, i.e., (21.3) instead of (21.2).

‘No Short-Selling’ Argument

The ‘no shorting’ argument is often wheeled out to claim that the usual arbitrage-free arguments underpinning Black ’76 don’t apply. Critics claim that the validity of Black ’76 depends on a no-arbitrage argument that itself depends on being able to sell the forward contract, but then claim that no-arbitrage does not apply because, e.g., one cannot sell the forward. An example is given by Tunaru:

> In our opinion, the formula (10) [i.e. Black ‘76] simply does not apply for house prices and this is unrelated to the Gaussian distribution assumption behind the GBM model. The forward contract on a house price cannot be calculated as in (11) [i.e. according to our equation (3.10): forward price = spot price $\times e^{(r-q)t}$], simply imitating the no-arbitrage formula for a stock paying dividend, where the dividend yield is replaced by the net rental rate. That formula cannot work *because currently we cannot shortsell the value of a house*. Hence, the no-arbitrage principle does not apply here to lock in the forward price as in the case of corporate stock. (Tunaru, 2019, p. 14, our emphasis)

But *of course* you can sell a house! ‘Can’ just means ‘it is possible’ and there are no legal restrictions against doing so. As we explained in the previous chapter, anyone can enter into an agreement whereby one party agrees to deliver an individual property of specified size, location and standard to the other party, at an agreed date and for an agreed price. The Over-the-Counter (OTC) market makes these trades all the time. If no individual property is available, the parties could contract to buy or sell the value of, say, the Halifax index, times some currency multiplier, at some agreed date in the future. Parties who wish to make such trades can usually find investment bankers willing to be their counterparties.

Another version of the ‘no shorting’ argument is that the usual arbitrage-free assumptions don’t apply, because you can’t (delta) hedge your property exposure. However this version is also false, because it is actually quite straightforward to delta-hedge a property exposure and explained in the previous chapter.

To recap. You have a collection of properties which will be coming onto your books at times and amounts corresponding to your longevity modelling. You work out your deltas to the market across the decrements and sell the rights to future possession to some interested counterparty. So you agree to deliver a fixed amount of properties at each decrement year, according to a pre-agreed standard (n bedrooms, y location etc), with contractually agreed penalties for late delivery or non-delivery. The only difference between this arrangement and a standard reversionary contract is that the identity of the properties is not known in advance. For example, suppose you expect to receive 50 four bedroom properties in the South East in year 24, and your delta is 70%. Then you commit to deliver 50 x 70% = 35 properties in that year, in return for cash. If it turns out only, say, 30 properties exit, then you defer the delivery of the remaining 5 until the following year, with an agreed penalty, perhaps one year’s rental equivalent, or settle the whole amount for cash. The key is working out the cash amount the hedger would be willing to pay in return for receiving the 35 properties.

Indeed, sometimes you don't even need hedging to carry out valid Black '76 valuation. Consider a deep-out-of-the-money option, illustrated by the left-hand-side of Figure 21.3, where ERM_t is equal to L_t, the PV of the loan value. Or consider a deep-in-the-money option, illustrated by the right-hand-side of Figure 21.3 where ERM_t approaches or is equal to the PV of the forward at t.

Figure 21.3: ERM Valuation and Forward House Price Volatility

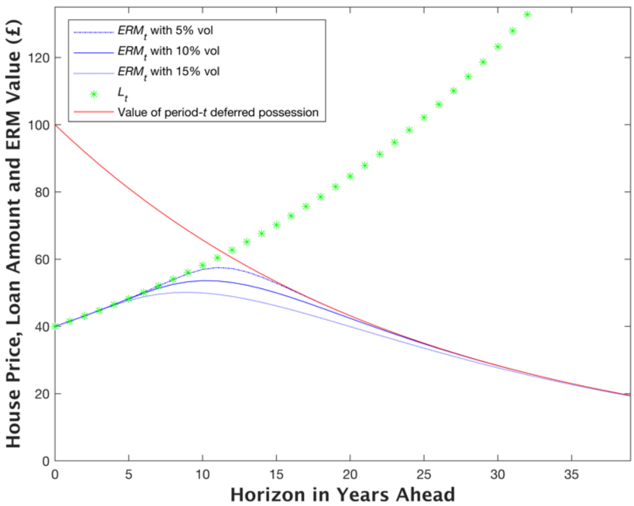

Notes: ERM_t is the value of the Equity Release Mortgage for decrement t with the relevant volatility and L_t is the value of the loan for decrement t. Based on the baseline assumptions: male aged 70, LTV=40%, r=1.5%, l=5.25% and q=4.2%. Exit probabilities are based on M5-CBD model projections using England & Wales male deaths rate data spanning years 1971:2017 and ages 55:89.

In both cases, the value of $NNEG_t$ does not depend on whether the option can be hedged.

Yes, you might say, but you have not shown that this argument also applies to the central region of the Figure where the value of ERM_t is below the minimum of the loan and deferred house price curves. This objection is true, but now allow the volatility to go towards zero and the central region disappears. Black '76 is then valid across the entire horizon spectrum. It follows that, at most, any concerns about the possible impact of imperfect hedging opportunities on the option value must translate into the calibration of the volatility parameter, so Black '76 *is still valid* or at least approximately so provided one handles the volatility calibration with appropriate care.

Ah, you might respond, the red upper bound (or deferment house price) curve in the above figure still depends on the ability to short-sell the right to deferred possession. However, this response is not correct. Recalling

(3.10) $$R_t = current\ house\ price\ \times e^{-qt}$$

we can always obtain a deferment price if we have the spot price and a reasonable estimate of the deferment rate q. The former is easy to ascertain and we discussed the calibration of the deferment rate in Chapters 8. Even though q is unobservable, the accounting fiduciary principle comes into play (see Chapter 26 below), that not only allows but mandates that even when a market price or other relevant entity is unobservable, an instrument should be priced *in a way that reflects the assumptions that market participants would use when pricing the instrument*. Market participants are by definition independent of one another, knowledgeable or have a reasonable understanding of the instrument, and willing and able to enter into a transaction.[101] In pricing a deferment, a market participant would therefore consider the cost of losing the income for the deferment period, and would adjust the price of immediate possession accordingly (see also SS 3/17, p.12). That this exercise involves some judgment goes without saying, but that judgement has to be reasonable and there are bounds on what a reasonable judgement can be.

'Rental Rate is Irrelevant Because a House is a Consumption Asset'

A final argument is that the standard forward valuation (e.g., equation (3.9) which states that forward price = current house price $\times\ e^{(r-q)t}$) is inappropriate because the rental rate is irrelevant to those who purchase houses to live in them. To quote Tunaru again (p. 30):

> The concept of rental yield has been introduced into real estate valuation by analogy with the link between dividends and share prices. However, it can be argued that the buyer of a house is not the equivalent to an investor buying a house as an investment asset. For the majority of buyers, houses play the role of a consumption asset and not that of an investment asset.[102] There is no evidence that rental yields are driving future house prices so the expected house prices at various future long horizons cannot be determined with growth models in the same way expected share prices may be determined with growth models linked to dividends.

One response is that the ERM valuation question applies to *institutional investors*, namely ERM lenders, who want to acquire residential property exposure, and who are not using the property as a consumption asset, but rather as an investment asset. Our point is that it is the institutional investors who matter here, because they are the ones who issue NNEGs, not the retail investors.

In any case, why would the *ordinary* buyer approach the valuation of a deferred possession any differently from an institutional buyer? Suppose we are currently locked into a leasehold but are looking to move into a property in 2 years time after our tenancy agreement runs out. You are looking to move out of your property at the same date, but need money now. So we agree to pay you money now in order to possess the property at

[101] See IFRS 2013, para 87 on unobservable inputs, and Appendix A on definitions.

[102] A related argument we have heard is that while the utility of not having to pay rental may be relevant to 1st time buyers trying to escape rental into owner-occupier, once people have a home of their own, they don't think about rental yield when looking at moving house "up the ladder." This type of argument focusses on how people think of their own home, i.e. as somewhere to live, and not as an investment. However, we would argue that the psychology of individual home owners is irrelevant, if only because the relevant owner is the ERM lender, who has a pile of property with rights to deferred possession on their books.

that future date, but the price we pay now would not be the price of immediate possession. Why should we pay the full price of the property now if we cannot move in or rent it out for a full two years? And how we would value the deferred possession? Answer: at some discount to the value of current possession, where the discount would be driven by the rental rate.[103]

Tunaru writes that "there is no evidence that rental yields are driving future house prices" but that claim is irrelevant even if it is true. It is not future house prices we are concerned with, but rather the price *now* of a contract for possession in the future. We are comparing the *current* prices of two different contracts, one for immediate, the other for future possession, and we are not speculating about future house prices, which are irrelevant here.

Of course, anticipated future prices may be relevant if it is the *decision* that is deferred. Do we buy now or do we wait for two years? If we expect prices to rise, we will pay for immediate or deferred possession. If we expect them to fall, we will defer our *decision* to buy, whether that be immediate possession or deferred possession. Decision is different from possession, however. If we pay now for deferred possession, we are already exposed to future house prices. If we expect prices to fall, we will pay neither for immediate nor deferred possession. Rather, we will wait, i.e. we will defer our decision.

All these points follow from elementary pricing economics.

[103] Or still another version of the 'householder is different' argument. Andrew Rendell at the ARC ERM Launch Event on 28th February 2019 argued that we must consider the utility to the occupier of living in the house, without saying how we might measure that utility. But what if we ask you to vacate your house for 2 years, so you no longer have that utility? What is the cost of it? Surely the cost of renting for 2 years.

Chapter Twenty-Two: The Discounted Projection Approach

The standard approach used by ERM actuaries in the UK is the Discounted Projection (DP) approach, sometimes also called the 'real world' approach. This approach is based on the use of a projection of future house price growth to value the NNEG. In particular, it replaces the forward house price as the underlying in the MC approach with some 'expected' future house price. Equivalently, it replaces the forward rate f in the MC approach with some assumed rate of future house price growth hpi.

Since $f = r - q$, replacing f by hpi gives

(22.1) $$hpi = r - q.$$

Given also that hpi has been specified and r can be easily calibrated from the spot rate curve, (23.1) implies that we can back out the following implied q:

(22.2) $$q = r - hpi.$$

To give an example, if we set $r = 1.5\%$ and use the 4.25% hpi assumed recently used by Just Group[104] then we can back out q as

(22.3) $$q = 1.5\% - 4.25\% = -2.75\%.$$

The negative sign in front of the '2.75%' on the right-hand side of (23.3) is not a typo. For this calibration, and indeed for any calibration in which hpi exceeds the risk-free rate, the DP approach produces a *negative* q net rental rate.

One can implement this DP approach by taking an otherwise sound MC (e.g., Black' 76) calibration and replacing the forward rate with an assumed future house price growth rate hpi. So one replaces the q rate one would otherwise use (e.g., a sensible q rate of around 3% or more) with an implied q rate equal to $r - hpi$ (e.g., a q rate of about -2.75% as in the Just example).

To illustrate the impact of this approach, Table 22.1 shows the NNEG and related valuations for our baseline calibration, obtained using each approach:

[104] The firm reported using this number in both its 2016 and 2017 *Annual Reports* (see pp. 163 and 110 respectively). The same number also appears in its 2018H1 results (p. 18).

Table 22.1: Baseline ERM/NNEG Valuations: Market Consistent vs. Discounted Projection Approaches

Approach	L	$NNEG$	ERM
Market consistent	£74.84	£32.19	£42.66
Discounted projection	£74.84	£4.37	£70.47

Notes: L is the present value of the loan component of the Equity Release Mortgage, $NNEG$ is the present value of the NNEG guarantee, and ERM is the present value of the Equity Release Mortgage. Based on the baseline assumptions: male aged 70, LTV=40%, r=1.5%, l=5.25%, q=4.2% for the MC approach and q=-2.75% for the DP approach, and σ=14.8%. Exit probabilities are based on M5-CBD model projections using England & Wales male deaths rate data spanning years 1971:2017 and ages 55:89.

In this case, which is not untypical of others we have looked at, the DP approach gives NNEG valuations that are close to an order of magnitude lower than those produced by the MC approach. The result is a considerable overvaluation of the ERM, in this case by 57.3%.

But ask yourself: do the DP valuations even *look* right? If you believe them, then you have to believe that the 'true' $NNEG$ is only 4.37/74.84 = 5.8% of L. This $NNEG/L$ ratio looks awfully low when you consider the spread between the loan rate and the risk-free rate, which is 5.25% - 1.5% = 3.75%. If the loan has so little risk, then why is the spread so high?

Here we see in a nutshell the valuation problems entailed by the use of the DP approach.

Origins and Overview of the Discounted Projection Approach

So where does this approach come from? In 2005, the IFoA published a report on NNEG valuation (IFoA, 2005). This report confirmed that it was reasonable to use Black-Scholes methodology when seeking to obtain a 'market consistent' NNEG valuation whilst noting that it "is not without its difficulties and shortcomings." We have no argument with that assessment. However, it also noted that, "Others may however, prefer to approach the assessment of the NNEG using more of a "real world" stochastic modelling approach," whatever that might be, and they did not explain. And so we have the juxtaposition of BS as a reasonable approach to 'market consistent' NNEG valuation, versus an alternative unspecified 'real world' approach that is pulled out of thin air and gives a different valuation.

Two years later, the IFoA issued another report on NNEG valuation, Hosty *et alia* (2007). This report started with some concerns about the decline of profitability and its impact on the development of the ERM market:

> the competitive environment that has driven product innovation has … resulted in lower product margins. This is all good for the consumer, but it is increasingly difficult for providers to reach target returns on capital, and this is deterring some prospective new entrants. One of the purposes of this paper is to investigate the profitability of typical schemes in the market at present, and so to address the question of whether competition has forced the market

> to function at non-profitable levels. … We will aim to provide a rational pricing methodology which can be adopted by any organisation active in the market, and we hope that this can support the market as it expands over the coming years.
>
> There is now concern that providers *may not be able to offer a product profitably at current margins*. Some competitive pressure is clearly a good thing, as it will force providers to find more efficient ways of providing their product to consumers. In the equity release market, too much competitive pressure may be a bad thing. (pp. 1-2, our emphasis)

To cut to the chase: their main concern is that overly high NNEG valuations might undermine the ability of firms to meet their profit targets. We are sure they are right, but the question is, if their concerns are correct, how to reconcile their commercial concerns with their obligations (under actuarial and accounting standards, see Chapters 25 and 26 below) to provide unbiased and fair value valuations.

If there wasn't a conflict between these two objectives (i.e., profits and unbiased/fair value), then we wouldn't even be having this conversation. The fact that ERM industry leaders emphasise the conflict between the two objectives and their preference for commercial considerations over unbiased/fair valuation indicates that these objectives are in conflict. Otherwise they would argue for unbiased/fair valuations and wouldn't need to argue for the primacy of commercial considerations instead.

They then examine what they understand the "market consistent" approach to be. They do not define the term "market consistent" however and the nearest we get to an explanation is that this approach is based on an

> approximate market consistent basis similar to the pricing of options on stocks. … The main challenge with a market consistent basis is the fact that there is no underlying market to speak of. Accordingly we have tried to create a proxy market consistent basis using techniques that are standard in similar markets, specifically Black Scholes style modelling. (p. 26)

The counter-argument is that there is always an underlying market! Almost all property transactions are forwards, admittedly, short-maturity forwards but there are no legal barriers to longer maturity forwards and ERM firms could always approach investment banks for quotes. Whether ERM firms wish to trade at those rates is another matter. We take their point about 'proxy' valuations, but as we discussed in Chapter 9, we can obtain empirically grounded proxy valuations from valuations in the leasehold market.

Hosty *et alia* then explain what they mean by proxy valuation:

> Using a risk neutral basis, *house price inflation* should be linked to the return on long term risk free instruments (i.e. government stocks) less an assumption for rental income (net of expenses). (p. 26, our emphasis)

But this way of proxying valuations is based on a howler of an error. The error is that it confuses the future price of spot possession with the current price of deferred possession.

This error is hugely material, since the one variable (the future price of spot possession) usually goes up over time, whilst the other (the current price of deferred possession) falls with maturity.

To illustrate the magnitude of this difference, consider the plots in Figure 22.1:

Figure 22.1: Future vs Deferment House Prices

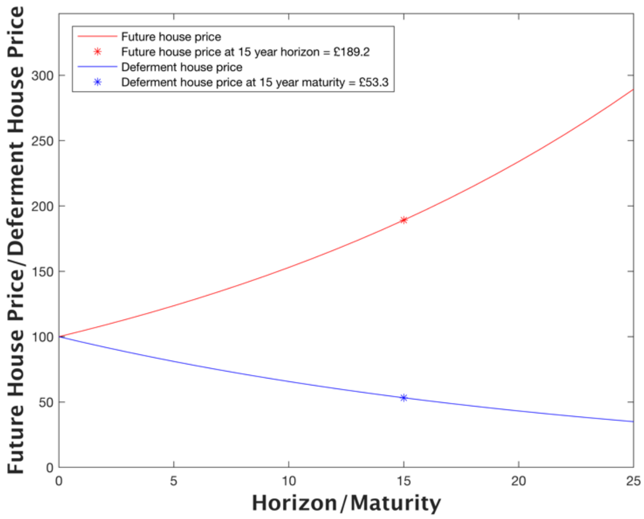

Notes: Based on current house price = £100, *hpi*=4.25% and *q*=4.25%.

The blue plot gives the correct price to use in the option pricing formula, i.e., the spot price of deferred possession. For a maturity of, say, 15 years, and our recommended deferment rate of 4.2%, the deferment house price (the price of a spot contract for possession in 15 years) is £53.3. The red plot gives the incorrect price that the Hosty *et alia* argument implied should be used, i.e., the future house price. This particular plot is based on an assumed 4.25% hpi rate. The expected future price in 15 years is £189.3.

So based on these calibrations, the Hosty *et alia* argument implies that we use an underlying value of £189.3 in the option pricing equation when the correct value is £53.3. If you believe that the Hosty approach is right, then you are believing that an asset, a forward worth £53.3, is actually worth £189.3, in which case let's do a trade. Thus, the root issue with the DP approach is a misinterpretation of the forward contract that leads to a large over estimation of the contracts' value.

The Hosty approach then produces a NNEG of £3.00 if we make our other baseline assumptions when Black '76 correctly applied would give us a NNEG of £31.42. The Hosty approach thus leads to a NNEG valuation that is 9.5 % of the Black '76 valuation.

The Hosty *et alia* use of the incorrect term 'house price inflation' instead the correct term 'forward house price' suggests that they consider that the future house price (or hpi rate, depending on the formulation one wishes to use) should go into the BS or Black '76 model, but neither the future house price the expected future house price nor the expected hpi rate belong in those models. You can input them if you insist, but you shouldn't, because the model gives you have no leave to. These variables are irrelevant in any BS-family option price model. Instead, BS or, more precisely, Black '76 tells us that the underlying variable that *should* go into the option pricing equation is the forward price, in this case, the forward house price. The use of term 'house price inflation' in this context suggests a serious misunderstanding of how BS option pricing works and leads to a major under-estimation of the NNEG value.

But Hosty *et alia* go on to make plain that they do not like the MC approach:

> In reality the absence of an underlying market means that this proxy market consistent approach is only of limited academic value … (p. 27, our italics)

By "absence of an underlying market" they mean the absence of a liquid market in which the option can be hedged using e.g. a zero-arbitrage trading strategy. But as we have explained in Chapter 20, it is perfectly feasible to apply the MC approach in the UK property market context, and whatever the difficulties of doing so, those difficulties in no way give Hosty *et alia* licence to replace the forward house price in the option pricing equation with some guess-estimate of the future house price. The "only of limited academic value" jibe is presumably meant to suggest that the MC approach – or "proxy market consistent approach" as they put it – is of no practical 'real world' use and perhaps to hint that practitioners should be looking for a more 'real world'-friendly alternative? Again, we disagree. The MC approach is not only feasible, but has no feasible alternative.

Then they make a further criticism of the MC approach:

> For providers attempting to price the NNEG on a market consistent basis there is insufficient product margin in order to provide a competitive product unless they have strong competitive advantages in one or more of the other cost areas. (p. 30)

Whether or not this claim is true, this statement begs the central issue, i.e., whether the MC-based valuations are reliable, and Hosty *et al.* provide no convincing scientific grounds to question them.

So Hosty *et ali*a's main objection to MC valuation boils down to it giving valuations that they don't like. But remember the problems that Equitable Life got into 20 years ago when it was discovered to have been undervaluing its long-term guarantees!

Section 7.3.2 examines their preferred alternative, an "insurance pricing basis using "real world'' assumptions." What these assumptions might be they do not explain; nor, do they explain how this "real world" approach might be consistent with a very un-real world negative net rental rate. In fact, they don't explain what their "real world" approach even is.

Section 7.3.2 consists of only 143 words and is here reproduced almost in full:

> 7.3.2 *"Real world: assumptions"*
> The alternative method we have used is to calculate the option cost using "real world" basis. The methodology we have used is as follows:
> – Use the log normal model as before (with same volatility).

> – A best estimate of 4.5% p.a. for HPI in the future (see Section 4.4). This is then the mean return under the model.
> – We have assumed that a real world discount rate of 4.75% per annum.
> – We have not assumed a "mean reversion'' so that the random walk in each future period is applied independently of the position is [sic] preceding periods. The authors acknowledge that use of a "mean reversion" approach is equally valid. …
>
> As can be seen [from Table shown], the resulting costs are significantly below those assessed using our proxy market consistent basis.

So not a word of explanation as to why we should regard this 'real world' approach as reliable, but the phrase that jumps out is "A best estimate of 4.5% for HPI in the future," i.e., the RW approach is based on a guess about future HPI!

Based on the limited information provided, their 'real world' approach would appear to be similar to the MC approach, but with the forward house price replaced by some assumed expected future HPI.

We now see the seed germinate. The 2005 IFoA report introduced the Trojan Horse of house price inflation, but at least did the calculations correctly. This error could be forgiven as an innocuous terminological one, except that the passage quoted opens the door to full-scale misuse and seems to confirm that the Hosty *et alia* 2007 'real world' valuation approach is based on exactly that error. The inclusion of HPI is no longer a mere mislabeling, but a bedrock principle of the RW approach.

To spell it out, HPI is now a key input *in its own right*.

Which points confirm that this approach is inconsistent with option pricing theory and therefore wrong.

Section 7.3.3 clarifies the authors' views on which approach is to be preferred. We reproduce part of it here:

> *7.3.3 Market consistent or real world?*
> On our proxy market consistent approach we have derived a cost for the NNEG which would render the product non-profitable, whilst real world modelling has produced a significantly lower cost.

This statement has major repercussions. If NNEG valuations on a MC basis would make ERMs unprofitable and if there is no justifiable alternative to NNEG valuations on an MC basis, then doesn't that make the ERM sector unprofitable? And if the sector only appears to be profitable because the 'real world' NNEG valuations make it appear so, then doesn't that mean that the profits that the firms have been making may have been more apparent than real?

The importance of commercial considerations as a reason for preferring this approach was confirmed by Tom Kenny at the 28 February 2019 Staple Inn launch event for the Tunaru report. Mr Kenny was the chair of the event, and is Director of Actuarial &

Underwriting, Retirement Lending at Just Group plc in his day job: "clearly if we move down a purely market consistent route … it's going to be extremely expensive," he said.

However, the issue is not whether the firms' NNEG valuations would go up if they used another approach. The issue is whether firms are using the right approach in the first place. If firms are using a valuation approach that greatly undervalues their NNEGs, then they have greatly under-estimated their costs and those costs are already being borne by firms and their investors, regardless of whether firms acknowledge that fact. Firms should be facing up to this problem instead of denying it. For their part, analysts should be wondering how big this problem might be and asking themselves about the potential impact on firms' financial conditions. Under-valued costs mean hidden losses and over-estimated capital, potentially on a large scale.

Imagine if Bosch are under-valuing the guarantees they issue with their washing machines. Their management then discover that the costs of replacing or repairing their washing machines are going to be higher than they had expected, but they don't yet know how much higher. The problem might only be a small problem but then again it might not. So what is the most appropriate response from the management when they are informed of it? Should they deny it on the grounds that they wouldn't like it if their guarantees turned out to be more costly than they had thought or should they look into the issue with a view to fixing the problem before it gets any worse? We would have thought that the answer to that question was obvious, but then why would the answer be any different if it was ERMs rather than washing machines whose guarantees were being under-valued? And if the losses involved might potentially be on a large scale, then doesn't that reinforce the need to address the problem as a matter of some urgency, lest a potentially large problem grow into a larger problem down the road if nothing is done about it?

Problems with the DP Approach

Consider that the Discounted Projection approach:

- has never been convincingly justified by those who advocate it;
- is being promoted by practitioners with a vested commercial interest who are promoting it for openly commercial reasons and are dismissive of the only approach that is scientifically respectable because they do not like the valuations it produces;
- has not been endorsed by a recognised independent expert;
- does not appear in the corpus of recognised scientific research journals that are subject to rigorous peer-review;[105] and

[105] Admittedly, Hosty *et alia* (2007) was later published in the *British Actuarial Journal*, but it is not clear whether *BAJ* articles (or for that matter, any articles and reports published by the IFoA, e.g., such as the Tunaru report) are subject to "rigorous peer review" and the only thing that is clear about the review process, whatever that might be, is that it is unclear. To quote from the BAJ's website, "*British Actuarial Journal* contains the sessional research programme of the Institute and Faculty of Actuaries along with transcripts of the discussions and debates. It also contains Presidential addresses; memoirs and papers of interest to practitioners." Nothing here about rigorous scientific peer review.

- is contradicted by alternative approaches such as Black '76 that are used and taught all over the world and have been published in top tier academic journals, albeit that their applications are sometimes still controversial.

There are two root problems with the DP approach. The first is that it is based on a forecast, in this case a forecast of future house price growth. However, it is a serious mistake to base an option valuation on a forecast, because basic option pricing theory (e.g., Black-Scholes or Black' 76) tells us that the value of an option is not dependent on any forecast, let alone some guess value pulled out of thin air. Instead, the option should be priced using current variables only – admittedly subject to some judgements about calibration but experts know how to handle these calibration issues. It follows that any approach that does depend on a forecast must be invalid and when a firm says that it bases its NNEG valuations on a forecast – any forecast – then we know that the firm must be getting it wrong.

The second root problem is that it is wrong on principle to replace the forward house price in the put pricing equation with the expected future house price (or equivalently, to replace the forward rate with the expected future house price inflation rate). This error can produce results that are known to be impossible. As Table 22.2 shows, the DP approach can produce results that violate two different sets of impossibility bounds: the Principle II bounds and Principle III bounds examined in Chapter 19 above):

Table 22.2: Baseline ERM and NNEG Valuations: Discounted Projection Valuations vs Bounds

Approach	$\boldsymbol{NNEG}$	$\boldsymbol{ERM}$
Discounted projection	£4.37	£70.47
	$\boldsymbol{NNEG}$ **lower bound**	$\boldsymbol{ERM}$ **upper bound**
PRA Principle II bounds	£28.09	£46.75
PRA Principle III bounds	£13.15	£61.69

Notes: $NNEG$ is the present value of the NNEG guarantee, and ERM is the present value of the Equity Release Mortgage. Based on the baseline assumptions: male aged 70, LTV=40%, r=1.5%, l=5.25%, q=-2.75% for the DP approach and q=4.2% for the bounds, and σ=14.8%. Exit probabilities are based on M5-CBD model projections using England & Wales male deaths rate data spanning years 1971:2017 and ages 55:89.

The DP valuations violate all these bounds. The DP NNEG valuations fall below the lower bounds, and the DP ERM valuations exceed the upper bounds.

So what do these violations actually imply? Well, the bounds are specified in terms of the forward and loan value decrements F_t and L_t. Assuming these to be correct (and why shouldn't we?) then if ERM decrements ERM_t violate their bounds, then the ERM valuation will be impossibly high and the NNEG valuation impossibly low. Since these DP valuations are known to be impossible, then no auditor can sign off on them because fair value principles do not allow impossible values.

In short, if we wanted a one sentence assessment of the validity of the DP approach, all we need to know is that it produces valuations that violate bounds that cannot be violated. It is hard to see how any approach can get much more wrong than that.

Thus, the correct approach is to start with the forward price = $Se^{(r-q)t}$ (see equation (3.9), which then gives us the discounted forward price or deferment price = Se^{-qt} (equation (3.10)).

The DP approach incorrectly treats the 'forward price' as the projection price = $Se^{hpi\times t}$, which then gives the 'discounted forward price' or 'discounted projection' price = $Se^{(hpi-r)t}$.

The DP is approach is wrong on principle and (as noted in Chapter 3) confuses the future price of spot possession with the current price of future (=deferred) possession.

It will give the wrong answers in general except in the special case where $hpi = r - q$.

Chapter Twenty-Three: The Tunaru Report

On 19 February 2019, the Actuarial Research Centre[106] published Professor Radu Tunaru's much awaited report on NNEG valuation.[107] His report had been commissioned by the Association of British Insurers and the IFoA in the hope that it would resolve some of the NNEG controversy.

The IFoA press release accompanying the Tunaru report reassured the public that the "valuations arising from insurers' current models and bases are sufficient," as if to suggest that the Tunaru report puts any concerns about these models to rest. Tunaru does no such thing.

Instead, Tunaru offers an approach that is free of the DP approach's most glaring weaknesses – the reliance for option pricing on forecasts and on absurdly low and typically negative net rental rates – but generates NNEG valuations that are of the same order of magnitude as that approach. Therefore it too fails the sniff test, in that it produces $NNEG/L$ ratios that are too low to justify the spread between lending and risk-free rates. Nonetheless, we can see why the industry would welcome it. From their point of view, Tunaru offers another way to skin the cat without getting the flak about negative deferment rates.

To quote from the transcript of the Staple Inn launch event on 28 February 2019:

> Gareth Mee (EY): Excellent new research published today on Equity Release Mortgages. Really moves the industry forwards in terms of understanding and risk management.

It seems to us however that the celebration might be premature and we are not the only ones to have our doubts. Similar concerns have been expressed by Tony Jeffery and Andrew Smith in their equity release report to the Society of Actuaries in Ireland on 28 March 2019, and by a number of actuaries who have corresponded with us privately.

Excessive Parameterisation: the ARMA-EGARCH Model

Professor Tunaru opens his report with an oft-cited quote from George Box: "All models are wrong but some are useful."[108] He uses this quote as a way to introduce his preferred model, the ARMA-EGARCH model – or to give it its full name, the Autoregressive Moving Average Exponential Generalised Autoregressive Conditional Heteroskedastic model.[109]

---

[106] Disclosure: one of us (Dowd) is currently working with the ARC on its project on the modelling, measurement and management of longevity and morbidity risk.

[107] R. S. Tunaru and E. Quaye "UK Equity Release Mortgages: a review of the No Negative Equity Guarantee." Actuarial Research Council and Institute and Faculty of Actuaries.

[108] G. E. P. Box (1976) "Science and Statistics," *Journal of the American Statistical Association* 71 (356): 791-799.

[109] See, e.g., R. F. *Engle (1982)* "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation," *Econometrica* 50 (4): 987–1007; T. Bollerslev (1986) "Generalized

This model has been around for a long time however and is well known to time-series econometricians and applied economists, who use it principally to model *short-term* volatility dynamics in asset markets. It is not obvious however why one would *want* to apply a model of short-term volatility dynamics to a long-term problem such as NNEG valuation. Just as 'is' does not imply 'ought', 'can' does not imply 'should'.[110] The ARMA-GARCH model adds needless complexity to a set of problems that that are already much misunderstood.

The ARMA-EGARCH has a lot more parameters than Black '76, and on the subject of choosing between different models Professor Box went on to state

> Since all models are wrong the scientist must be alert to what is importantly wrong. It is inappropriate to be concerned about mice when there are tigers abroad. (Box, 1976, p. 792)

The key phrase is "importantly wrong", as opposed, e.g., to trivially wrong. His point was then when choosing between models, it is advisable to go with the more parsimonious one unless there is good reason to the contrary. To quote Box again (1976, p. 792):

> Since all models are wrong the scientist cannot obtain a "correct" one by excessive elaboration. On the contrary following William of Occam he should seek an economical description of natural phenomena. Just as the ability to devise simple but evocative models is the signature of the great scientist so overelaboration and overparameterization is

not to be recommended. So when considering adding new parameters to the model, one should establish a clear purpose in doing so. At a minimum, one would have to show that the existing simple model was inadequate in some respect and that fixing this inadequacy requires adding at least one more parameter. But Tunaru offers no such demonstration.

Let's try to reconstruct an argument that might underlie his thinking. A point that Tunaru stresses repeatedly is that property prices are autocorrelated, so we get at least one additional parameter from that. In our earlier analysis this parameter took the form of the Hurst Exponent, but it could equally take the form of an autocorrelation coefficient, a parameter more familiar to economists. We prefer to work with the former because it leads more naturally to the solution to the problems posed by autocorrelation (see Chapter 10 above), but in principle one could work with either. Even so, it still does not follow that the Black '76 model needs to be replaced by some alternative model that can handle autocorrelation, such as the ARMA-EGARCH model. Why? Because we can *still use* Black '76 in an autocorrelated context, i.e., because Black '76 can handle autocorrelation too. As we have stated before (see Chapter 10), we can handle an autocorrelated underlying in Black '76 provided that we adjust the volatility calibration fed into the model. So autocorrelation does exist and does have an impact, but does not require that

---

Autoregressive Conditional Heteroskedasticity," *Journal of Econometrics* 31(3): 307–327; and D. B. Nelson (1991) "Conditional Heteroskedasticity in Asset Returns: A New Approach," *Econometrica* 59 (2): 347–370.

[110] This model is not new to the NNEG literature. Its first appearance in an academic journal is J. S.-H. Li, M. R. Hardy and K. S. Tan (2010) "On Pricing and Hedging the No-Negative Equity Guarantee in Equity Release Mechanisms," *Journal of Risk and Insurance* 77 (2): 499–522.

we replace Black '76 with some other model with more parameters than Black '76. The Box/Occam's Razor principle then kicks in and we are led to the conclusion that we should go with Black '76 combined with our proposed autocorrelation 'fix'. Put differently, we don't need to throw the BS out with the bathwater; we just need to get the volatility calibration right. The ARMA-EGARCH model fails the Box test.

Inconsistent on Market Consistency

Professor Tunaru repeatedly criticizes the 'risk-neutral' pricing methodology that underpins standard option pricing formulas, e.g.,

> "In the absence of an underlying market, liquid and free of counterparty risk, it is not possible to have a direct risk-neutral approach." (p. 11)
>
> "Unfortunately, Black (1976) model cannot be applied in the current context for the NNEG market since there is no futures house price contract currently traded in the UK." (p. 14)
>
> "In our opinion, the [Black '76] formula … simply does not apply for house prices and this is unrelated to the Gaussian distribution assumption behind the GBM model. The forward contract on a house price cannot be calculated as in the [Black '76] simply imitating the no-arbitrage formula for a stock paying dividend, where the dividend yield is replaced by the net rental rate. That formula cannot work because currently we cannot shortsell the value of a house. Hence, the no-arbitrage principle does not apply here to lock in the forward price as in the case of corporate stock." (p. 14)

We have dealt with these and other criticisms of Black '76 (or the MC approach more generally) in Chapter 22.

As an alternative, he then proposes (pp. 16-17) an alternative risk-neutral pricing rule based on an Esscher transform.[111]

There is an inconsistency here. Having repudiated option pricing theory because of incomplete markets, inability to take short positions etc, he then recommends the Esscher transform risk-neutral pricing rule. However the arguments for the existence of any risk-neutral pricing rule still hang on much the same complete market / costless shorting assumptions he rejects elsewhere.

To spell it out, Tunaru dismisses Black '76 because it relies on the standard 'risk-neutral' pricing methodology, which relies on a bunch of 'unrealistic' assumptions that he does not like. He then invokes his preferred Esscher transform pricing methodology, but that pricing methodology relies on much the same set of assumptions that he has already

[111] See F. Esscher, (1932) "On the Probability Function in the Collective Theory of Risk," *Skandinavisk Aktuarietidskrift* 15 (3): 175–195; or H. U. Gerber and E. S. W. Shiu "Option Pricing by Esscher Transforms," *Transactions of the Society of Actuaries* 46: 99–191.

rejected. But if those underlying assumptions can be invoked to justify his preferred model, then they can also be invoked to justify Black '76, and in that case his arguments against Black '76 fall away. On the other hand, if those arguments cannot be invoked because markets are incomplete, then Tunaru can't invoke those arguments for the ARMA-GARCH model either. In that case he has no pricing methodology.

The bottom line is that either you believe in the applicability of a 'risk-neutral' pricing approach based on market completeness or you don't. If you do, you can use it and if you don't, then you can't use it. But you cannot reject one pricing approach because it relies on assumptions you don't like and then turn around and propose an alternative pricing approach that relies on the same assumptions that you have just rejected.

Insufficient Volatility

There are also issues about Tunaru's calibration and let's start with the volatility. Tunaru begins by setting out some Maximum Likelihood (MLE), Generalised Method of Moments (GMM) and Method of Moments (MM) volatilities. These he reports in his Table 1:

Table 1: Parameter estimates (annualised) for the GBM process applied to the monthly Nationwide index, between Jan 1991 and Sep 2016, and Halifax Jan 1983 and Dec 2014

	Nationwide		Halifax	
Method of Estimation	μ	σ	μ	σ
Maximum Likelihood (MLE)	5.36%	3.94%	5.80%	3.96%
Generalized Method of Moments (GMM)	3.33%	3.84%	6.45%	2.27%
Method of Moment (MM)	5.36%	3.94%	5.88%	3.96%

Notes: Tunaru (2019, Table 1).

We would describe these estimates as referring to annualized spot volatilities, as opposed, e.g., to any forward volatilities. On the basis of these estimates, he proposes a baseline volatility of 3.9%.

He then presents a set of comparable results for alternative sample periods and UK regions based on the Nationwide quarterly house price index, over a full sample period of 1971 to 2018. These results indicate that "a range of values between 3.85% to 6.5% seems representative for GBM volatility parameter" (p. 1, referring to his Table 2). One is immediately struck by the fact that his baseline volatility is just next to the minimum of the volatility range: one usually puts baseline parameters somewhere in the middle, because otherwise the range becomes redundant. It is then strange, in our view, that he does not revise his baseline volatility estimate upwards in light of the results in his Table 2.

And that, essentially, is it, as far as his volatility analysis is concerned.

He goes on to argue against de-smoothing for house prices. De-smoothing is an alternative approach to our Hurst exponent analysis for dealing with autocorrelated house prices.[112] Curiously, however, he accepts that desmoothing can be useful for CRE. He then explains, "This point [why one should not de-smooth for house prices] is important since in the NNEG literature 10% volatility is taken as indicative for the UK. Based on the results in Table 2 we can see that a value of 10% is already a very conservative stressed upwards estimate." (p. 21) Elsewhere (p. 1) he elaborates by stating that "10% or 13% is then more of a stressed scenario value."

In short, he takes the spot volatility of 3.9% and recommends that this volatility estimate be used as the single volatility input for all NNEG puts, and he regards a figure of 10% or over as more of a stressed value. Implementing this low volatility into his NNEG model then helps produce the (very) low NNEG valuations that he reports.

One error in this treatment is that it confuses the volatility of individual houses with the volatility of an index. As we have seen in Chapters 9 and 10 (and see especially Figure 9.2 reproduced below), the volatility of house prices is much greater than the volatility of the index:

Figure 9.2: Indexed vs. Achieved House Prices

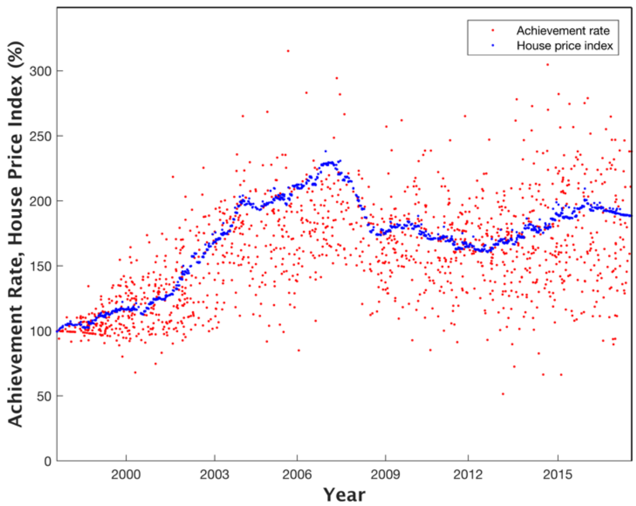

Source: SAMS.

So it seems to use that Tunaru should have picked a higher spot vol than the 3.9% that he selected based on his Table 2 results and then made a further adjustment to allow for volatility around the index.

But even without doing anything especially ambitious, he could have followed the PRA in its analysis underpinning section 2.16 in CP 13/18 and used some square root rule or, better, some Hurst Exponent extrapolation rule, to get from his spot volatility (which would have been > 3.9%) to some higher volatility.

In this context, it is also interesting that the PRA reports that the volatilities provided to it by firms were "generally in the range 10% - 15%," so his baseline volatility is way out of line with those estimates too, and what makes that all the more odd is that he

[112] For more on de-smoothing, see, e.g., P. Booth and G. Marcato (2004) "The Measurement and Modelling of Commercial Real Estate Performance," *British Actuarial Journal* 10(1): 5–61.

acknowledges that his baseline calibrations are also "selected based on discussions with experts working on ERMs and using public available tables from Legal & General, Just Group and Equity Release Council" (p. 27). Our point is the disconnect between the firms reporting 10% to 15% to the PRA and Tunaru going for 3.9% based on advice from presumably much the same set of experts.

The Rental Yield and the Tunaru Multiplier

A final issue is Tunaru's unusual way of estimating rental yields. Tunaru starts (p. 31) with an estimate of the mean gross rental yield of just under 5.2%. He then notes that less than 20% of properties are rented out and concludes:

> This means that a rough calculation would give a total rental yield, weighted by the 20% representing the actual renting market, of 1.03% (5.1776% × 20%) per annum. (Tunaru, 2019, p. 32)

This conclusion does not follow, however. As Andrew Smith observes in private correspondence:

> Th[is] rental yield analysis just seems wrong to me - arguing, as far as I can tell, that only 20% of the market is rented, so the rental yield on the market as a whole is a fifth of the yield on rented properties. That implies that owner-occupiers place no value on their right to occupy their own property. *I can't see any perspective from which this makes sense.* (Our emphasis)

It is also apparent from the discussion at Staple Inn that there was little support for the Tunaru multiplier there either. To be fair, Tunaru acknowledges the point. In a footnote he writes:

> It has been debated with other academics and market practitioners who are not entirely convinced about the weighting being applied. (Tunaru, 2019, p. 32, note 16)

We are not entirely convinced either. One of these was his own colleague, Dan Alai, who raised the issue at a Kent seminar on 28 January 2019:

> I was wondering why you multiply the rental yield by the proportion of properties that are rented out. In other words, why is 5.1776% divided by 5. I just do not see how it is relevant whether other properties are being rented out or not in determining the appropriate rental yield for a certain property. (Quoted in Tunaru, 2019, p. 74)

Dr. Alai is correct. The economic rental – that is, the use value, the value of the 'roof over one's head,' etc. – is still enjoyed by someone (or potentially enjoyable even if the property is void) regardless of whether the property is rented out or not. *What matters is that the economic rental is valuable, not whether the property is actually rented out.* Even if there is an owner-occupier, the property still has use value and the best estimate for the market value of that use value comes from the rents for similar properties currently prevailing in the property rental market, *not* those rents multiplied by 20%.[113]

Remember that what we are trying to do here is use the rental yield on rented properties to come up with an approximate calibration for the implied rental yield on ERMed properties. The proportion of properties that are rented out is irrelevant. Imagine, for example, that I have just ERMed my house, but the identical house (plus garden etc) next door has just been rented out at a given rental rate. Then I can estimate the value of the rental services on my house from the actual rental on the house next door, and my estimate of the value of those rental services will be 100% of the value of next door's rental. It is as simple as that.

In any case, the claim, indeed the whole report, completely fails to engage with the rationale given by CP 13/18 (para 3.16, p. 19), that the only difference between a contract for immediate possession and one for deferred possession, is *the value of foregone rights* (e.g. to rental income or use of the property) during the deferment period. You will pay less for deferred possession because you will lose the income that you could get by renting the property out or, alternatively, you will lose the use benefit that you could get by living in the property. Why would you rob yourself by pretending that you have only lost 20% of that income or use? You have lost *100%* of what you have lost.

If this argument is not clear then consider the following reductio argument: by Tunaru's logic, if the proportion of rented properties were to fall to say 0.001%, then the rental yield to be used in ERM calculations would be 5.1776% × 0.001% = 0.005%, effectively 0%. A rental yield of about zero cannot be correct. Why? Because the rental yield on my house cannot be virtually 0% of the rental yield on the identical house rented out next door. Why would you pay the same for a property that you could not take possession of for another 20 years, when you could buy a similar property now and have the use of it for the same 20 years? And if the Tunaru argument pushed to its limits gives an (even more) obviously incorrect result, then the argument itself must be faulty.

Tunaru then subtracts about 36% of gross rental to get the net rental – we have no strong argument with that calculation – and arrives at a net rental rate of 1.06% × 64% = 0.66%. Our best estimate was 4.2%.

---

[113] Others have said much the same thing. To quote David Rule's 'Dear CEO' letter of 3 April 2019: "One of the most financially significant parameters is the deferment rate, which the research estimates by considering rental yields. Several commentators have already challenged the research's judgement to multiply the rental yield by a factor (currently 20%) representing the proportion of properties rented out – to its credit, the research highlights a challenge to this judgement made by an academic reviewer. The PRA's own view is that the challenges are well-founded, the justification for the 20% factor is *not persuasive*, and that it is necessary to consider the benefits of owner-occupation on properties that are not rented out (such as those on which ERMs are written)." (Our emphasis) https://www.bankofengland.co.uk/-/media/boe/files/prudential-regulation/letter/2019/solvency-ii-equity-release-mortgages-part-2-apr-19.pdf See also Turnbull (2019), who suggests that the Tunuru multiplier "is best left behind" (C. Turnbull, "On the Actuarial Treatment of Equity Release Mortgages." https://www.linkedin.com/pulse/actuarial-treatment-equity-release-mortgages-craig-turnbull/ 13 June 2019).

Summary: Tunaru Does Not Work

The Tunaru approach produces NNEG valuations similar to those produced by the DP approach, but without the reliance on forecasts or incredible net rental rates that are well below zero percent. But it has major weaknesses: (1) It uses an overly parameterized model, when a simpler model such as Black '76 would have sufficed. Given the assumptions Tunaru makes and especially about autocorrelation in house prices, Black '76 would work, provided the volatility calibration was appropriate to the autocorrelation. (2) His arguments against Black '76 also undermine his own ARMA-GARCH model as well. (3) His recommended volatility is way too low and (4) his recommended net rental/deferment rate is based on an elementary error and is one fifth of what it should be. Strip away the first two errors and he may as well have used Black '76. Strip away the second two errors and he would have ended up with Black '76 based on our baseline calibrations.

Chapter Twenty-Four: Just Group's Deferment Rate Calibrations

In the previous chapter we stated that Just Group had used a 4.25% hpi assumption when valuing its NNEG. When we made a presentation on equity release to the London School of Economics on 1 October 2018 one equity release analyst in the audience suggested that we had made the number up.[114]

The facts are easily verified, however: the firm reported using this number in both its 2016 and 2017 *Annual Reports* (see pp. 163 and 110 respectively) and the relevance of this number in these reports is also clear, because it is the IFRS reports that shareholders would be interested in. The same number also appears in its 2018H1 results (p. 18). However in their 2017 SFCR, the firm reports an explicit q rate of 0.5% (p. 54). The latter accompanies an almost £1 billion hit to their balance sheet that is offset behind transitionals (Buckner, 2018a, b).

Page 83 of the firm's 2017 Solvency and Financial Condition report reconciles the statutory with the regulatory balance sheet. The almost £1bn figure appears as the change in 'other valuation differences', from end 2016 to end 2017. However, this almost £1 billion loss does not make a significant impact on capital because it is largely offset by an increase in the PRA transitional arrangement, which is an entry on the asset side of the regulatory balance sheet that can be used to create extra regulatory capital. It is a puzzle why this latter item (which is meant to be slow-moving and declining over time) should have increased so much over just one year. It has been put to us that the firm increased this item merely to hide the hit to its capital, but we find it difficult to believe that a reputable firm like Just would have resorted to such dissimulation, so there must be some innocent explanation that we are unaware of.

The situation for 2016 is relatively straightforward. To quote its 2016 *Annual Report*:

> When calculating the value of the no-negative equity guarantee on the lifetime mortgages, certain economic assumptions are required within the variant of the Black-Scholes formula. …
>
> In the absence of a reliable long-term forward curve for UK residential property price inflation, the Group has made an assumption about future residential property price inflation. This has been derived by reference to the long-term expectation of the UK retail price inflation, "RPI", (consistent with the Bank of England inflation target) plus an allowance for the expectation of house price growth above RPI (property risk premium) less a margin for a combination of risks including property dilapidation and basis risk. This results in a single rate of future house price growth of 4.25%. (p. 163)

The natural reading is that they are using the Black 76 formula (which takes forward prices, not spot) using an hpi of 4.25%. This reading suggests that they are taking the Black forward rate, which should be equal to risk free minus the deferment rate:

[114] The seminar is reported in Dowd, 2018b.

(24.1) $$f = r - q$$

and replacing it (incorrectly) with the forecast hpi, i.e.,

(24.2) $$hpi = r - q$$

as per the bad old 'discounted projection' approach that no-one should be using. Rearranging, we get an implied q:

(24.3) $$q = r - hpi$$

If we assume that $r = 1.5\%$ then we get

(24.4) $$\text{implied } q = 1.5\% - 4.25\% = -2.75\%.$$

We can then criticise this implied q value as making no sense.

The situation for 2017 is more involved, however. Their 2017 *Annual Report* states

> The return on equity release assets is adjusted to allow for the risks associated with these assets – namely, the potential shortfall resulting from the No-Negative Equity Guarantee ("NNEG"). The Group calculates the shortfall in respect of the NNEG using a variant of the Black-Scholes option pricing model. Inputs required (e.g. current house prices, future house price growth and house price volatility) are derived from available market data. (p. 51)
>
> In the absence of a reliable long-term forward curve for UK residential property price inflation, the Group has made an assumption about future residential property price inflation based upon available market and industry data. These assumptions have been derived with reference to the long-term expectation of the UK retail price inflation, "RPI", (consistent with the Bank of England inflation target) plus an allowance for the expectation of house price growth above RPI (property risk premium) less a margin for a combination of risks including property dilapidation and basis risk. An additional allowance is made for the volatility of future property prices. This results in *a single rate of future house price growth of 4.25%*, with a volatility assumption of 12% per annum. (*ibid*, p.110)

This passage is consistent with the previously quoted passed from their 2016 *Annual Report*. Going through the same calculations as before, we then get the same implied $q = 1.5\% - 4.25\% = -2.75\%$. So far, so same.

Then the fun starts. In their 2017 SFCR, the firm claimed to be using a deferment rate of 0.5%:

> As at 31 December 2017, the Board considers the Matching Adjustment in the Group's balance sheet in respect of LTM notes satisfies the principles of SS3/17 giving rise to an implied property volatility of 12% and *a positive deferment rate of 0.5% on a risk neutral basis*. (2017 SFCR, p. 54)

The issue then is how to reconcile this implied $q = -2.75\%$ with the explicit $q = 0.5\%$ that they also claim to be using. The difference between the two is enormous.

Furthermore, the average q rate we are discussing here is a slow-moving variable, which cannot move much from one period to another. A jump of 325 basis points from -2.75% in one year to 0.5% in the next year is highly implausible, even leaving aside the fact that both q rates are way out of line with the empirical evidence set out in Chapter 9.

Now for the awkward bit. In its 2018H1 results (p. 18) the firm offers the following treatment of an implied HPI vs an 'actual' or explicit HPI:

> Implied HPI = actual HPI – volatility/dilapidation – effect of capital requirement – effect of securitisation = 4.25 % - 3% - 1.5% - 1.4% = -1.65%

which they round to -1.7%.

So the firm has gone from an explicit $hpi = 4.25\%$ to an implicit $hpi = -1.7\%$, a difference of close to 600 basis points!

The derivation of the implied HPI is also problematic. The adjustment for 'volatility' is odd, given that the Black formula *already* includes an explicit treatment of volatility, namely the direct input for volatility (which the firm tells us is 12%). If the firm is using a 4.25% hpi input for the forward rate calculation and a volatility input for the put valuation, then it would be wholly incorrect to include an additional volatility 'adjustment' as well. Likewise it is spurious to include the capital requirement, because the calculation is for the amount of capital *available*, not the capital *required*. The NNEG calculation is an input determining the amount of capital available only. So there appears to be some obvious double counting. The 'effect of securitisation' item is also strange and we have no idea what it is or why it is there.

We now rearrange (24.5) as

(24.6) $$r = q + hpi$$

and substitute $q = 0.5\%$ and hpi=-1.7% into (24.6) to obtain

(24.7) $$r = 0.5\% - 1.7\% = -1.2\%!$$

So in making an explicit assumption of $q = 0.5\%$ and going from explicit $hpi = 4.25\%$ to an implicit $hpi = -1.7\%$, the firm is also implying an astonishing $r = -1.2\%$, but that can't be right either.

What seems to have happened is this: In 2016, the firm used an expected hpi rate of 4.25% to model its NNEG, equivalent to using an implicit q rate of -2.75% or thereabouts. As we have repeatedly stated, this approach is manifestly wrong, because the q rate should be much higher. In 2017, the firm *again* used the 4.25% expected hpi rate of 4.25% to model its NNEG, but this time it introduced a series of (mostly inappropriate) extra items driving a wedge between this hpi rate and an implied hpi rate that is only consistent with the firm's assumed q rate of 0.5% if $r = -1.2\%$, which confirms that the firm's reconciliation of its explicit and implied HPI rates makes no sense. The derivation of the implied expected hpi rate and the corresponding 0.5% q rate is thus totally half-baked.

Table 24.1 summarises the firm's NNEG valuation approaches for 2016 and 2017 in terms of their (a) explicit parameter assumptions, (b) their implied parameters and (c) their errors.

Table 24.1 Key Parameters of Just's NNEG Valuation: 2016 vs 2017

	2016	2017
Explicit param	$hpi = 4.25\%$	$hpi = 4.25\%$ $q = 0.5\%$
Implicit param	$q = -2.75\%$	$hpi = -1.7\%$ $r = -1.2\%$
Error	$q \ll 0$	$r = -1.2\%$

In short, the firm has a highly original approach to its NNEG modelling that defeats our efforts to make sense of it.

Chapter Twenty-Five: Actuarial Standards

The IFoA have a fair amount of material on their website about actuarial standards and regulation. Here are some quotes from their website (but the italics are ours):

https://www.actuaries.org.uk/about-us
"Under our Royal Charter *we have a duty to put the public interest first*"

https://www.actuaries.org.uk/upholding-standards
"We regulate actuaries *in the public interest*."

https://www.actuaries.org.uk/about-us/stepping-out-shadows
"One third of the UK public say they understand our traditional role in navigating financial risk, but we believe that it's *our ethics and professionalism that set us apart*." We agree.

https://www.actuaries.org.uk/about-us/our-brand
"Our vision is for the Institute and Faculty of Actuaries (IFoA) to serve the public by ensuring that where there is uncertainty of future outcomes, *actuaries are trusted and sought after for their valued analysis and authority"*

"Integrity

- We are: *Doing the right thing* for the organisation, our members, the profession *and the public interest*
- By being:
    - Honest
    - Accountable, and
    - Professional."

https://www.actuaries.org.uk/upholding-standards/standards-and-guidance/actuaries-code
The principles of the **Actuaries' Code** include, and we quote:

1. **Integrity**: members will act honestly and with the highest standards of integrity
2. **Competence and care**: members will perform their professional duties competently and with care
3. **Impartiality**: *members will not allow bias, conflicts of interest, or the undue influence of others to override their professional judgement*
4. **Compliance**: members will comply with all relevant legal, regulatory and professional requirements, take reasonable steps to ensure they are not placed in a position where they are unable to comply, and will challenge non-compliance by others

The conflicts of interests page states:

"As one of the five key principles of the Actuaries' Code, impartiality is placed in sharp focus in the context of professional conflicts of interest, *actual or perceived*"

To quote the industry manual on NNEG valuation, Hosty *et alia* (2007): "For providers attempting to price the NNEG on a market consistent basis there is insufficient product margin in order to provide a competitive product …" (p. 30.). Section 7.3.3 explains that under a market consistent approach the product would not be profitable, whilst the discounted projection (aka real world) model "has produced a significantly lower cost" and is therefore to be preferred. Our point is that we cannot rule out the possibility that some people *might* perceive a conflict of interest here.

On the question of honesty, there is also a legal standard namely, what any ordinary person would reasonably regard as dishonest. Lord Lane CJ set out the well known 'Ghosh Test' or *two limb approach* to the issue of dishonesty in R v Ghosh 1982. The test is (i) whether according to the ordinary standards of reasonable and honest people what was done was dishonest, and (ii) If so, did the defendant realise that what was done was *by those standards* dishonest.

Justice Cooke applied the test in the well-known case against Tom Hayes, the trader who was convicted for the manipulation of LIBOR.[115] Cooke ruled that the standard for dishonesty is absolute, and cannot change by reference to market standards or market ethos, standard practice in an industry or any common understanding amongst employees.

There is no authority for the proposition that objective standards of honesty are to be set by a market. Quite the contrary:

> The history of the markets have shown that, from time to time, markets adopt patterns of behaviour which are dishonest by the standards of honest and reasonable people; in such cases, the market has simply abandoned ordinary standards of honesty. Each of the members of this court has seen such cases and the damage caused when a market determines its own standards of honesty in this way. Therefore to depart from the view that standards of honesty are determined by the standards of ordinary reasonable and honest people is not only unsupported by authority, but would undermine the maintenance of ordinary standards of honesty and integrity that are essential to the conduct of business and markets.

Actuarial standards are summarised in the following four documents, from which we reproduce key passages.

The first is APS X1:

[115] https://www.judiciary.gov.uk/wp-content/uploads/2015/12/r_-v_tom_alexander_william_hayes_redacted_approved.pdf

APS X1: Applying Standards to Actuarial Work[116]

"*8.2. Members must be able to justify the standards applied (and/or not applied) to their Actuarial Work, if reasonably called upon to do so.*

9.1. *A failure to comply with this APS may result in a finding of misconduct in terms of the IFoA's Disciplinary Scheme*."

APS X2: Review of Actuarial Work[117]

The second is APS X2, which applies to the review of actuarial work, and which is relevant to the IFoA/ABI working party's work on NNEG valuation:

"1.3. In considering for the purposes of paragraphs 1.1 and 1.2 whether and to what extent Work Review should be applied to a piece of work (including whether and to what extent Work Review should be in the form of Independent Peer Review), Members should have regard to all of the relevant circumstances, including the following:

1.3.1. the degree of difficulty of the piece of work and its complexity

1.3.2. the significance of the piece of work, including any financial, reputational or other consequences for the person(s) for whom the work is produced

1.3.3. whether the circumstances of the piece of work make it more likely that errors could be made

1.3.4. the reasonable expectations of the person(s) for whom the work is produced;

1.3.5. the extent to which judgement and/or analysis is required

1.3.6. *the application of other quality assurance controls to the piece of work*;"

We have tried and tried to elicit information – hard information, as opposed to boilerplate waffle – about the quality assurance processes used by the IFoA/ABI working party's on NNEG valuation, but no-one in any position of responsibility will take responsibility, even in private, let alone in public. What was the independent scrutiny process, who were the senior figures in the IFoA or Actuarial Research Council who signed off, etc? So as regards the quality assurance in this case, we can't work out what it was and no-one in the know will tell us.

"1.3.7. *the desirability of assuring public confidence in the quality of the work in question*."

[116] https://www.actuaries.org.uk/documents/aps-x1-applying-standards-actuarial-work
[117] https://www.actuaries.org.uk/documents/aps-x2-review-actuarial-work

So the question is how failing to answer concrete questions about the quality assurance process helps to assure public confidence in the quality of the work in question.

Technical Actuarial Standard 100: Principles for Technical Actuarial Work, Financial Reporting Council December 2016.[118]

"Technical Actuarial Standard 100: Principles for Technical Actuarial Work (TAS 100) promotes high quality technical actuarial work. *It supports the Reliability Objective that "users for whom actuarial information is created should be able to place a high degree of reliance on that information's relevance, transparency of assumptions, completeness and comprehensibility, including the communication of any uncertainty inherent in the information*.

How users can place "a high degree of reliance" on valuations produced by an approach which produces valuations that are impossible and close to an order of magnitude too low?

1*. Judgement shall be exercised in a reasoned and justifiable manner*; material judgements shall be communicated to users so that they are able to *make informed decisions* understanding the matters relevant to the actuarial information."

What exactly is "reasoned and justifiable" about the DP approach? And in what sense are decisions based on impossible valuations to be considered informed?

2. *Data used in technical actuarial work shall be appropriate for the purpose of that work so that users can rely on the resulting actuarial information*.

"appropriate", "rely" ...

2.1 *Data shall be relevant* for the purpose of the technical actuarial work."

Actuaries are using assumptions about hpi to price the forward in their NNEG valuation models, but hpi is irrelevant. See, e.g., PRA SS 3/17 (p13, para 3.17) which states: "It is important to note that views on future property growth play no role in preferring one contract over the other. Investors in both contracts will receive the benefit of future property growth (or suffer any property depreciation) because they will own the property at the end of the deferment period. Hence expectations of future property growth are *irrelevant* ... "

So in what sense is an irrelevant variable relevant?

"3. *Assumptions used, or proposed for use, in technical actuarial work shall be appropriate for the purpose of that work so that users can rely on the resulting actuarial information*."

118 https://www.frc.org.uk/getattachment/b8d05ac7-2953-4248-90ae-685f9bcd95bd/TAS-100-Principles-for-Technical-Actuarial-Work-Dec-2016.pdf

How can it be appropriate to use an incorrect approach that depends on an irrelevant variable? And how are results based on an inappropriate assumption about an irrelevant variable reliable for users?

"Technical Actuarial Standard 200: Insurance," Financial Reporting Council December 2016[119]

The fourth is TAS 200, which applies to insurance.

"8. Measures, assumptions and judgements used to derive any estimates described as "best estimate", "central estimate" or other similar terms shall be neither optimistic nor pessimistic and shall not contain adjustments to reflect a desired outcome."

See the conflict of interest discussion above on the importance of profitability concerns.

Finally, some advice on what the IFoA might do:

> *If we make mistakes we want to put things right*. By monitoring any concerns raised, including any formal complaints, and by taking prompt corrective action where necessary, we seek to learn from where things have gone wrong and improve the standard of our service for future users.[120]

That quote comes from the IFoA's document "Putting things right."

[119] https://www.frc.org.uk/getattachment/c866b1f4-688d-4d0a-9527-64cb8b1e8624/TAS-200-Insurance-Dec-2016.pdf

[120] https://www.actuaries.org.uk/documents/putting-things-right.

Chapter Twenty-Six: Accounting Standards

This chapter sets out the accounting basics as they apply to equity release valuation.

Fair Value

For UK-recognised insurance firms, the valuation of the regulatory balance sheet is set out in the Solvency II regulations as transposed into UK law. This process is strictly rules based. The valuation of the statutory balance sheet, by contrast, is governed by accounting standards such as IFRS ('International Financial Reporting Standards') which tend to be more principles-based. Under modern accounting standards such as IFRS, valuations must be based on the principle of 'fair value'. IFRS defines a "fair value" price as:

> The price that would be received to sell an asset or paid to transfer a liability in an orderly transaction between market participants at the measurement date.[121]

IFRS does not define "fair," but the assumption is that a market participant, i.e. someone who is independent, knowledgeable, able and willing to enter into the transaction, would not be duped into an *unfair* transaction. So current market prices must deemed fair, because a market participant would not be duped into buying at greater than the market price, or be duped into selling at less than the market price. Consequently fair value equals market price, where the market price exists.

But what is fair value if the market price does not exist?

The answer comes from the Level 1/Level 2/Level 3 fair value hierarchy.

Level 1 fair value is the market price, where the market price exists.

Where no Level 1 fair values exist, i.e. where there are no market prices, IFRS uses Level 2 fair values: these are the prices of related instruments that can be used as proxies for unobservable values. An example in the equity release context would be the use of leasehold and freehold market prices as proxies for the values of the notional "leasehold" granted to the equity release borrower when an ERM is taken out.

Where no Level 2 prices are available, IFRS uses Level 3 or mark-to-model fair values, i.e., Level 3 involves the use of a model to obtain fair values. However, the model and its calibrations should still reflect "the assumptions that market participants would use when pricing the asset or liability, including assumptions about risk." In the equity release context, the natural example is a NNEG model. Such a model, which is by definition mark to model, would require under Level 3 to be calibrated using assumptions that market participants would make. One such assumption would be Principle II that no

[121] See, e.g., https://www.iasplus.com/en/standards/ifrs/ifrs13.

value of the ERM can exceed the value of forward contract (see (19.1) above); another is Principle III, that the deferment house value must be less than the current spot house value, reflecting the point that a market participant would want compensation for the income or use that was lost through deferment.

Whatever level is used, the underlying principle is always the same. To quote Financial Reporting Standard 102:

> 2.2 The objective of financial statements is to provide information about the financial position, performance and cash flows of an entity that is *useful for economic decision-making by a broad range of users* who are not in a position to demand reports tailored to meet their particular information needs.[122]

This information should enable users to take a neutral and objective view of the company and ensure that they are not being cheated.

Qualitative of this information include understandability, substance over form, completeness, comparability and timeliness, as well as (quoting FRS 102):

- Relevance: 2.5 The information provided in financial statements must be relevant to the decision-making needs of users. Information has the quality of relevance when it is capable of influencing the economic decisions of users by helping them evaluate past, present or future events or confirming, or correcting, their past evaluations.
- Materiality: 2.6 Information is material—and therefore has relevance—*if its omission or misstatement, individually or collectively, could influence the economic decisions of users taken on the basis of the financial statements*.
- Reliability: 2.7 The information provided in financial statements must be reliable. Information is reliable when it is *free from material error and bias and represents faithfully that which it either purports to represent or could reasonably be expected to represen*t. Financial statements are not free from bias (ie not neutral) if, by the selection or presentation of information, they are intended to influence the making of a decision or judgement in order to achieve a predetermined result or outcome.
- Prudence: 2.9 The uncertainties that inevitably surround many events and circumstances are acknowledged by the disclosure of their nature and extent and by the exercise of prudence in the preparation of the financial statements. Prudence is the inclusion of a degree of caution in the exercise of the judgements needed in making the estimates required under conditions of uncertainty, *such that assets or income are not overstated and liabilities or expenses are not understated*. However, the exercise of prudence does not allow the deliberate understatement of assets or income, or the deliberate overstatement of liabilities or expenses. *In short, prudence does not permit bias*.

[122] Financial Reporting Standard 102 The Financial Reporting Standard applicable in the UK and Republic of Ireland." Financial Reporting Council, September 2015. https://www.frc.org.uk/getattachment/e1d6b167-6cdb-4550-bde3-f94484226fbd/FRS-102-WEB-Ready-2015.pdf

The accountant is then hired by the management of the company to draw up accounts on the basis of these principles and in accordance with IFRS rules and existing law (e.g., the Companies Act). These accounts would be approved by the directors, who are deemed to have prepared the accounts, and then presented to the auditors ideally for sign-off. "An auditor is an independently qualified person who is appointed to give shareholders an independent, professional and informed opinion on the financial statements prepared by the directors"[123] and the

> auditor's objectives are to obtain reasonable assurance about whether the financial statements as a whole are free from material misstatement, whether due to fraud or error, and to issue an auditor's report that includes the auditor's opinion. Reasonable assurance is a high level of assurance .... *Misstatements* can arise from fraud or error and *are considered material if, individually or in the aggregate, they could reasonably be expected to influence the economic decisions of users taken on the basis of these financial statements*.[124] (Our emphasis)

Improving on Fair Value?

There are a number of common objections to fair value. The first relates to the issue of whether people should seek to 'improve' on the fair value/market value price. The short answer is "no."

For example, actuaries sometimes claim that current market values should be ignored because they are currently too low or too high, relative to the actuary's judgement of what the "long-term" price should be. As David Wilkie once put it, "The actuary is ... saying that the market has temporarily got it wrong, but that, in due course, it will get it right."[125] So the actuary is suggesting that the market price should be replaced by a non-market valuation based on actuarial judgement or, if you want to put it that way, by an implied (not even explicit!) "actuarial forecast."

Objection #1: At least if the forecast were explicit we could scrutinise the methodology on which it is based and come to an informed view of its merits. However, all we have to go on here is a nebulous 'actuarial judgement'. Basing a critique of market price valuations on an "I know it's wrong" gut feeling is wrong on principle.[126]

Objection #2: It is doubtful that a reliable forecast exists. We can predict eclipses, the reaction of hydrogen and oxygen to a flame, the acceleration due to gravity and so forth,

---

[123] http://www.corplaw.ie/blog/bid/337442/The-Responsibility-Of-Auditors. Dean link doesn't work.

[124] FRC "Description of the Auditor's Responsibilities for the Audit of the Financial Statements Applicable for Audits of Financial Statements for Periods Commencing on or after 17 June 2016." https://www.frc.org.uk/auditors/audit-assurance/auditor-s-responsibilities-for-the-audit-of-the-fi/description-of-the-auditor%E2%80%99s-responsibilities-for

[125] David Wilkie on p. 549 of the Discussion of A. C. L. Dyson and C. J. Exley, (1995) "Pension Fund Valuation and Investment." *British Actuarial Journal* 1(5): 965-977.

[126] It lacks objectivity and is scientifically indefensible: "we claim to know the answer but cannot tell anyone else how to derive it in advance" as Tim Gordon eloquently put it. See T. Gordon (1999) "The Price of Actuarial Values." Paper presented to the Staple Inn Actuarial Society (16 February).

but science hasn't found a way to predict the path of market prices. The problem is that the market price of an asset *itself* involves a forecast, by the market, of future cashflows, so in trying to predict where the market will be in a year's time, we are trying to forecast a forecast. Instead of trying to predict the result of the next election, it is like trying to predict what the *Times* will predict it to be. Good luck on that.

Also, either the market price is the best forecast, or it is not. If the former, we can't improve on it. If the latter, we have to forecast what the bad forecast will be in a year's time. But which bad forecast do we choose and how we select it?

To go to the heart of the matter, we can be pretty confident *that* the market price will change all the time, but the problem is that we don't know *how* the market price will change from one period to the next. The market valuation might not be very good, but it's the best we have.

Objection #3: Even if we had perfect foresight, such as God might have, we would still have no leave to mark the value of an asset to anything other than the current market price. It may be that the market is in some sense 'wrong'. Clearly the market price must be 'wrong' most of the time, because it is changing all the time. Even so, if we mark an asset on a firm's books at higher than the market on the grounds that we have perfect foresight, or better judgment than the market, then we are defrauding *prospective* shareholders of the firm, because they would pay more for shares than they would have paid had we marked the shares to market. If we mark the value at lower than the market price, because our flawless judgment values it at less, then we are defrauding *existing* shareholders, because their shares would be valued at less than they would have been had we marked the shares to market. If God were an accountant, He would not value an asset differently from its market value, despite being omniscient, for God is also Perfectly Good, and so would not get involved in false accounting.

Well clearly, if even God would not superimpose His Judgement over that of the market, then there isn't much of a case for anyone else to superimpose his or her judgement over that of the market either.

The Liquidity Premium Fallacy

The argument is often made that if a firm holds an illiquid asset, a bond, say, and intends to hold that asset to maturity, then the firm is entitled to mark up the asset to 'capture' the liquidity premium[127].

One problem with this argument is that all we know in practice is that the asset has a spread over the risk-free, but whether that spread is a risk spread or a liquidity premium, or so much of one and so much of the other, we do not know. We only observe the spread,

[127] Note that this notion of a liquidity premium is quite different from the one familiar to economists. The economic notion of a liquidity premium refers the price difference between two otherwise similar assets, where one asset is liquid (i.e., easy to sell) and the other is not liquid or may prove to be difficult to sell in the future.

not the liquidity premium, and any claim we might make about the size of the liquidity premium is merely a hypothesis.

A counterargument we have sometimes encountered is that liquidity premiums have been reliably estimated in a paper by Webber and Churm published in the 2007Q4 *Bank of England Quarterly Bulletin*.[128] This paper presents a chart purporting to show that modelled corporate bond spreads are about 50% of the observable spread, from which they infer that the residual is 'illiquidity risk' that could be captured by holding the asset to maturity. However, we have done our own reconstruction and we found that after using a different leverage parameter, almost all the difference between modelled and observed spread disappears.[129] So it appears that they had misspecified the equation and then misidentified the residual from their misspecified equation as 'liquidity risk'.

But for the sake of argument, let's treat the liquidity premium as if Webber and Churm had correctly captured it and let's even grant the argument that the firm is justified in marking up the asset to capture that premium. Therefore, when a prospective shareholder comes to purchase the shares, he pays the premium. However, when the same individual buys the same asset in the market, he will pay less, because the market price of the asset is marked down by the same premium. Since the market price is the fair value of the asset, he is paying too much when he buys the share and is therefore being defrauded.

An example of this 'liquidity premium argument' is to be found in remarks made by Andrew Rendell at the 28 February Staple Inn event for the Tunaru report:

> … if you have a corporate bond, is the economic worth *to the insurer* the same as it is to everybody else, arguably it isn't, and the reason for that being that *a typical market participant* will put a discount to the price that they would be prepared to pay for it, because that corporate bond has risks around liquidity, and it has risks around price volatility over the duration of the asset.
>
> The insurer says, "Well I don't care about that, because I'm going to hold my asset to maturity, and therefore I don't need that discount, so the corporate bond is worth more to me than it is to a typical participant." …
>
> So the question then is how does that map through to the ERM, in particular the property side of ERMs, so … what's the economic worth of that property?
>
> … you are not going to see any cash out of that asset for some time, you can't sell it, so to many investors that would be quite a significant disadvantage, but maybe less so *to an insurer* that has long term liabilities and it can *just wait for that value to emerge*.[130]

There are three problems here. [1] Mr. Rendell talks about the economic worth *to the insurer*, but who is the insurer? The insurer isn't a person. The insurer is a company which

---

[128] L. Webber and R. Churm (2007) "Decomposing Corporate Bond Spreads." *Bank of England Quarterly Bulletin* 2007(Q4): 533-541.
[129] See D. Buckner (2019) "From the Postbag – Merton Model Part II." *The Eumaeus Project* (19 June).
[130] https://youtu.be/DdLRqcIvR20?t=3794

has shareholders, so we should be asking about the economic worth to the shareholders. So what is the economic value *to the shareholders* of an asset where you have to "wait for that value to emerge" as he puts it? Well, the economic value to shareholders is just the same as if they weren't shareholders at all, but were buying in the open market. If the management of the insurance company decides that a prospective shareholder must pay *now* for the value that *will* emerge later, so that they can get a whacking great bonus, then they have defrauded the prospective shareholder, who must pay more than the market price. [2] As for "waiting for the value to emerge", the management (or existing shareholders) haven't waited at all. They have crystallised the value now and thereby robbed future shareholders of the rewards they were waiting patiently for. Those who believe in "waiting for the value to emerge," should wait for the value to emerge. [3] Finally, it is the current or prospective shareholder or investor, not just the company, to whom the accountant has a fiduciary duty to report a fair value.

If we still haven't persuaded you, consider the following argument.

We buy the right to possession of a property whose estimated vacant value is £1m, i.e. a freehold encumbered by a leasehold. The lease is for 99 years. Based on available data (e.g. from the Land Tribunal) the value of the deferment would be about £30,000.

We sell the same right to you for £1m because you think there is a massive illiquidity premium that can be marked up now.

We have then made an immediate massive profit (£970k) from this deal. We have just earned an illiquidity premium that would have otherwise have taken us 99 years to earn, and you have paid away the same thing! You will earn nothing from the liquidity premium, which has all gone to us. Why should you pay the premium to someone else?

You may argue that it is different with an insurance company, because an insurance company is there for the long term and has the patience to wait 99 years. Maybe so, but why would the insurance company mark up *now* a liquidity premium that must be earned over 99 years? You can only obtain the liquidity premium after 99 years, and if you sell the asset, you only get the current market price which does not include any liquidity premium. But above all, why would *anyone* pay £1m when they could pay £30k instead for the same asset in the market?

The fallacy is not in asserting that there may be an illiquidity premium that can be 'captured' by holding the asset to maturity. The existence and extent of any such premium is an empirical matter and (as the Webber-Churm example shows) is difficult to pin down. Instead, the fallacy lies in the belief that a liquidity premium to be earned in the future should be marked up now, above the currently prevailing market price of the asset.

Or consider this final argument. Suppose there exists a bond type asset that offers an illiquidity premium. So let's set up a company where we borrow long dated liabilities at risk free and invest the proceeds in these long-dated illiquid assets with their (certain) illiquidity premium. Persuade shareholders/PRA etc that an illiquidity premium that exists that 'it can't be arbitraged out'. Create a pile of equity by discounting liabilities at risk free + premium, pay yourself a lot of dividends or sell the company, and retire to the beach.

Congratulations! You have just arbitraged out the illiquidity premium, which no longer exists.

The 'Buy n' Hold' Fallacy

A related (and common) argument is that an insurer can mark up higher yielding assets because the insurer is (a) going to hold the assets to maturity, so (b) will not be exposed to changes in market value, and (c) will realise the additional return.

This argument might seem plausible but quickly unravels on examination:

First, how do we *know* that the insurer will hold the assets to maturity? We don't. It may *intend* to hold them to maturity, but it could be forced to sell later. Who knows?

Second, even if we grant that the asset will be held to maturity, it is (usually) false to claim that the holding entity is not exposed to changes in the market value of the asset. If the asset is a bond, the bond might default before maturity. If it is an equity, the maturity (of the annuity) may coincide with a downturn in the market, and if it is an ERM, maturity could be in the aftermath of a Japan style housing decline. The only assets that are not subject to risk of loss (in nominal terms) would be gilts, but gilts don't offer a return higher than risk-free.

Thus the argument is internally contradictory: if an asset offers a higher return than risk-free, then that return must be risky, so there is risk of loss at maturity; but if the does not offer higher than risk-free, because it is risk-free, then there is no higher return to realise. So you can have (b) or you can have (c), but you can't have both (b) and (c).

Prudential Regulation and Fair Value

Another class of objections to fair value is that fair value does not apply because prudential regulators have different objectives or are working to different standards. As an example, consider the following letter of 19 March 2019 that we wrote to Hans Hoogervorst, the chair of the International Accounting Standards Boards (IASB):

"Dear Mr Hoogervorst,

One of us (Buckner) wrote to you on 9 January 2019 raising the concern that some life insurance firms are valuing embedded guarantees in a way *radically* different from accepted financial theory, and different (in our view) from the way they would and should be valued under broad IFRS principles, i.e. how a market participant would value them.

The IASB response on 4 March 2019 to this concern was a puzzling one. The reply claims that 'differing valuations for prudential purposes are largely driven by the different objectives of the two measurement bases'.

> Prudential regulators focus on measures of regulatory capital that absorb losses, whereas accounting standard setters are concerned with "financial information about the reporting entity that is useful to existing and potential investors, lenders and other creditors in making decisions about providing resources to the entity" (IFRS Conceptual Framework OB2). Those different objectives and approaches to measurement will sometimes give rise to different valuations for the same instrument.

We have worked for more than 40 years between us in the areas of capital measurement and capital management and the view that there are differing valuations for prudential purposes is news to us. To be sure, prudential regulators are concerned with the adequacy of any given amount of capital. If one insurer has £10bn of low quality assets, another insurer has the same amount of high quality assets, and both have £9bn of pension liabilities at the same duration, then both have the same *amount* of capital, namely £1bn.[131] But the *adequacy* of that amount of capital is a different matter, and the objective of the prudential regulator is to assess capital adequacy by means of established techniques such as value at risk, stress testing etc. There is no equivalent of capital adequacy or capital requirement in IFRS, however.

The issue raised in the letter of 9 January was not about the *adequacy* of capital, but rather about the *amount* of capital, and in particular, about how that amount is *measured*. In this case the objectives of regulators and of accounting standard setters are identical, namely to value assets and obligations as a market participant would value them, and we are astonished that the IASB would go on record to claim any different. Insurers even provide a reconciliation of regulatory and statutory balance sheets in their financial reports, from which it is clear that, while there are minor differences, the purported objective of measurement is the same, i.e. arm's length valuation, fair value measurement etc. …

The IASB response goes on to say that accounting standard setters are concerned with "financial information about the reporting entity that is useful to existing and potential investors, lenders and other creditors in making decisions about providing resources to the entity." But can you please explain how is it useful to market participants if firms value assets and obligations in ways that a market participant would not? How then would investors know whether the reported values were fair/useful/reasonable or even legal, or not? If valuations are higher than those that a market participant would place on them, then prospective investors are being unfairly disadvantaged; if valuations are lower than those that a market participant would place on them, then existing investors are being unfairly disadvantaged. The only fair way to avoid either party being disadvantaged is to make the same valuations that a market participant would make.

Faithfully, etc"

We are still awaiting a reply.

---

131 We leave out details such as regulatory capital including subordinated debt. Such details are irrelevant in the present case.

The Fiduciary Principle

We then come back to the fiduciary principle. Even where market prices do not exist, accounting principles say that the accountant *should* value economically similar assets in the same way and imply that valuation should reflect rational investor preferences. To quote, e.g., IFRS 13 B14a: "Cash flows and discount rates *should* reflect the assumptions that market participants would use when pricing the asset or liability." An accountant or auditor or some other person, who has an obligation of trust towards a less knowledgeable investor, must value an asset or liability as a rational knowledgeable investor (or market participant, or knowledgeable, willing independent person) would. This principle applies regardless of whether the accountant, auditor or whoever has private views about valuation that differ from fair value valuations. It also applies even if he or she has superior knowledge of the future: even if one had a perfect crystal ball, one is not allowed to use it to provide valuations that differ from fair value ones.

False Accounting

Those who report valuations are also required to desist from false accounting, the practice of which is a criminal offence. To be precise, the offence of false accounting is a subclass of the offence of theft[132] and is created by section 17 of the Theft Act of 1968 which states:

> 17.-(1) Where a person dishonestly, with a view to gain for himself or another or with intent to cause loss to another,-
>
> - *destroys, defaces, conceals or falsifies any account or any record or document made or required for any accounting purpose*; or
> - in furnishing information for any purpose produces or makes use of any account, or any such record or document as aforesaid, which to *his knowledge is or may be misleading, false or deceptive in a material particular*;
>
> he shall, on conviction on indictment, be liable to imprisonment for a term not exceeding seven years.
>
> (2) For purposes of this section a person who *makes or concurs in making in an account or other document an entry which is or may be misleading, false or deceptive in a material particular*, or *who omits or concurs in omitting a material particular from an account or other document, is to be treated as falsifying the account or document.* (Our emphasis)

We emphasise two points about this definition. The first is that to be guilty of the offence of false accounting it is not enough to report information that is or might be misleading.

---

[132] The offence of theft is defined under Section 1 of the Theft Act of 1968: "A person is guilty of theft if he dishonestly appropriates property belonging to another with the intention of permanently depriving the other of it …" See http://www.legislation.gov.uk/ukpga/1968/60/pdfs/ukpga_19680060_en.pdf.

One must also do so *dishonestly*, i.e., knowing that one's behaviour is dishonest by reasonable standards, and one must do so with the *intent* of personal gain or to deprive someone else of what is lawfully theirs. The other point is that the phrase "conceals or falsifies *any* account or *any record or document made* or required for *any* accounting purpose …" is highly encompassing, and would include public accounts, internal books and other documents such as, e.g., spreadsheets and internal memos. It also potentially includes knowingly making on-the-record statements that are anything but the full truth.

Consider also the following National Fraud and Cybercrime Reporting Centre statement on false accounting fraud:

> False accounting fraud happens when company assets are overstated or liabilities are understated in order to make a business appear financially stronger than it really is.
>
> False accounting fraud involves an employee or an organisation altering, destroying or defacing any account; or presenting accounts from an individual or an organisation so they don't reflect their true value or the financial activities of that company. …
>
> Some examples of false accounting fraud include:
>
> - an employee making inflated expenses claims
> - a customer or an employee falsifying accounts in order to steal money
> - an employee using false accounting to cover up losses built up through trading or fraudulent activity. …
> - at the extreme end of the scale, the fraud may mean that a company has incurred serious financial losses and/or is trading while insolvent. [133]

There are also related offences such as conspiracy (where two or more people plan to engage in another offence, e.g., false accounting)[134] and aiding, abetting, counselling and procuring the commission of another offence, e.g., false accounting.[135]

[133] /www.actionfraud.police.uk/a-z-of-fraud/false-accounting-fraud
[134] http://serious-crime-solicitors.co.uk/conspiracy.php
[135] http://www.legislation.gov.uk/ukpga/Vict/24-25/94/section/8

Chapter Twenty-Seven: Recommendations for Good Valuation Practice

Any reasonable approach to NNEG and ERM valuation must use a model that is fit for purpose and be based on reasonable (i.e., plausible and defensible) calibrations. These requirements narrow the field to some form of MC approach. The alternatives to an MC approach are the DP and Tunaru approaches, but neither of these meet the requirements of being fit for purpose or being reasonably calibrated.

We can implement an MC approach using a rehedging algorithm, Black '76 or the PRA's Principle II bounds. The latter two are the easiest to implement and difference between these two are:

- The bounds-based NNEG valuations will be lower than the Black '76 valuations and the bounds-based ERM valuations will be higher than ERM valuations based on Black '76.
- The Black '76 valuations will depend on the volatility calibration, but bounds-based valuations can be obtained without using any volatility calibration.

We emphasise however that any of these MC approaches is reasonable.

On the calibration, our recommended calibrations are:

- Loan to value ratio based on 'age minus 30' rule. This calibration is justified in Chapter 4.
- Risk-free rate r = 1.5%. This calibration is justified in Chapter 5.
- ERM loan rate $l = 5.25\%$ for current conditions. This calibration is justified in Chapter 6.
- Deferment rate $q = 4.2\%$. This calibration is justified in Chapter 8.

The volatility calibration is more involved. In principle one could use any of a range of volatilities, but the key is to choose volatility calibrations that are consistent with one's underlying position on the rehedging frequency, explicitly if one uses a rehedging approach to value the NNEG, and implicitly if one uses Black '76. Assuming one uses Black '76, then in principle one would use the Black '76 puts with volatilities for each decrement drawn from the volatility term structure we identified in Chapter 10.

To put the valuation formulas into mathematics, and using obvious notation:

(27.1) $$NNEG = \sum_t[exit\ prob_t \times NNEG_t] = \sum_t[exit\ prob_t \times put_t]$$

where

(27.2) $$put_t = e^{-rt}[K_t N(-d_{t2}) - F_t N(-d_{t1})]$$

(27.3) $$d_{t1} = [ln(F_t/K_t) + \sigma_t^2 t/2]/(\sigma_t\sqrt{t})$$

(27.4) $$d_{t2} = d_{t1} - \sigma_t\sqrt{t}$$

where σ_t is the volatility of the forward house price at maturity t.

However, since we are all used to working with a single volatility rather than a term structure, it is convenient to use a single volatility for all puts. The use of a single volatility is acceptable, *provided* one uses a single volatility calibration that is appropriate to the borrower's age and gender. To obtain such a calibration, one could use the expected volatility obtained by weighting each volatility by its exit prob or decrement probability. The formula for the expected volatility σ^e is then

(27.5) $$\sigma^e = \sum_t [exit\ prob_t \times \sigma_t]$$

Table 27.1 shows these expected volatilities against borrower age: for males 21.7% for age 55, 14.8% for age 70 (and hence our earlier baseline single volatility recommendation for males aged 70) and so on. These expected volatilities give results that are very close to the results that one would have obtained had one used the full volatility term structure.

Table 27.1: Expected Volatilities for Different Ages

Borrower Age	**Expected Volatility (Males)**	**Expected Volatility (Females)**
55	21.7%	22.9%
60	19.1%	20.0%
65	16.8%	17.8%
70	14.8%	15.7%
75	13.2%	13.9%
80	12.0%	12.5%
85	11.2%	11.5%
90	10.7%	10.8%

Notes: Based on the assumptions: LTV based on 'age-30' rule, r=1.5%, l=5.25% and q=4.2%. Exit probabilities are based on M5-CBD model projections using England & Wales deaths rate data spanning years 1971:2017 and ages 55:89.

Table 27.2 shows the resulting NNEG and ERM valuations for the Black '76 and PRA Principle II valuation approaches for a borrower aged 70:

Table 27.2: ERM and NNEG Valuations: Male Age 70

Valuation Approach	***NNEG***	***ERM***
Black '76 (using vol term structure)	£32.3	£42.5
Black '76 (using expected vol)	£32.2	£42.7
PRA Principle II bounds	£28.1	£46.8

Notes: $NNEG$ is the present value of the NNEG guarantee, and ERM is the present value of the Equity Release Mortgage. Based on the baseline assumptions: male aged 70, LTV=40%, r=1.5%, l=5.25% and q=4.2%. Exit probabilities are based on M5-CBD model projections using England & Wales male deaths rate data spanning years 1971:2017 and ages 55:89.

Table 27.3 gives the same valuations for age 70 as percentages of the initial loan amount, £40:

Table 27.3: ERM and NNEG Valuations as Percentages of Loan Amount: Male Age 70

Valuation Approach	*NNEG*	*ERM*
Black ’76 (using vol term structure)	80.8	106.3
Black ’76 (using expected vol)	80.4	106.7
PRA Principle II bounds	70.2	116.9

Notes: As per Table 27.2.

Table 27.4 gives the same results as in Table 27.3 for ages spanning 55 to 90:

Table 27.4: ERM and NNEG Valuations as Percentages of Loan Amount: Male Age 70

Valuation Approach	*NNEG*	*ERM*
Age 55		
Black ’76 (using vol term structure)	235.2	91.2
Black ’76 (using expected vol)	235.9	90.5
PRA Principle II bounds	213.3	113.1
Age 60		
Black ’76 (using vol term structure)	170.8	96.9
Black ’76 (using expected vol)	170.8	96.7
PRA Principle II bounds	154.0	113.7
Age 65		
Black ’76 (using vol term structure)	120.0	102.1
Black ’76 (using expected vol)	119.9	102.1
PRA Principle II bounds	106.9	115.3
Age 70		
Black ’76 (using vol term structure)	80.8	106.3
Black ’76 (using expected vol)	80.4	106.7
PRA Principle II bounds	70.2	116.9
Age 75		
Black ’76 (using vol term structure)	51.8	109.0
Black ’76 (using expected vol)	51.3	109.6
PRA Principle II bounds	43.2	117.7
Age 80		
Black ’76 (using vol term structure)	31.6	110.2
Black ’76 (using expected vol)	31.1	110.7
PRA Principle II bounds	24.7	117.1
Age 85		
Black ’76 (using vol term structure)	18.4	110.0
Black ’76 (using expected vol)	17.9	110.5
PRA Principle II bounds	13.0	115.3
Age 90		
Black ’76 (using vol term structure)	10.3	108.9
Black ’76 (using expected vol)	10.0	109.2
PRA Principle II bounds	6.4	112.9

Notes: As indicated, otherwise as per Table 27.2.

Figures 27.1 and 27.2 give the corresponding plots.

Figure 27.1: NNEG/Loan Ratios Vs Borrower Age

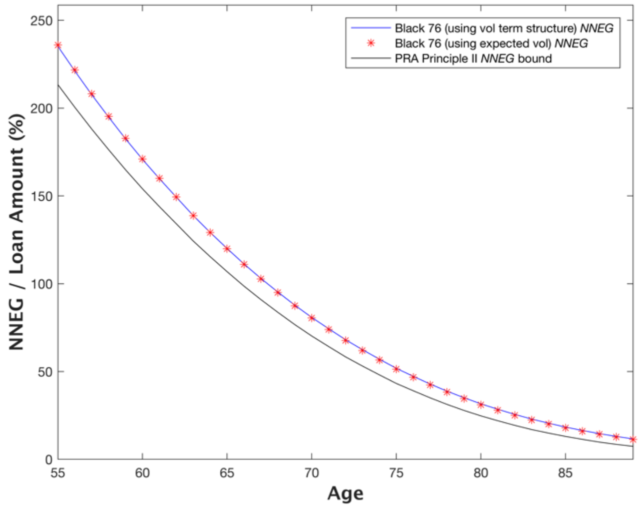

Notes: As indicated, otherwise as per Table 27.2.

Figure 27.2: ERM/Loan Ratios Vs Borrower Age

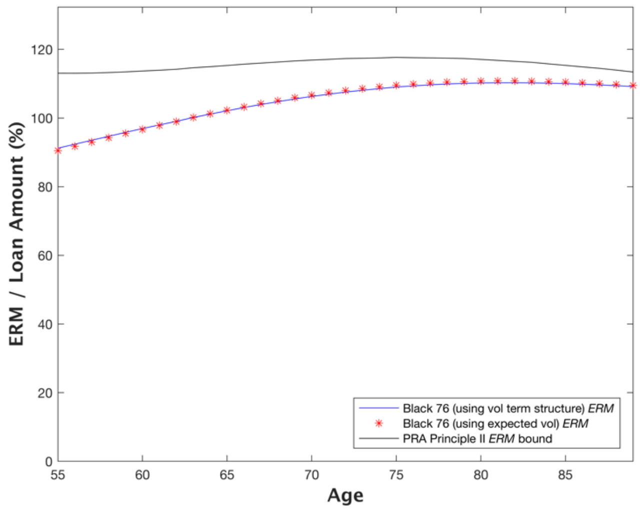

Notes: As indicated, otherwise as per Table 27.2.

One notices that the ratios of $NNEG$ to loan amount fall sharply with age. However, what is most significant is the low ratios of ERM to loan amount. Even the upper bounds are below 120% for any borrower age. For Black '76 they are below 100% for ages under 65 and barely touch 110% for the peak profitability age range which is 75 to 85. We might add that these numbers are gross of operating costs. These results suggest that ERM loans to younger borrowers are loss-making and that ERM loans to older borrowers are less profitable than is commonly thought.

These results suggest that lenders might do better if they stop lending to younger borrowers and the question then arises whether eking out low-return short-term ERM loans to older borrowers is the best possible use of shareholders' capital. As far as shareholders are concerned, it would appear that the term 'equity release' may be more appropriate than anyone has yet realised.

Chapter Twenty-Eight: Recommendations for Governance and Disclosure

Poor valuation practices are endemic in the equity release sector, but the IFoA has done nothing to condemn such practices. Indeed, the IFoA itself is on record as having endorsed a number of errors on NNEG valuation. This is hardly a healthy state of affairs.

The IFoA needs to put this situation right, and promptly. We recommend that it issue a corrective statement along the following lines:

> The IFoA acknowledges and regrets that serious errors have been made by a number of equity release actuaries in the valuation of No-Negative Equity Guarantees.
>
> The IFoA recognises that there are a variety of possible approaches to NNEG and ERM valuation, but it affirms the market consistency principle: any valid approach must be market consistent.
>
> The IFoA regards approaches that are non-market consistent as not meeting technical actuarial standards. In particular, it regards the so-called 'Real World' or 'Discounted Projection' approach as inherently fallacious and advises that practitioners should refrain from using it.

The regulator, the PRA, needs to take action too. We recommend that the PRA issue a new Policy Statement stating that equity release poses a prudential problem and confirming that it regards only MC approaches as being consistent with good actuarial practice. The PRA should also make clear what it regards good practice to be. Besides market consistency, good practice should cover issues such as acceptable modelling practices, plausible calibration and the importance of disclosure and transparency.

In addition, the PRA should make clear that good practice valuation methods should not be swayed by notions of profitability. If firms can't make the profits they want, or can't make profits at all, then they are free to exit the industry, but that shouldn't be the PRA's concern. The PRA's concern should be to ensure that the valuations are done properly. To quote an authoritative source on this very issue:

> No man can serve two masters: for either he will hate the one, and love the other; or else he will hold to the one, and despise the other. Ye cannot serve God and Mammon. (Matthew 6:24)

Regarding the proposals in CP 3/19, the PRA should monitor the net rental rate rather than the real interest rate, and it should disclose its methodology so that it is open to outside scrutiny.

The PRA would still face the problems associated with firms' gaming the regulatory system, so it needs to take action to minimise firms' scope for gaming. It could do so in the following ways:

- The PRA should impose a market consistency requirement on the deferment rate, with firms to provide a plausible defence of their assumptions. With the current rate of gross rental yields (2019 around 5%) we would find any rate of less than 3% difficult to justify.
- Firms should offer the PRA a plausible market-based defence of their volatility analysis and assumptions, demonstrating consistency with observed values of achievement dispersion, interest rate and deferment rate volatility, and correlation.

We understand that the PRA already requires firms to report NNEG and ERM bounds based on its Principles II and III, and will not accept any valuations that violate any of these bounds.

Thirdly, we have some advice for accounting and audit practice based on our evidence to the Brydon Review on the quality and effectiveness of audit:[136]

- Auditors should be encouraged to ensure, as far as possible, that shareholders have the information they need to determine the adequacy of a firm's capital.
- Firms should be required to report all the major risk factors that might impede the quality of capital. Some firms already do so, but it is doubtful that they capture the whole truth. In our experience firms will often choose to ignore or to understate the most material risks.
- Auditors should be encouraged to determine whether any material risk to capital has been omitted from the financial statements. This task should not be difficult. Given that it is easy to identify books with excess or unusual returns, the auditor should by default declare such books high risk, and check that the risk has been communicated clearly in the statements.
- Risks hidden in the maturity structure should not be concealed in the financial statements. Many firms report the structure of debt up to 5 years, with an aggregate of debt longer than 5 years. Yet the sensitivity to debt is (roughly) proportionate to the term! Our research has uncovered some staggering risk sensitivities concealed in this way.
- Material risks should not be concealed in some 'other' category. Carillion's accounts represented an early payment facility as 'other creditors', meaning it was not incorporated in a debt to earnings ratio presented to lenders.[137] Eumaeus has reviewed another firm which reports 'other valuation differences' of nearly £1bn, more than half of its reported capital![138]

As for the industry, good luck, but it is essential to get the valuations right.

[136] https://www.gov.uk/government/publications/the-quality-and-effectiveness-of-audit-independent-review
[137] 'Carillion', Business, Energy and Industrial Strategy and Work and Pensions Committees 2018, p. 43
[138] See e.g. http://eumaeus.org/wordp/index.php/2018/10/19/past-and-present-tens.