Using RBOs and Megafunds to Hedge Longevity Risk and Specialty Drug Costs

Roger M. Stein

Roger M. Stein
teaches finance and machine learning at New York University Stern School of Business in New York, NY.
rstein@stern.nyu.edu

KEY FINDINGS

- Research-backed obligations (securitizations backed by portfolios of biomedical or other research, also called RBOs, bio bonds, translational-bionotes, etc.) can be useful for hedging certain types of health- and life-related costs, which we demonstrate in simulation experiments.
- Such hedges may be best characterized as natural diversifying trades. While not perfect, they nonetheless confer superior performance to the hedger under most states of the world, regardless of portfolio outcomes.
- RBO-based hedging strategies align the interests of patients and benefit providers in accelerating the development of new, life-extending therapies for society's most challenging diseases.

ABSTRACT

Longevity risk, the risk that a member of a pension plan or an annuity holder lives much longer than anticipated, can cause shortfalls in funding for writers of insurance policies, annuities, and pension benefits. Reimbursement risk, the risk faced by health benefit providers in paying for expensive new drugs, can similarly cause shortfalls for health insurance firms and other benefit providers, particularly as the costs of specialty drugs increase. These risks can be acute when scientific breakthroughs increase new drug development rates. An emerging asset class, research-backed obligations (RBOs) (Fernandez et al. 2012; Hull et al. 2019), has the potential to provide a natural hedge for a broad class of such risks. This article demonstrates how to calculate simplified, first-order portfolio-specific hedge ratios. The author presents simulation results that demonstrate the benefits of such hedges for a hypothetical health insurance portfolio facing rapid increases in reimbursement risk due to the high cost of new specialty drugs.

Recent advances in science and medicine have rapidly changed the prospects for those afflicted with a range of diseases—many of them life threatening. For example, until about 2011, treatments for hepatitis C were able to clear only about 40% of all patients with the virus, while today the cure rate exceeds 90% (FDA 2014). For patients with certain (but not all) forms of breast cancer, the introduction of the drug Herceptin in a chemotherapy and endocrine therapy regimen appears to significantly increase both the survival rate and the longer-term prognosis, with over 97% of patients surviving five years later and over 80% showing no signs of recurrence eight years after treatment (O'Sullivan et al. 2015).

Society as a whole benefits from, and indeed rejoices in, the development of therapies that extend the lives of those afflicted with cancer and other diseases. However, in some financial contexts, this exuberance is tempered by an analysis of the financial impact of extended life expectancy on certain classes of financial obligations. In this article, we discuss these financial instruments and the implications of increased longevity on them. We also suggest a novel approach to mitigating them and provide some simplified first-order measures.

Although the results we show here are simplified, in that they ignore higher-order (Jensen) effects, they are nonetheless useful in providing intuition for the impact of including RBO equity in portfolios of health-sensitive liabilities. Practical applications would incorporate more detail about the liability portfolios and adjustments for these effects.

Pension Payers Face Increasing Longevity Risk

One such class of longevity-sensitive financial obligations is pension obligations that require the obligor to make regular pension payments to retired employees in defined-benefit pension funds. Defined-benefit pension plans commit to provide a fixed monthly payment to pensioners from the time they retire until the ends of their lives. To the extent innovations in science and medicine materially extend the lifespans of these pensioners, all else equal, the funds risk a mismatch between their assets and liabilities, since the number of years over which the funds are committed to pay pensions increases, on average.

Example 1. The Cost of a Life-Extending Drug for a Pension Fund

To give a sense of the order of magnitude of the financial impact of even moderate increases in longevity, consider the following simplified hypothetical case:

For the reporting year ending March 31, 2019, the New York Employees Retirement System paid pension benefits to 481,795 retired state employees, with the average payment being $24,345 for regular employees and $52,804 for retired police and firefighters (New York State Office of the State Comptroller 2020).

Imagine that the plan begins the year fully funded, but at the beginning of the year a new drug is introduced that can extend the life expectancy of a patient with a certain chronic disease by one year.

Imagine further that just one-half of one percent (0.5%) of retired NY State employees were afflicted with this disease. By a coarse calculation,[1] the pension fund would now face an unplanned shortfall of about $137 million.

For reference, note that about 72% of the pension recipients in the New York State plan were reported to be 65 years old or older (New York State Office of the State Comptroller 2020), and it has been estimated that about 92% of US residents over the age of 65 have at least one chronic disease; that more than 40% have three or more; and that more than 25% suffer from some form of disability in daily life (Hung et al. 2011). ■

Health Benefits Providers Face Increasing Reimbursement Risk

For related reasons, health insurance firms and other health benefit providers are being called upon more frequently to reimburse patients for expensive new specialized treatments that are being developed at an accelerated rate.

[1] For expository purposes, we ignore the exact distribtuion of ages and salaries.

While the societal impact of treating these diseases is overwhelmingly positive, the financial impact of these new and often very expensive drugs can be material for organizations that must provide patient reimbursement. This has become a growing concern for such organizations: a recent industry report noted that "[T]he major contributors to rising...spend for specialty medications are brand inflation and the accelerating development of expensive, highly targeted therapies" (Express Scripts Lab 2015).

Example 2. The Cost of a Life-Extending Drug for a Pension Fund

To give perspective on the order of magnitude of the financial impact of the introduction of a new drug on health benefit providers, consider again the case of the drug discussed in Example 1.

For convenience, we will again consider the case of Americans 65 years old or older who have private insurance.[2] The US Census estimates that in 2019 approximately 58% of Americans over the age of 65 were covered by some form of private health insurance (US Census Bureau 2020a). Based on the estimated population (US Census Bureau 2020b), this implies approximately 17.4 million individuals with private insurance who are 65 or older.

Imagine that, as in Example 1, just one-half of one percent (0.5%) of these individuals have a specific fatal disease that can be treated with a new drug that will extend their life for one year. Assume that the new drug costs $50,000 per year (which is lower than the typical specialty drug price). By a coarse calculation,[3] and assuming that only 50% of the patients with the disease would be eligible for the drug, the aggregate cost to the private insurance sector of the new drug would be about $2.2 billion in the first year.

For reference, note that the NAIC estimates that the sector-wide aggregate profit for the health insurance industry was approximately $23B (NAIC 2019), which would imply that the costs under the analysis above represent just under a 10% reduction in total profit for the industry. ■

Research-Backed Obligations (RBOs) as Hedges

Clearly, the uncertainty around the costs of such new scientific advances is of concern both to institutions that must reimburse patients for new life-extending therapies and to institutions that must continue to pay retirement and other lifetime benefits to individuals whose lifespans are extended through the advent of these new drugs.

In what follows, we discuss one approach to partially hedging these types of risks using a newly introduced class of security called a research-backed obligation (RBO, bio bond, translational bio note), which involves the securitization of early-stage biomedical research projects (Fernandez et al. 2012; Hull et al. 2019).

While previous articles in the literature on RBOs focused on the supply side of the market for RBO securities—that is, the viability of structuring such securities and their likely payoff behavior—our approach was the first to explore the demand side of the market: the specific needs of investors that RBO securities may be used to address, and how investments in RBOs may be sized to meet these specific needs (see Stein 2016).

[2] We do this for expediency in order to make use of the setup of Example 1. In later examples, we will focus more directly on specialty drugs, the majority of which are developed to treat rare diseases.

[3] For expository purposes, we ignore the exact distribtion of ages, mortality rates, as well as coinsurance, discounts, and so on.

The structure of RBO securities provides a natural cashflow profile that enables diversifying trades for longevity and reimbursement risks. Such a trade resembles a form of (imperfect) hedge for longevity and/or reimbursement exposures. The diversifying trade gains value when new, life-extending drugs are developed, which can offset the costs to the payer of such medical advances. We use simple analytic models to demonstrate more rigorously that this new asset class (RBOs) can potentially be used to offset many forms of longevity risk that arise as the result of scientific breakthroughs. We provide an example demonstrating how RBO equity could be used as a longevity hedge by an underwriter of annuities, which represent a particularly simple class of lifespan-dependent liabilities.

We acknowledge that many of the drivers of longevity risk are not directly related to the development of specific drugs for specific diseases, but are rather the results of changes in general health measures (e.g., changes in diet, exercise, better access to health services, etc.). In recognition of this, we next extend our approach to accommodate hedges for reimbursement risk of the sort faced by health insurance firms that must repay patients for the cost of newly developed drugs. In contrast to the longevity case, these risks are quite directly tied to the introduction of new therapies. Our most widely applicable results are thus likely to be those relating to reimbursement risk.

While our analytic results provide intuition for the risk mitigation approach, we stress the stylized nature of our closed-form analytic models. In particular, we note that for simplicity, and to develop intuition, we treat many of the quantities in our derivations as deterministic, while they might more realistically be taken as stochastic.

We address this, albeit only partially, through simulation experiments. For example, our simulation results suggest a correlation of 0.66 between the return to investors on RBO equity and the reimbursement shortfall experienced by a health insurer. Even under extremely conservative assumptions, this correlation estimate is still 0.34, implying that RBO equity offers substantial risk mitigation benefits. In about 87% of all simulation paths, the diversifying position in RBO equity reduces the RBO shortfall, even after accounting for the cost of the trade construction costs.

Although we make a number of stylized assumptions in our modeling, many of these may be made more realistic through reference to application-specific data sets. However, because practical application of our results does require simulation (closed form solutions do not appear to exist for the joint portfolio dynamics), many of these extensions may be incorporated into the simulation framework as needed. Our main results are from simulation experiments.

This article is organized as follows: The next section offers background on longevity risk and RBOs, reviews the mathematics of annuities, and provides a stylized representation of an RBO portfolio. Then we use these constructs to derive "hedge ratios" based on the characteristics of an annuity underwriter's portfolio. We next extend this approach to reimbursement risk; then we use a more realistic simulation of an RBO portfolio and cashflows to provide a sense of how reimbursement "hedge ratios" might be estimated in practice. In the next section, we discuss our results and present a number of important limitations in our current model and suggest avenues for future work; then we conclude in the final section.

ANNUITIES, LONGEVITY RISK, AND RESEARCH-BACKED OBLIGATIONS

Background

Longevity risk. This work is related to a number of research streams within the longevity risk literature. Of particular interest are results that discuss two topics:

(a) pricing and mitigating of the longevity risk of individual instruments; and (b) valuing and constructing portfolios or securitizations of longevity-risk-sensitive instruments. While it is beyond the scope of this article to review these literature streams in detail, we do discuss several relevant examples to provide context. More-detailed treatments can be found in original publications and references therein.

The advent of a secondary market (Smith and Washington 2006), and later a securitization market for logevity-risk-sensitive instruments led to increased interest in pricing and risk mitigation for such exposures (e.g., Perera and Reeves 2006), as well as the introduction of new derivative instruments for hedging (e.g., Mott 2007).

Many of these results focused on the risks of longevity extension associated with general, non-specific, increases in population longevity. Stone and Zissu (2006) introduce measures of duration and convexity as analogs to the fixed income equivalents, using time as a proxy for expected increases in population longevity. They subsequently extended this approach to accommodate computation of percentage changes (elasticities) of pools of life settlements (Stone and Zissu 2008).

Later authors considered various forms of portfolio stratification as a means to better measure the impact of longevity risk. Brockett et al. (2013) demonstrated how information available to investors can be used to adjust standard mortality tables to value life settlements and also presented a number of deterministic and stochastic models for valuing these products, in the context of a single contract. Parnes (2018) explored a different dimension by developing a multistate model for decomposing mortality risk among subpopulations of healthy and of already terminally ill individuals, and demonstrated how this may be used to provide a more nuanced view of the aggregate portfolio longevity risk.

In many cases these results are useful to practitioners and academic researchers. However, the bulk of this work focused on changes in value related to population increases in health and longevity, rather than on specific subpopulations that may experience substantial increases in longevity due to medical and scientific discoveries. We discuss the implications of this in the next section.

RBOs. Research-backed obligations (RBOs) can provide one potential means to address risks associated with medical innovations that extend lifespans. The original motivation for RBOs was to address funding shortages in the early stages of drug research. But RBOs also offer natural longevity payoffs (as opposed to synthetic payoffs, such as those related to mortality indexes) that can offset the increases in liabilities from life-extending innovations in medicine.[4]

An RBO pools ownership in many individually risky, but largely uncorrelated, drug development projects and thereby reduces the overall risk of the portfolio (see Fernandez et al. 2012). Early articles in this literature showed that portfolios containing early-stage cancer therapy candidates could be structured in such a way as to support debt issuance and yield returns consistent with those targeted by large institutional investors. Subsequent extensions of this approach (Fagnan et al. 2014) considered its application to candidate therapies for orphan diseases. Orphan diseases enjoy more favorable regulatory treatment and shorter approval times; they are also, in many cases, monogenic diseases making targeting more precise and more likely to succeed. The authors showed these portfolios can be much smaller than those originally described for cancer therapies (by virtue of the much higher success probabilities of candidate therapies for genetic diseases), and their profit profiles correspondingly more favorable.

[4] Certain classes of life insurance policies also may be structured to provide natural hedges under specific conditions (see Milevsky and Promislow 2001). In these cases, firms may seek to hedge the *residual* longevity risk on their portfolios.

Annuities. While our main results relate to reimbursement risk, it is instructive to begin with the case of longevity risk, from which much of our machinery required for analyzing reimbursement risk derives. In addition, in cases in which an annuity underwriter is concerned about mitigating the longevity-related impact of new drug developments, our longevity analysis provides a practical mechanism for analyzing how to use RBOs to mitigate the risk.

We begin by reviewing notation and mathematical formulations and also outline the key results from previous work on RBO structures. We focus on the analysis of annuities, because of their natural relationship to longevity-related obligations.[5] However, defined-benefit pension plans effectively underwrite functionally equivalent contracts in the form of promises to pay fixed pension benefits.

Annuities and the Impact of Life-extending Therapies on Annuity Cost

It is well known that an annuity that pays a fixed periodic cashflow, C, over T periods may be valued as

$$A(C,r,T) \equiv A = \frac{C}{r}\left(1 - \frac{1}{(1+r)^T}\right) \tag{1}$$

where

C = the one period constant cashflow,
r = the discount rate, and
T = the number of periods over which C will be paid.

It will also be convenient to distinguish between different values of T, the maturity of a specific annuity, in different contexts.

$\Delta_j \equiv$ years lost from a diseased individual's lifespan due to disease j,
$D_j \equiv$ set of individuals with disease j, and
$\Pr(i \in D_j) \equiv$ the probability that individual i has disease j.

If we assume that all diseases have the same Δ_j, then depending on the context, we can define the *remaining lifespan* (r.l.) for an individual in any of the following ways:

$T_i \equiv$ the specific r.l. of individual i;
$T_h \equiv$ the r.l. of a healthy individual;
$T_d \equiv$ the r.l. of a diseased individual
$= T_h - \Delta_j$; and
$T_u = \Pr(i \in D_j)T_d + (1 - \Pr(i \in D_j))\ T_h$
$\equiv$ the unconditional r.l. of an individual of unknown disease status

where $\Pr(i \in D_j)$ is the probability[6] that individual i has disease j.

If a therapy for disease j is discovered at time $\tau_j < T_d$, and this therapy extends life by Δ_j periods for a diseased individual, then we can derive the exact form of the change in the value of the annuity as a result of the new therapy, which turns out to be particularly compact due to the structure of Equations 2 and 3.

Specifically, for an individual with lifespan T_i and an intervention that extends life by Δ years,

[5] Annuities remain a popular retirement-planning financial product. According to one estimate, just under $242 billion of new annuities were purchased in the United States in 2019, bringing the total outstanding deferred annuity assets to just under $3.15 trillion (Insurance Information Institute 2019).

[6] Note that for convenience, and without loss of generality, we assume that this reflects the probability that an individual has the disease *and would respond to a new drug in development*.

$$\Delta A_i = A(C, r, T_i + \Delta) - A(C, r, T_i) \tag{2}$$

$$= \frac{C}{r}\left(1 - \frac{1}{(1+r)^{T_i+\Delta}}\right) - \frac{C}{r}\left(1 - \frac{1}{(1+r)^{T_i}}\right) \tag{3}$$

$$= \frac{C}{r}\left(\frac{(1+r)^{\Delta} - 1}{(1+r)^{T_i+\Delta}}\right) \tag{4}$$

If an annuity is written to individual i who has disease j and who has a lifespan of $T_i = T_d$, but a new therapy is introduced that will extend life for this individual by Δ_j years, the expected change in cost to the underwriter is Equation 4, evaluated at $T_i = T_d$ and $\Delta = \Delta_j$:

$$\Delta A_d = \frac{C}{r}\left(\frac{(1+r)^{\Delta_j} - 1}{(1+r)^{T_d+\Delta_j}}\right). \tag{5}$$

Note that we have introduced the notation A_d to denote the value of an annuity written to a diseased individual. We will correspondingly denote by A_h and A_u the values of annuities written to healthy individuals and individuals of unknown disease status, respectively.

The more natural case is one in which we do not know i's disease status so that $E[T_i] = T_u$. If a therapy for an arbitrary disease j^* is discovered at a future time, the expected change in T_i as a result of the new drug discovery is given as

$$\mathrm{E}[\Delta] = \Pr(\tau_{j^*} < T_d)\Pr(i \in D_{j^*})\Delta_{j^*} \tag{6}$$

In words, the expected change in value is the change in value of the annuity for a diseased individual who receives the treatment, weighted by the compound probability that individual i has the disease and that the therapy is introduced during the lifespan of the diseased individual:[7]

$$E[\Delta A_u] = p_{\tau_j} p_{D_j} \Delta A_d + (1 - p_{\tau_j} p_{D_j})\Delta A_h \tag{7}$$

$$= p_{\tau_j} p_{D_j} \frac{C}{r}\left(\frac{(1+r)^{\Delta_d} - 1}{(1+r)^{T_d+\Delta}}\right), \tag{8}$$

where $p_{\tau j} = \Pr(\tau_{j^*} < T_d)$ and $p_{Dj} = \Pr(i \in D_{j^*})$.

Equation 8 gives the expected change in annuity value for an annuity holder i of unknown disease status subject to the introduction of a new therapy for disease j^*.

For completeness, note that the expected cost of the annuity to the annuity writer is given as

$$E[A_u] = p_{\tau_j} p_{D_j} A_d + (1 - p_{\tau_j} p_{D_j})A_h \tag{9}$$

[7] We emphasize that, for simplicity, we are treating most quantities as deterministic, rather than stochastic or estimated, respectively. This ignores other sources of uncertainty in several of the fundamental quantities, including uncertainty (and related effects) about the realized value of Δ_j^*; uncertainty about the change in the specific individual's lifespan, T_i, due to the advent of the therapy; and uncertainty associated with estimates, T_j, of the lifespans of diseased and healthy individuals themselves, respectively. The simple first-order derivations we present in this article do not account for this. A more involved treatment would accommodate these (e.g., via the delta method). We later incorporate only the first of these additional sources of uncertainty in our simulations.

While an underwriter can adjust the price of an annuity based on actuarial tables that account for the historically small, but non-zero, probabilities of new therapies that extend the lives of their clients, a larger concern is the possibility that many new therapies become viable at the same time—either by chance, or due to a change in the scientific process or knowledge base.

For example, if T_h is the average lifespan of a healthy individual, known to have no diseases, and T_j is the average lifespan of an individual known to have disease j, then (ignoring individuals with multiple diseases and mangling slightly the notation) a first-order estimate of $T_i = T_u$ for an individual of unknown disease status (assuming $T_j < T_u$, i.e., that disease shortens lifespan) is

$$T_u = T_h - \sum_j (1 - \Pr(\tau_j < T_i))\Pr(i \in D_j)\Delta_j \tag{10}$$

The risk to an underwriter, which we call *scientific longevity risk* or longevity risk for short, is that Pr ($\tau_j < T_d$). Said differently, as the probability of new drugs being developed increases, individuals are more likely to enjoy the lifespan of a healthy individual (T_h) rather than suffering from the shortened lifespan of a diseased individual (T_d). For large portfolios, even modest shifts in lifespan can have material financial consequences, while more substantial shifts can be catastrophic financially.

(We have provided a stylized example in Appendix D of the supplemental information, online, that uses the above representations to calculate the impact of new therapeutic innovations on the value of an annuity and on portfolios of annuities.)

For completeness, we now formalize the expected change in the value of the portfolio of annuities, $\mathrm{E}[\partial P_A / \partial k]$, for a single new drug approval. If we denote the number of drugs approved for the life of the cohort as k, and, for convenience, assume that for all diseases, j, $\Pr(i \in D_j) = p_D$ and $\Delta_j = \Delta_D$ (i.e., both quantities are constant across diseases), then for a portfolio of N annuities

$$\begin{aligned}
\frac{\partial A_d}{\partial k} &= \frac{C}{r}\left(\frac{(1+r)^{\Delta_D} - 1}{(1+r)^{T_d + \Delta_D}}\right) \\
\frac{\partial A_h}{\partial k} &= 0 \\
\mathrm{E}\left[\frac{\partial A_P}{\partial k}\right] &= \sum_{i=1}^{N}\left[p_D \frac{\partial A_d}{\partial k} + (1 - p_D)\frac{\partial A_h}{\partial k}\right] \\
&= (p_D \times N) \times \frac{C}{r}\left(\frac{(1+r)^{\Delta_D} - 1}{(1+r)^{T_d + \Delta_d}}\right)
\end{aligned} \tag{11}$$

Equation 11 provides an intuitive representation of the impact of a new drug on the portfolio of annuities. The expected change in the value of the portfolio of annuities is proportional to the expected number of diseased individuals in the portfolio whose lives are extended as a result of the new therapy.[8]

Payoffs for Portfolios of Translational Therapies

As the RBO market begins to develop for institutional investors, it may present a unique opportunity for firms facing longevity risk.

Since the equity tranches of RBO transactions gain value as more therapies are developed, they behave *somewhat* like a hedge against classes of longevity risk that

[8] See Footnote 8 for caveats.

are associated with new life-extending therapies. RBOs are unique in this regard: their portfolios provide natural risk-mitigation for longevity risk related to the introduction of novel therapies, in contrast to longevity swaps, for example, which are based on actuarial indexes rather than real assets.

In this section, we consider a simple analytic model for the returns on a portfolio of translational drug projects.

To build intuition, the model described in this section reflects a stylized, but informative, view of the portfolio behavior of a basket of drug development projects. We base our model on the motivational model introduced in and extended in Fernandez et al. (2012) and Hull et al. (2019).

Assume that each successful therapy has a total NPV of $\$V$. If there are k successful therapies in the portfolio, then the total value of the portfolio, V_P is given simply as

$$V_P = kV, \tag{12}$$

and the change in the portfolio value for an additional successful drug is

$$\frac{dV_P}{dk} = V. \tag{13}$$

If we assume further that each candidate drug in the portfolio has a success probability of p_S and a failure probability of $1 - p_S$, and that the number of successes across the portfolio, k, follows a binomial distribution:

$$p(k < k^*) = \sum_{j=0}^{k^*-1} \binom{n}{j} p_s^j (1-p_s)^{n-j}, \tag{14}$$

then the mean value of the portfolio is simply

$$\mathrm{E}[V_P] = n \times p_S \times V \tag{15}$$

and an investor wishing to receive $V_I \leq E[V_P]$ in expectation could simply purchase a V_I / V_P share of the portfolio, assuming that the payoffs on success of the drug trial process were uncorrelated with systematic market factors, and assuming the fund could be made arbitrarily large (to reduce the idiosyncratic risk of the individual projects).

In the absence of correlation[9] and path dependence (we introduce these real-world features in our simulations), the probability of the portfolio generating $\$F$ or more cashflow ($F \leq nV$) is simply 1 minus the probability in Equation 14 with $k^* = \lceil F/V \rceil$, where $\lceil x \rceil$, indicates the largest integer greater or equal to x. The quantity $\lceil F/V \rceil$ is the minimum number of projects that must succeed in order to be confident at a $1 - p\left(k < \lceil F/V \rceil\right)$ probability level that there will be at least $\$F$ of cashflow generated.

[9] In the case of a single common factor $F \sim N(0, 1)$ across candidate drugs, for example, we may modify p_s in Equation 14. When $F = F_k$,

$$p_s = N\left(\frac{N^{-1}(p) - \sqrt{\rho}F_k}{\sqrt{1-\rho}}\right),$$

where F_k is an arbitrary value of F of interest and ρ is the pairwise correlation between two candidate therapies in the portfolio. To compute the full probability, we integrate over F. See Hull et al. (2019) for details.

This observation forms the basis of debt issuance by the fund: To the extent the fund wishes to use leverage, debt issued with a face value of $\$F$ will have a probability of default of $PD = p(k < \lceil F/V \rceil)$. Such debt would need to offer at least a market rate of interest but also serves to make the equity investment both riskier and more profitable, presumably offering a better return than the debt.

Thus, the optimization of the capital structure of the levered fund becomes a decision problem analogous to that of determining the optimal capital structure of a firm (or any structured financing), which is familiar to students of financial economics. More-detailed discussions of this stylized model, and the impact of leverage on the risk and return of equity, can be found in Fernandez et al. (2012) and Hull et al. (2019).

HEDGING SOME FORMS OF LONGEVITY RISK

Longevity risk can be seen as a systematic exposure for financial institutions that underwrite lifetime-payment contracts such as pensions or annuities (Aro 2014). A primary source of longevity risk has historically been largely statistical: Existing models for longevity forecasts (Lee and Carter 1992) have tended to consistently underestimate future lifespans (IMF 2012, Ch. 4). By one estimate, for example, the impact of a one-year shock to longevity would result in a corresponding increase in US pension liabilities of 2%–3% or approximately $84 billion (Kisser et al. 2012).

However, another source of longevity risk is improvements in population health outcomes due to better diet and nutrition, increased exercise, early detection of some treatable conditions, and so forth. In general, these factors are more difficult to hedge. In recognition of the financial impact of longevity risk, and in order to seek capital relief under some regulatory regimes (e.g., Solvency II), some firms have recently begun structuring insurance contracts to protect against longevity-related increases in liabilities. In one transaction from this emerging sector, BT Group launched its own insurance firm and, in collaboration with Prudential, created longevity insurance for about £16 billion of its approximately £47 billion in pension liabilities (Aglionby and Cumbo 2014). Prudential also underwrote a $33 billion annuity to GM to effectively buy out the pension obligations for 118,000 GM workers (Olsen 2012).[10]

Importantly in the context of this article, a subset of health improvement-related increases in lifespan is tied to medical advances that extend life through the introduction of new or better drugs—a smaller component of longevity risk but a substantive source of extreme changes in longevity. Thus, it represents a form of longevity tail risk: even modest changes to mortality rates can have material impacts on the lifespans of the older individuals most likely to be the source of longevity risk. It has been historically difficult to hedge this form of medical longevity risk. However, the introduction of RBOs as an asset class provides new opportunities.

RBO Equity as a Natural "Hedge" for Scientific Longevity Risk

Assume an annuity writer maintains a portfolio in which a certain percentage of annuity-holders are likely to die of a specific set of diseases, and therapies for some of those diseases are being developed within an RBO portfolio.

[10] For a more detailed discussion of the structure of various longevity hedging instruments, see BIS (2013).

In principle, an annuity underwriter would want to purchase an amount of RBO equity, E_{RBO}, such that the additional cost of longer lifespans on the part of annuity holders, $E[\Delta A_i]$ is offset by the expected rise in the value of the RBO equity investment. Said differently, when the success rate on the portfolio is higher than the expected rate, the annuity writer would like the portion of the profit from those successes to offset the increased costs of paying the annuities in the portfolio that result from the longer life expectancy of the annuity holders who receive those successful therapies.

Thus, if $E[V_P] = \Sigma_{c=1}^{n} p_s V$ is the expected payout for the portfolio, then $\Delta V_P = V_P - E[V_P]$ is the unexpected loss or gain on the RBO portfolio. The underwriter would like to maintain a hedge such that

$$H^* \times \Delta V_P = \Delta A_P \tag{16}$$

so (to a first approximation)[11]

$$H^* = \frac{\Delta A_P}{\Delta V_P}. \tag{17}$$

H^* is simply the hedge ratio for the annuity portfolio against the RBO equity. Of course, ex ante the values of ΔA_P and ΔV_P are not known. However, we have already estimated the change in A_P and V_P with respect to successful drug approvals; these are (11) and (13). It is sufficient for these changes to be hedged. We thus have

$$H = \frac{\partial A_P / \partial k}{dV_P / dk} \tag{18}$$

or

$$H = \frac{(p_D \times N) \times \frac{C}{r}\left(\frac{(1+r)^{\Delta_D} - 1}{(1+r)^{T_d + \Delta_d}}\right)}{V} \tag{19}$$

Thus, an underwriter seeking to hedge longevity risk due to the introduction of new therapies, to a first approximation, would purchase $\$H$ of equity in an RBO for every dollar of the annuity portfolio, assuming an annuity portfolio of size N and annuity value A_u:

$$E_{RBO} = N \times A_u \times \frac{(p_D \times N) \times \frac{C}{r}\left(\frac{(1+r)^{\Delta_D} - 1}{(1+r)^{T_d + \Delta_D}}\right)}{V} \tag{20}$$

Example 3. Using RBO Equity to Hedge Longevity Risk in a Cohort of Underwritten Annuities

Consider the highly stylized example of an insurance firm that sells annuities to a cohort of 100,000 40-year-old individuals, with each annuity paying $10,000 per year for the remainder of the client's life. We assume that healthy individuals live to age 85 and that a diseased individual's life-span, T_d, is 60 years, but that the advent

[11] See caveats in Footnote 8.

of a therapy for the disease would extend life by $\Delta_d = 25$ years. We assume that the risk-free rate r is 5%.

We also assume a total of 100 diseases of interest and that $\Pr(i \in D_j) = 1/100$, $j = 1 \ldots 100$ so that there are, on average, 1,000 individuals with various diseases in the cohort; and that diseases are uniformly distributed across the diseased subpopulation, so that on average, 10 individuals will have each disease.

We assume that an RBO is available, supported by a portfolio of 40 candidate drugs that each has a 5% probability of success. Successful drugs are sold out of the portfolio for a value of $V = \$250$ million, so the expected value of the portfolio is $500 million. If the target total return on the fund were 25%, and the fund issued no debt, this would imply (ignoring discounting) $400 million of equity available at the launch.[12]

If the underwriter wished to hedge the longevity risk of the annuities, the implied first-order hedge ratio is

$$
\begin{aligned}
K &= 400,000,000 \equiv \text{Initial drug portfolio capital} \\
H &= \frac{(p_D \times N) \times \frac{C}{r}\left(\frac{(1+r)^{\Delta_d}-1}{(1+r)^{T_d+\Delta_d}}\right)}{V} \quad \text{(Eq. 19)} \\
&= \frac{0.0001 \times 100,000 \times \frac{10,000}{1.05}\left(\frac{(1.05)^{25}-1}{(1.05)^{20+25}}\right)}{250,000,000} \\
&\approx 0.002124744 \\
E_{RBO} &= K \times H \\
&\approx 400,000,000 \times 0.002124744 \\
&= 849,898.
\end{aligned}
$$

It is useful to observe that the underwriter might not wish to pay for a hedge against the introduction of fewer new drugs than it has already budgeted for in its actuarial price assumptions; if the underwriter already priced in the cost of the mathematical expectation of five new therapies, it would prefer to hedge only the new drug introductions for drugs after the fifth new drug. Using straight equity there is no natural mechanism for doing this. However, the financial engineering techniques that characterize the fuller RBO model (tranching, cashflow triggers, etc.) make such hedges feasible, at least approximately. ■

There are a number of reasons to believe that a hedge of the sort described in the previous example would be imperfect, due to the mathematical properties of the hedge ratio and its components on the one hand, and the structure of the drug development and discovery market on the other. This latter effect may be exacerbated by realistic limitations on the size of an RBO portfolio.

While a more detailed analysis would account for the mathematical nuances (see Supplemental Information, online, for more details), the structural ones may make perfect hedging practically impossible at this time. This does not mean that practically useful hedges cannot be constructed using RBO equity, but rather that they may remain imperfect for the near term. We discuss potential limitations and remedies below.

[12] For purposes of this simple example, we assume that the fund can acquire a portfolio of compounds and pay for required clinical trials and expenses out of this initial investment. In later simulations, we accommodate much more detailed and realistic assumptions, as well as debt issuance by the RBO.

EXTENDING THE FRAMEWORK TO "HEDGE" REIMBURSEMENT RISK FOR HEALTH INSURERS

We now extend the framework described above to consider hedge construction for reimbursement risk. Such hedges would be attractive to institutions facing uncertain future payouts due to new drug discoveries. Unlike in the case of longevity risk on an annuity portfolio, one of the biggest drivers of reimbursement risk *is* the introduction of new and expensive therapies for diseased individuals in the insurer's portfolio. Furthermore, assuming a fixed prevalence rate for a given set of diseases, the financial impact of these new drugs is closely tied to the number of new drugs introduced and their costs, both of which are proxied for by the RBO portfolio.

Reimbursement Risk

Reimbursement risk is a growing problem for providers of health benefits, such as health insurance firms and reinsurance companies (Fernandes 2015). Although insurance providers can accommodate modest shifts in cost structure through annual changes in member premiums, a more pressing concern is the potential for very large changes in expense base associated with the introduction of many specialty drugs in a short period of time.

Consider the following recent developments:

- In 2012, four drugs were approved with annual per-patient costs of over $200,000 (Herper 2013).
- A study of 47 oral oncologics found that patient spend on this subsegment increased by 37% between 2006 and 2011 (Conti et al. 2014), and in 2019 alone, patient spend on oncologics rose almost 13%, marking the fifth year in a row in which prices rose by double digits (Nass and Simorellis 2019).
- Specialty medications represented almost half of the total 2019 drug spend by patients, though only about 2% of patients used specialty drugs (Express Scripts Lab 2020; Tichy et al. 2020).

Against this backdrop, the pace of new specialty drug development also appears to be increasing. In a survey of its members, PhARMA, an industry association of biotech and biopharmaceutical research firms, reported that over 70% of the drugs in the current development pipeline are so-called "first-in-class" medications—those that use different mechanisms than any currently approved drugs (PhARMA 2015). This suggests that first-in-class drugs enjoy little or no initial competition and thus are likely to be priced aggressively.[13]

As the drug development landscape shifts, such shocks are also becoming increasingly likely. The 2014 approval of the Hepatitis C drug Harvoni® (ledipasvir/sofosbuvir) by the US Food and Drug Administration was a notable example. A single course of treatment with Harvoni spans 8–24 weeks, and fully resolves (cures) Hepatitis C in the majority of patients. Despite Hepatitis C affecting about 3 million

[13] There has also been an increased focus on so-called "orphan diseases"—those that affect fewer than 200,000 Americans. (Similar designations have also been instituted in other jusridictions in Europe and Asia.) Many of these diseases are genetic in nature and research on their therapies may be expected to benefit from recent advances in genetic sequencing. For example, although it has taken 25 years to sequence about half of all genetic diseases, Boycott et al. (2013) assert that the rapid increase in the pace of genetic mapping is the result of dramatic improvements in DNA-sequencing technologies. The impact for patients with these diseases is likely to be faster development of more-targeted and effective treatments for the over 7,000 orphan diseases that affect about 1 in 10 individuals (PhARMA 2013).

patients, the high pricing of this drug has been more similar to orphan disease therapies than to drugs targeting broader populations. The drug is priced at over $1,000 per tablet, implying a monthly cost of over $30,000 (Express Scripts Lab 2015).

Such sudden increases in costs are not easily transferred to members through single-year premium increases, which would require drastic reimbursement cost shifts well beyond the typical inflation-based year-on-year premium increases tolerated by insurance buyers. So the goal of many reimbursement paying agents is less to entirely eliminate the costs associated with new drug introduction (which must be managed over multiple years), but rather to smooth out large spikes in costs and incorporate them into premium levels in an orderly fashion.

Modeling Reimbursement Risk

Reimbursement risk is the risk that the reimbursement agent will be required to reimburse patient subpopulations for treatment costs associated with one or more newly discovered therapies for a set period of time or for the patient's lifespan.

Assume the cost of a single therapy to the health benefit provider is C_E per year while the drug enjoys exclusivity, and C_O thereafter. Thus, if the drug enjoys exclusivity for Y_E years, and the patient is reimbursed for T years, then if the new drug is approved today, the total cost is

$$C = \sum_{t=1}^{Y_E} NPV(C_E, t) + \sum_{t=Y_E+1}^{T} NPV(C_O, t). \tag{21}$$

We may value these cashflows as two annuities. The present value of an annuity maturing at time T that begins payments at time $\tau > 1$ may be seen as the value of an annuity that begins making payments at time 1 and maturing at time T less the value of a second annuity that begins making payments at time 1, but with maturity $\tau - 1$.

The present value of this pair of annuities is

$$A(C, r, \tau, T) \equiv A = \frac{C}{r} \frac{(1+r)^T - (1+r)^{\tau-1}}{(1+r)^{T+\tau-1}}. \tag{22}$$

Then, assuming the patient will require the therapy for the duration of the insured term, the total per-patient reimbursement cost to the agent for a drug approved at time t_a with Y_E years of exclusivity is

$$\begin{aligned} A_r &= A(C_E, r, t_a, T_E) + A(C_O, r, T_E + 1, T) \\ &= \frac{C_E}{r} \frac{(1+r)^{T_E} - (1+r)^{t_a-1}}{(1+r)^{T_E+t_a-1}} + \frac{C_O}{r} \frac{(1+r)^T - (1+r)^{T_E}}{(1+r)^{T+T_E}}. \end{aligned} \tag{23}$$

where

$C_E \equiv$ the cost of the drug during exclusivity
$C_O \equiv$ the cost of the drug after the expiration of exclusivity
$t_a \equiv$ the time at which the new drug is approved
$Y_E \equiv$ the number of years of exclusivity for the drug
$T_E \equiv$ the time of the expiration of exclusivity
$= t_a + Y_E - 1$ and,
$T \equiv$ the expiration of the reimbursement obligation

Example 4. The Expected Cost to a Health Insurer of the Discovery of a Therapy for a Juvenile Disease

Consider the case of a health insurance firm that insures a family with a newborn infant. Assume that the infant is diagnosed with a rare and terminal condition which, if untreated, would likely be fatal by age five. However, the year after the baby is born, a new drug for this condition is approved with an exclusivity period of seven years. During this period, the per-patient annual cost of the drug is $150,000; then it will drop to $20,000. (Consider this pricing to represent a blend of pricing for biologics and small molecules.) Assume also that the insurance will expire on the child's 18th birthday. For simplicity, we ignore deductibles, copayments, and so on Then we have

$C_E = \$150{,}000$
$C_0 = \$20{,}000$
$t_a = 2$
$Y_E = 7$
$T_0 = 8$ and
$T = 18$

If the interest rate r is 5%, then by Equation 23, the cost of this new drug for each child in the insurers portfolio diagnosed with this disease is

$$A_r = \frac{150{,}000}{0.05}\frac{(1.05)^8 - (1.05)^1}{(1.05)^9} + \frac{20{,}000}{0.05}\frac{(1.05)^{18} - (1+0.05)^8}{(1.05)^{26}} = \$931{,}153$$

More realistically, if the insurance firm calculated the cost of the reimbursement on the day the drug was approved ($t_a = 1$), then the cost per patient with the disease would be

$A_r = \$986{,}020$.

If there were 20 million insured individuals in a firm's portfolio, of which 6% were below the age of five, there would be 1.2 million children under the age of five in the portfolio. If the incidence rate of the disease were $1/10{,}000$, there would be an expected 120 children with the disease, and the present value of the expected total cost to the insurer would be

$$\begin{aligned} \mathrm{E}[A_P] &= 986{,}020 \times 120 \\ &= \$118{,}322{,}434 \text{ total reimbursement cost.} \end{aligned}$$

For comparison, in 2014, Aetna provided coverage to about 23.5 million medical members and realized a net income of approximately $2.2 billion (Aetna 2015). The cumulative cost of the increased reimbursement expense for the single drug in this stylized example would represent over 5% of Aetna's annual net income. ■

RBO equity can provide a useful, and more direct, hedge for such drug reimbursement risk than for longevity risk. For practical applications, given the many stochastic and path-dependent features in drug development and health outcomes, simulation may be used to estimate hedge ratios for reimbursement risk; we describe this approach in the next section.

SIMULATIONS

While the results of the previous sections are illustrative, they cannot be applied directly to determine hedge ratios using real RBO equity. In practice, neither portfolios

of annuity-like instruments nor portfolios of early-stage research compounds behave deterministically: the lifespans of diseased and heathy individuals are stochastic, as are the number of patients with the disease in a specific carrier's portfolio and the cost to reimburse individual patients for new drugs under different plans and formularies. Furthermore, the valuations, costs, and trial durations of compounds in the RBO portfolio exhibit stochastic behavior, as do the success probabilities for each compound, which, along with other quantities, are estimated with error. The uncertainty around the development path for a new drug, and for its monetization, is further complicated by the fact that new drug compounds in RBO portfolio are often sold before approval, using a structured payment agreement that withholds portions of the payment until set development milestones are achieved.

This makes it difficult to estimate the hedge ratios in analytic form. We therefore estimate hedge ratios through simulation, as would be done in practice.

We focus on risk mitigation strategies for reimbursement risk because longevity-exposed annuity portfolios are driven by many factors beyond drug development. This makes analysis more complex from an actuarial perspective, requiring more numerous and involved assumptions. That said, extending our simulation framework to portfolios exposed to longevity risk is conceptually straightforward.

We consider an RBO collateralized by a portfolio of therapies for rare genetic diseases. Genetic diseases often affect very young children, and the drugs that treat them are often very costly. Our simulation considers a portfolio of orphan diseases as described in Fagnan et al. (2014), but with updated simulation parameter assumptions that reflect more recent published research on genetic disease clinical success rates (Hay et al. 2014), as well as discussions with scientific and industry experts. We also extend the simulation structure to accommodate a more realistic representation of the drug development process.

Simulation Design

We calculate hedge ratios numerically through simulation. In addition to simulating the evolution of the RBO, as in previous work (Fagnan et al. 2014), we also simulate the co-evolution of a moderately sized portfolio of healthcare benefit recipients. We use the jointly simulated behavior of these two portfolios to calculate the appropriate hedge ratio for each simulation path. We then use the distribution of these hedge ratios to select an appropriate hedge ratio policy for the portfolio.

The details of the joint simulation algorithms are given in Appendix E of the Supplemental Information (online).

Our simulations are CPU intensive, using multithreaded algorithms. In order to reduce the number of simulations required and the size of our simulated portfolio, we select a somewhat higher than average prevalence rate. While this inflates our results, we note that the size of the portfolio we simulate is less than 1% of that of a large insurance firm. More generally, our assumptions have been made for demonstration purposes only; in practice, in addition to consideration of higher-order effects, a detailed analysis of the specific RBO portfolio compounds and the payer's specific member-portfolio would be required.

Simulation Results and Discussion

Exhibit 1 presents summary statistics for the simulation results.

The exhibit shows the shortfall to the reimbursement agent (S_f) given various hedging strategies using RBO equity. In the lower portion of the table, we present an analysis of the behavior of hedges calculated by using the median, 75th, and 95th percentiles of H_o. For reference, we also present the results on the unhedged portfolio.

EXHIBIT 1

Economics of the Reimbursement Hedge

RBO	
Expected IRR	0.227
RBO bond probability of default	<1bp
Mean number of new drugs	5.1
Insurance portfolio	10,000
Portfolio size	~$14,801,156
Mean shortfall ($\bar{S}_f$)	

Hedging Policy →	No Hedge	Median H_o	75th pctle H_o	95th pctle H_o
Hedge amt per insured (one-time)	$0	$691	$930	$1,508
Risk reduction:*				
Reimbursement shortfall (S_f) ($MM) (negative numbers are profits)				
Mean	14.8	0.3	–4.8	–16.9
95th pctile	30.5	11.2	6.8	0.2
99th pctile	40.3	17.8	12.8	5.9
Pct time upfront hedge cost < S_f		83.5%	71.5%	42.0%
Pct time better off w/hedge (includes cost of hedge)		86.7%	86.7%	86.7%
Pct time hedge eliminates (or better) S_f (includes cost of hedge)		11.9%	25.2%	52.3%

NOTES: This exhibit shows the results of a portfolio simulation. Hedge amounts are calculated per insured individual in the portfolio. S_f is the reimbursement-related shortfall. Risk reduction shows the mean, 95th, and 99th percentile shortfalls after adjusting for the RBO hedge. The lower half of the table shows the percentages of times when the upfront cost of the hedge was lower than the shortfall, the hedger had lower costs as a result of the hedge, and the hedge completely eliminated the shortfall or resulted in additional positive cashflow to the hedger. *Correlation between RBO sale value of individual portfolio asset and cost per-patient assumed to be 60% ($\rho_{VC} = 0.6$). See Exhibit 2.

The first set of results shows how much the hedge would cost the provider. We show this on a per-insured cost basis (see the discussion in the next section). We assume that RBO equity is bought "at par," that is, that one dollar of equity may be purchased for $1. The remainder of the exhibit shows the performance of the hedges in absolute terms and then in terms of the percentage of realizations for which the hedger is better off having used the hedge (after including the cost of the hedge).

Note that the Risk Reduction section shows the raw shortfall, not including the initial cost of the hedge itself. Here negative numbers indicate positive cashflow (negative costs) beyond the reimbursement shortfall. The section below this reports results after including the initial cost of the hedge.[14] A more detailed discussion of the results can be found in Appendix A, in the supplemental information, online.

In general, regardless of the hedging policy, our results suggest that this form of risk mitigation is prudent. In about 87% of cases the hedger is better off for having hedged, as the ROE is generally positive for the RBO equity. Not surprisingly, the more

[14] Note that here we show the results of applying the hedge ratios calculated using the simulation to the simulation paths themselves. In principle, there is a risk of "look-ahead" bias since the sample used to calculate the ratio is also the sample we are using to report the results. However, because the simulation parameters are the same, there is very little difference between doing this and rerunning the simulation using different seeds. To explore this we reran the simulations with new seeds, but using the hedge ratios calculated above. The results were almost identical. A more practical issue in assessing the sensitivity of the results to hedge ratio selection would relate to the simuimulation parameters themselves, about which there is both sampling and structural uncertainty.

EXHIBIT 2

Impact of Correlation between Compound Valuation (V_j) and Per-Patient Cost (C_O^j, C_E^j)

	No Hedge	ρ_{VC} = 0.0	ρ_{VC} = 0.3	ρ_{VC} = 0.6
Risk reduction:*				
Reimbursement shortfall (S_f) ($MM)				
(negative numbers are profits)				
Mean	14.8	–6.0	–5.4	–4.7
95th pctile	30.5	10.4	8.3	6.8
99th pctile	40.3	20.2	16.3	12.8
Cor(ROE_{RBO}, S_f)		0.34	0.56	0.66

NOTES: This exhibit shows the performance of the hedging strategy under different assumptions about the correlation between the exit valuation of a candidate drug sold out of the RBO portfolio and the drug's subsequent per-patient cost in the future. ρ_{VC} denotes the assumed value of this correlation in the simulations, ROE_{RBO} represents the ROE on the RBO portfolio, and S_f represents the shortfall due to increased reimbursement costs.

conservative the hedge (i.e., the higher quantile of H_o chosen), the more it reduces the extreme losses in the tail of the shortfall distribution. However, the cost of the hedge scales in the level of H_o. In addition, the percentage of instances in which the cost of the hedge is greater than the realized shortfall on the portfolio increases. Due to the positive ROE on the portfolio, the hedger typically makes a profit.

Our results suggest that there is significant variability in the preferred hedge ratio. It is also clear that even the larger hedges (i.e., those based on large values of H_o) are not complete, in that for every hedge level there remain cases that are not fully hedged by the RBO equity. We discuss this in more detail below.

The relationship between the RBO equity performance and the hedge effectiveness is nuanced. Consider that one of the primary sources of variability in the simulation of the reimbursement portfolio is the number of diseased individuals; this is, in turn, a function of both the number of diseased individuals and the number of those for whom one of the therapies would have been appropriate. In general, the RBO equity position gains value when new therapies are developed. So all else equal, the number of therapies developed should offset, proportionally, the increase reimbursement-linked shortfall.

However, when the number of diseased individuals treated is much larger, or the cost of certain drugs very much higher, than expected (from the initial Phase II valuation of the compound when it was sold out of the RBO portfolio), the hedge may not fully reduce the shortfall. Conversely, unless there are an unusually large number of individuals with the disease or unexpectedly high costs (or both), even modest hedges appear to improve performance in most cases.

Sensitivity of Results to Assumptions about Correlations between Compound Valuations at Sale and Subsequent Per-Patient Costs

We make a basic assumption that the price at which the collateral manager sells a compound out of the RBO in Phase II is related to unbiased expectations for the future cashflows the drug will generate, should it be approved, given a specific prevalence rate for the disease. Thus, for any target population affected by a disease, higher valuations on the compounds at exit should imply higher per-patient pricing of the approved drug (all else equal). In our baseline simulation, we assumed that the correlation between the valuation at exit and the per-patient cost, which we denote $\rho_{VC} = \text{Cor}(V_j, C_E^j)$, is 0.6. (See Exhibit 2.)

To give a sense of how the performance of the hedging strategy changes with respect to the level of this correlation, we also simulated the same portfolios using different parameter values for ρ_{VC}. Then, because we include the case of $\rho_{VC} = 0$, we may also use the results of these simulations to decompose (coarsely) the drivers of the correlation between the RBO portfolio performance and the effectiveness of the hedge.

Note first that the strategy improved performance regardless of the assumed values of ρ_{VC}. The portfolio that includes RBO equity performed better than the unhedged portfolio in the large majority of cases. In particular, the hedged portfolio always

performed better than the unhedged portfolio (even after factoring in the cost of the hedge) in the quantiles in which reimbursement costs are extremely high.

Furthermore, we see that even when we assume a correlation of zero between the Phase II sale price and future per-patient reimbursement costs, the realized value of Cor(ROE_{RBO}, S_f) is 34%. This represents more than half of the total correlation between ROE_{RBO} and S_f that we observed in our base case ($\rho_{VC} = 0.6$). Thus, it seems that substantial correlation is induced simply because both the total cost of reimbursement for the benefits provider and the value of RBO equity position depend on the number of approved compounds.

Finally, and not surprisingly, the hedge is more effective as the correlation between the RBO exit values and per-patient costs of approved drugs increases. Although the mean value of the hedge declines as the correlation increases, its effectiveness in lowering the tail risk increases.

Imprecision of Hedge Due to Incompleteness and Non-Exclusivity of RBO

Even if the hedge ratio is perfectly matched to the number of potential diseased individuals treated, the hedge still may not be perfect, for a number of reasons.

Ignoring the higher-order statistical issues (see Footnote 8), the most prominent structural reasons for the imprecision in the strategies are: 1) in most cases, the drugs in the RBO portfolio are sold before approval, so their final value (and thus likely costs) are only partially known; 2) the portfolio underlying the RBO likely does not contain drugs for all possible diseases that may impact insured individuals, and may omit therapies for diseases that affect many of them; and 3) even if therapies for the diseases targeted in the RBO portfolio do become available during the lifespan of the diseased individuals, the portfolio (and thus the hedger) only benefits financially if the therapy that comes to market for the patients in the insurance portfolio is the one in the portfolio, rather than a therapy developed outside the portfolio for patients in the insured portfolio.

Therapeutic drugs are sold from the portfolio prior to final approval. In general, drugs are sold out of the RBO portfolio after one or more successful clinical trial phases, not at the very end of the approval process. So in some cases drugs that exit the portfolio never actually reach patients. This creates the potential for a mismatch between the cashflows of the RBO portfolio and the ultimate reimbursement experience of the provider. To some degree, this mismatch is minimized by the fact that for a large majority of drugs sold out of the portfolio, substantial payments are received only when trial milestones are met.

The mismatch also creates uncertainty about the degree to which the cashflows from the sale of the drug out of the portfolio will match the ultimate costs of the drug.

Therapeutic drugs for diseases not targeted in the RBO portfolio. Sometimes a therapy comes to market for a disease not targeted by a compound in the RBO portfolio. This new therapy, against which the reimbursement agent is not hedged though the RBO, may require repayment without providing an offsetting increase in RBO equity value for the underwriter.

A natural way for RBO portfolio managers to address this is to explicitly disclose the diseases for which their portfolio is seeking therapies. This can be done confidentially, and with appropriate safeguards. Investors then can calculate hedge ratios based only on the targeted diseases, hedging out the risk from potential therapies for those specific conditions, while pursuing alternative risk mitigation for the remainder of the portfolio.

Even if the portfolio is not disclosed, or if the portfolio is being dynamically constructed over time, RBO equity investments still can be used to offset potential risk of breakthrough scientific advancements. The underwriter can use RBO equity to

hedge the general scientific "marketplace" rather than to offset specific positions, albeit with a more obvious mismatch.

This type of mismatch is not uncommon in other asset markets. For example, an investor in a portfolio of small US oil and coal mining exploration firms might hedge the risk using put options on the iShares S&P Global Energy Index Fund (IXC). While it is unlikely that all (or in some cases any) of the firms in the portfolio would be represented in the index, the investor might reason that the hedge provides protection against adverse price movements across the sector. In the same way, RBO equity exposure may provide protection against risk associated with more general types of scientific breakthroughs, regardless of whether the RBO portfolio contains all (and only) those in the underwritten cohort.

Such strategies must be constructed carefully, however. In the case of genetic diseases, for example, the inter-drug scientific correlation is typically considered to be quite low, at least initially.[15] This is due in part to the nature of these therapies, which tend to target single genes rather than more complex combinations of pathways (as is sometimes the case in certain types of cancer) or other broader systemic environments. For such portfolios, only very major breakthroughs (e.g., on the order of the Human Genome Project) might influence outcomes. This limitation is not unique to RBOs.[16]

Therapeutic drugs not in the portfolio but for diseases targeted in the RBO portfolio. Consider the case in which the RBO portfolio manager correctly hedges a disease by investing in a portfolio containing a drug that targets it. If a different therapy for the disease makes it to market first, the hedge is worse than ineffective; it is likely costly, as the value of the drug in the portfolio will go down at the same time the reimbursement cost goes up.

This can be addressed to some degree through a slightly more complicated treatment of $\partial A_P / \partial k$, by defining $\Pr(\theta_j \in P | \tau_j, \leq T_d) = p_\theta$ as the probability that the therapy brought to market for condition j is the one in the RBO portfolio, given that a therapy is introduced during the lifespan of the affected individuals in the underwritten population and adjusting the value of the payoff, accordingly.

Future Research

Although our results suggest substantial benefits to using RBO securities for hedging historically hard-to-hedge risks, our analysis is still stylized in a number of respects and thus can be fairly criticized along a number of dimensions. These critiques suggest a number of opportunities for future research.

Our current work is limited by our assumptions regarding the actuarial properties of various at-risk portfolios and the manner in which these are modeled. For example, our model did not contemplate stochastic interest rates when valuing annuity-like exposures at the horizon of the RBO securities' payoffs. Our simulations did not discount annuity cashflows dynamically or attempt to exactly synchronize the valuation times of the annuities with the arrival times of residual "milestone" payments on drugs sold out of the portfolio. We did not consider the potential correlation between the cashflows on the RBO equity tranche and the interest rates used to discount the

[15] It is not uncomon for a genetic breakthrough that initally targets a specific disease to then find applications across a variety of other condititions.

[16] There is evidence, for example, that the current generation of index-based longevity hedging vehicles may suffer from material basis mismatches. In one study, using an augmented Lee and Carter (1992) model, Li and Hardy (2011) estimate that for a model Canadian pension plan, the efficiency of a longevity hedge constructed using *q*-forward contracts (a form of longevity swap based on the LifeMetrics indices) would be 56% of longevity VaR and about 82% of longevity risk.

annuity cashflows, and for convenience, we assumed homogeneity across prevalence rates and other features of the drugs in the portfolio.[17]

There also is a substantial scope for extending our actuarial framework. For example, with respect to longevity risk, recent work on stochastic longevity models (such as Cairns et al. 2006) suggest a number of more refined approaches to modeling (and simulating) lifetimes for annuity valuation.

We chose our examples for exposition rather than realism, and more-realistic examples are certainly possible, subject to the analysis of historical data on pension and insurance portfolios and based on a more detailed cross-sectional analysis of current ones.

In particular, the assumptions we used in our longevity examples about the impact of a new drug on a patient's lifespan are somewhat arbitrary; while some drugs may extend life beyond our assumptions, others may offer a much shorter extension.[18] These assumptions, and those in our reimbursement examples, could be tailored to a specific portfolio of diseases. This is also true of our representation of the dynamics of the processes representing both annuity writing and reimbursement, which we have simplified somewhat in our examples and simulations.

We also have left the optimization of hedge strategies for future work. More-efficient hedge ratio calculations could be developed if more-specific information were available about the health states of the members of the hedger's portfolio and the specific diseases targeted by the therapies in the RBO portfolio. With more-detailed data on the specific diseases in the portfolio, a hedger can more precisely determine the expected cost of a hedge and factor this into the optimization of the hedge ratio. Similarly, more-detailed information about, for example, the insurance premiums paid by individuals in the portfolio, in the case of reimbursement risk, or the contributions of pension fund members, in the case of longevity risk, make more-exact calculations of net hedging costs feasible.

This last topic also suggests an interesting opportunity for customizing RBO portfolios to appeal to investors with differential hedging needs. For example, based on portfolio composition, health insurers may be more interested in therapies for diseases that affect key demographics in their portfolios while pension funds may prefer portfolios that focus on late-in-life diseases.

Finally, in the mathematics we used to motivate our examples, we often ignored higher-order effects related to nonlinearities and uncertainty in expectations.

CONCLUSION

The equity tranche of a research-backed obligation (RBO) gains value when individuals live longer. Entities exposed to certain types of scientific longevity risk, including defined-benefit pension plans and insurance firms that underwrite annuities to individuals at risk for certain diseases, may find RBO equity to be an attractive natural hedging vehicle.

Using simplified mathematics, we have demonstrated that the equity of a research-backed obligation has the potential to mitigate longevity risk and reimbursement risk. We have proposed an approach to estimating an appropriate hedge ratio.

[17] A cursory analysis of the relationship between the yield on five-year treasuries and the value of the iShares Nasdaq Biotechnology Index Fund suggests negative correlation—that is, discount rates tended to be higher at times when biotech valueations were lower.

[18] Furthermore, some drugs in the portfolio may be selected by the manager for benets other than simply extending life: some drugs may dramatically improve the quality of a patient's life without necessarily always extending it. (Presumably, if longevity hedging were the target application for the RBO equity, the portfolio's investment guidlines might limit such purchases.)

To do so, we first introduced an analytic model for the case in which all quantities required to calculate hedge rations were deterministic and known. Because many of these quantities (e.g., RBO portfolio exit valuations, numbers of diseased individuals, per-patient costs of approved drugs, etc.) are, in fact, stochastic, estimated and/or path dependent, we proposed a simulation framework for estimating hedge ratios at various levels of confidence. We observed that in principle the hedges can be made arbitrarily effective in expectation, but that care is required to determine the appropriate hedge size.

Our results demonstrate that it is feasible to use RBO equity to hedge reimbursement and longevity risk—but the practical effectiveness depends on the unique characteristics of the member portfolio of the insurance company and the drug portfolio underlying the RBO. In our simulation experiments, we used assumptions from prior work regarding these characteristics; other researchers or portfolio managers may substitute their own. Our objective was to demonstrate the approach, rather than to determine whether it would work for specific pairs of drug and insurance portfolios.

Managing longevity and reimbursement risks continues to be growing challenges for pension funds and insurance firms. The advent of RBO securities provides one possible mechanism for hedging specific dimensions of this risk. As the RBO market continues to develop, we can envision RBO portfolios designed to meet the needs of specific institutions.

At least in the short term, the size of the RBO market likely will remain small relative to liabilities of pension plans. It is therefore likely that RBOs will be more useful for health benefit payers hedging reimbursement risk than for pension plans hedging longevity risk.

Finally, health benefit payers who wish to both hedge reimbursement risk and match the timing of their future funding needs will find specific budgeting and reserve protocols useful. The arrival of the cashflows from RBO portfolio exits likely will predate the approval of the drugs that exit. As a result, the health benefits payer will receive the cashflow before it is needed to reimburse patients, and so it must be reserved or otherwise accounted for in a manner to ensures it is available at the time of shortfall.

An RBO's ability to pool multiple smaller drug development projects and create portfolios large enough to be investable represents a significant step forward in reimbursement and longevity risk mitigation. In addition to being financially attractive, it offers the welcome social benefit of channeling much-needed funding to drug development in key areas affecting public health, thereby increasing the potential for more and better therapies for millions of patients in need.

In the era of impact investing, many investors consider how to achieve social ends through their investments. RBOs can be especially useful, providing a vehicle by which investors can "do well by doing good."

ACKNOWLEDGMENTS

Much of this work was done while the author was at the MIT Laboratory for Financial Engineering. I am grateful to Zvi Bodie, Jeffrey Bohn, David Fagnan, Jose-Maria Fernandez, John Hull, Ronald Khan, Neil Kumar, Andrew Lo, Seth Merrin, Moshe Milevsky, Robert Pozen, Brian Stephenson, and Dylan Tyson for suggestions and detailed comments on earlier drafts of this article. The article benefited greatly from the comments of an anonymous reviwer. David Fagnan also provided extensive programming support for some portions of the Julia implementation. All errors are, of course, my own. An earlier version of this article (Stein 2016) was circulated under a different title.

REFERENCES

Aetna. 2015. "2014 Aetna Annual Report, Financial Report to Shareholders." Annual Report. https://www.annualreports.com/HostedData/AnnualReportArchive/a/NYSE_AET_2014.pdf.

Aglionby, J., and J. Cumbo. 2014. "BT Pension Fund Strikes Deal to Hedge Longevity Risk." *Financial Times*. July 4. https://www.ft.com/content/ea371a34-0342-11e4-9195-00144feab7de.

Aro, H. 2014. "Systematic and Nonsystematic Mortality Risk in Pension Portfolios." *North American Actuarial Journal* 18 (1): 59–67.

BIS. 2013. "Longevity Risk Transfer Markets: Market Structure, Growth Drivers and Impediments, and Potential Risks." Joint Forum of the Basel Committee on Banking Supervision, Bank for International Settlements. https://www.bis.org/publ/joint31.pdf.

Boycott, K. M., M. R. Vanstone, D. E. Bulman, and A. E. MacKenzie. 2013. "Rare-Disease Genetics in the Era of Next-Generation Sequencing: Discovery to Translation." *Nature Reviews Genetics* 14 (10): 681–691.

Brockett, P. L., S.-I. Chuang, Y. Deng, and R. D. MacMinn. 2013. "Incorporating Longevity Risk and Medical Information into Life Settlement Pricing: Incorporating Longevity Risk Into Life Settlements." *Journal of Risk and Insurance* 80 (3): 799–826.

Cairns, A. J. G., D. Blake, and K. Dowd. 2006. "A Two-Factor Model for Stochastic Mortality with Parameter Uncertainty: Theory and Calibration." *The Journal of Risk and Insurance* 73 (4): 687–718.

Conti, R., A. Fein, and S. S. Bhatta. 2014. "National Trends in Spending on and use of Oral Oncologics, First Quarter 2006 Through Third Quarter 2011." *Health Affairs* 33 (10): 1721–1727.

Express Scripts Lab. 2015. "The 2014 Drug Trend Report." Express Scripts Lab Technical Report. https://corporate-site-labs-dev.s3.amazonaws.com/2019-09/Express%20Scripts%202014%20Drug%20Trend%20Report.pdf.

——. 2020. "The 2019 Drug Trend Report". Express Scripts Lab Technical Report.

Fagnan, D. E., A. A. Gromatzky, R. M. Stein, J.-M. Fernandez, and A. W. Lo. 2014. "Financing Drug Discovery for Orphan Diseases." *Drug Discovery Today* 19 (5): 533–538.

FDA. 2014. "FDA Approves First Combination Pill to Treat Hepatitis C." Press release, October. https://www.hiv.gov/blog/fda-approves-first-combination-pill-to-treat-hepatitis-c.

Fernandes, D. 2015. "Health Care Costs Forecast to Rise 7 Percent." *Boston Globe*, February 5. https://www.bostonglobe.com/business/2015/02/05/health-care-costs-rising-driven-drugs-economy/XYjQjZmbynA5m83bO6PkLJ/story.html.

Fernandez, J., R. Stein, and A. Lo. 2012. "Commercializing Biomedical Research through Securitization Techniques." *Nature Biotechnology* 30: 10.

Hay, M., D. W. Thomas, J. L. Craighead, C. Economides, and J. Rosenthal. 2014. "Clinical Development Success Rates for Investigational Drugs". *Nature Biotechnology* 31 (1).

Herper, M. 2013. "Inside the Pricing of a $300,000-A-Year Drug." *Forbes*, January 3. https://www.forbes.com/sites/matthewherper/2013/01/03/inside-the-pricing-of-a-300000-a-year-drug/?sh=4fd4266a4917.

Hull, J., A. W. Lo, and R. M. Stein. 2019. "Funding Long Shots." *Journal of Investment Management* 17 (4): 1–33.

Hung, W. W., J. S. Ross, K. S. Boockvar, and A. L. Siu. 2011. "Recent Trends in Chronic Disease, Impairment and Disability Among Older Adults in the United States." *BMC Geriatrics* 11 (1): 47.

IMF. 2012. "Global Financial Stability Report." World Economic and Financial Surveys, April. https://www.imf.org/en/Publications/GFSR/Issues/2016/12/31/Global-Financial-Stability-Report-April-2012-The-Quest-for-Lasting-Stability-25343.

Insurance Information Institute. 2019. "Annuities." http://www.iii.org/fact-statistic/annuities.

Kisser, M., J. Kiff, E. S. Oppers, and M. Soto. 2012. "The Impact of Longevity Improvements on US Corporate Defined Benefit Pension Plans." IMF Working Paper WP/12/170. https://www.imf.org/external/pubs/ft/wp/2012/wp12170.pdf.

Lee, R., and L. R. Carter. 1992. "Modeling and Forecasting US Mortality." *Journal of the American Statistical Association* 87 (419): 659–671.

Li, J., and M. Hardy. 2011. "Measuring Basis Risk in Longevity Hedges." *North American Actuarial Journal* 15 (2): 177–200.

Milevsky, M. A., and D. Promislow. 2001. "Mortality Derivatives and the Option to Annuitize." York University Finance Working Paper MM08-1, August. York University. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=289551.

Mott, A. R. 2007. "New Swaps to Hedge Alpha and Beta Longevity Risks of Life Settlement Pools." *The Journal of Structured Finance* 13 (2): 54–61.

NAIC. 2019. "US Health Insurance Industry 2018 Annual Results." NAIC Technical Report. https://naic.org/documents/topic_insurance_industry_snapshots_2018_health_ins_ind_report.pdf.

Nass, D., and A. Simorellis. 2019. "Global Oncology Trends 2019." *Therapeutics, Clinical Development and Health System Implications Institute Report.* IQVIA Institute for Human Data Science, May.

New York State Office of the State Comptroller. 2020. "2019 New York State and Local Retirement System Comprehensive Annual Financial Report." https://www.osc.state.ny.us/retirement/resources/2019-comprehensive-annual-financial-report.

Olsen, K. 2012. "GM-Prudential Annuity Deal Seen Setting Off Race to Gain Piece of DB Derisking Action." Pensions & Investments.com, June 11. https://www.pionline.com/article/20120611/PRINT/306119975/gm-prudential-annuity-deal-seen-setting-off-race-to-gain-piece-of-db-derisking-action.

O'Sullivan, C. C., I. Bradbury, C. Campbell, M. Spielmann, E. A. Perez, H. Joensuu, J. P. Costantino, S. Delaloge, P. Rastogi, D. Zardavas, K. V. Ballman, E. Holmes, E. de Azambuja, M. Piccart-Gebhart, J. A. Zujewski, and R. D. Gelber. 2015. "Efficacy of Adjuvant Trastuzumab for Patients with Human Epidermal Growth Factor Receptor 2–Positive Early Breast Cancer and Tumors 2 cm: A Meta-Analysis of the Randomized Trastuzumab Trials." *Journal of Clinical Oncology* 20 (33): 24.

Parnes, D. 2018. "Analyzing the Different Modules of Longevity Risk in Death Bonds." *The Journal of Alternative Investments* 21 (3): 35–44.

Perera, N., and B. Reeves. 2006. "Risk Mitigation for Life Settlements." *The Journal of Structured Finance* 12 (2): 55–60.

PhARMA. 2013. "Medicines in Development: Rare Diseases: A Report on Orphan Drugs in the Pipeline." Technical Report.

——. 2015. "2015 Profile: Biopharmaceutical Research Industry". Technical Report. PhARMA (Pharmaceutical Research and Manufactures of America). http://phrma-docs.phrma.org/sites/default/files/pdf/2015_phrma_profile.pdf.

Smith, B. B., and S. L. Washington. 2006. "Acquiring Life Insurance Portfolios: Diversifying and Minimizing Risk." *The Journal of Structured Finance* 12 (2): 41–45, 9–10.

Stein, R. M. 2016. "A Simple Hedge for Longevity Risk and Reimbursement Risk Using Research-Backed Obligations." MIT Sloan School of Business working paper 5165-16, February. http://www.rogermstein.com/wp-content/uploads/SteinLongevity.pdf.

Stone, C. A., and A. Zissu. 2006. "Securitization of Senior Life Settlements: Managing Extension Risk." *The Journal of Derivatives* 13 (3): 66–72,11.

——. 2008. "Using Life Extension-Duration and Life Extension-Convexity to Value Senior Life Settlement Contracts." *The Journal of Alternative Investments* 11 (2): 94–108, 6.

Tichy, E. M., G. T. Schumock, J. M. Hoffman, K. J. Suda, M. H. Rim, M. Tadrous, J. Stubbings, S. Cuellar, J. S. Clark, M. D. Wiest, L. M. Matusiak, R. J. Hunkler, and L. C. Vermeulen. 2020. "National Trends in Prescription Drug Expenditures and Projections for 2020." *American Journal of Health-System Pharmacy* 77 (15): 1213–1230.

US Census Bureau. 2020a. "2019 American Community Survey (ACS)." Technical report. https://www.census.gov/programs-surveys/acs.

——. 2020b. "Annual Estimates of the Resident Population for Selected Age Groups by Sex for the United States: April 1, 2010 to July 1, 2019." Annual Estimates of the Resident Population for Selected Age Groups by Sex for the United States: April 1, 2010 to July 1, 2019.