

Optimal annuity demand for general expected utility agents

Carole Bernard^{a,b,1}, Luca De Gennaro Aquino^{a,*}, Lucia Levante^c

^a Department of Accounting, Law and Finance, Grenoble Ecole de Management, F-38000 Grenoble, France

^b Department of Economics and Political Sciences, Vrije Universiteit Brussel (VUB), Belgium

^c Dipartimento di Metodi e Modelli per l'Economia, il Territorio e la Finanza, Università di Roma "La Sapienza", Italy

ARTICLE INFO

Article history:

Available online 9 July 2020

Keywords:

Expected utility theory
Annuity equivalent wealth
Longevity risk pooling
Life-cycle model
Annuity puzzle

ABSTRACT

We study the robustness of the results of Milevsky and Huang (2018) on the optimal demand for annuities to the choice of the utility function. To do so, we first propose a new way to span the set of all increasing concave utility functions by exploiting a one-to-one correspondence with the set of probability distribution functions. For example, this approach makes it possible to present a five-parameter family of concave utility functions that encompasses a number of standard concave utility functions, e.g., CRRA, CARA and HARA. Second, we develop a novel numerical method to handle the life-cycle model of Yaari (1965) and the annuity equivalent wealth problem for a general utility function. We show that the results of Milevsky and Huang (2018) on the optimal demand for annuities proved in the case of a CRRA and logarithmic utility maximizer hold more generally.

© 2020 Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

Longevity risk refers to two components: systematic longevity risk and idiosyncratic longevity risk (Hanewald et al., 2013). The former is non-diversifiable and linked to unexpected changes in mortality rates for aging cohorts; the latter is instead linked to the individuals' life expectancy and can be reduced by pooling the risk among agents with different ages, incomes, health conditions and wealth (it is also called "personal longevity risk" by Milevsky (2018)). Individuals face such longevity risk when reaching their retirement age to ensure they have enough wealth to avoid outliving their own savings. In this paper, we mainly focus on the benefits for agents of insuring this second type of longevity risk by purchasing life annuities.

As already pointed out by Yaari (1965), a rational consumer should invest 100% of her wealth in life annuities if she aims at maximizing her expected lifetime utility, assuming that she has access to a frictionless market in which life annuities are offered at fair prices and that she has no bequest motives. A lifetime annuity is an insurance contract between an insurance company and a policyholder in which the latter pays a one-time premium, or a periodic flow of premiums, to obtain in return a monthly

or annual amount of money until her time of death. This product allows retirees to manage the risk of outliving their savings and to provide them a regular flow of income after reaching retirement age. Despite the latest improvements within the market, the demand for annuities continues to be low, resulting in an intense debate which goes by the name of "annuity puzzle". It is therefore necessary to understand this phenomenon by analyzing all the reasons behind the annuity demand and measuring the gain in terms of utility for a potential annuitant.

The approach proposed in Milevsky and Huang (2018) is inspired by the classical life-cycle utility framework of Yaari (1965). The authors consider a retiree both in the presence and in the absence of a preexisting pension income, and where the value of the longevity risk pooling is measured starting from the so-called *annuity equivalent wealth* (AEW), which is defined as "the additional amount of wealth that a consumer would need if he did not have access to an annuity market, in order to achieve the same lifetime expected utility level that he could achieve by using that wealth to purchase a nominal annuity" (Mitchell, 2002). Under some assumptions, Milevsky and Huang (2018) obtain a closed-form expression to evaluate the value of longevity risk pooling δ (which is directly related to the AEW), with fixed life annuities and assuming a *constant relative risk aversion* (CRRA) utility function. Our goal is to check the robustness of their results by changing the type of utility function used in the calculations, that is, by considering more general concave increasing utility preferences.

We contribute to the literature as follows. Firstly, we propose a novel way to test the robustness of results obtained with a

* Corresponding author.

E-mail addresses: carole.bernard@grenoble-em.com (C. Bernard), luca.degennaroaquino@grenoble-em.com (L. De Gennaro Aquino), levante.1538395@studenti.uniroma1.it (L. Levante).

¹ C. Bernard acknowledges funding from the FWO (Odysseus project GOF7316N).

specific utility function. Specifically, we introduce a new family of utility functions, parametrized by a continuous probability distribution function, that encompasses all concave increasing utility functions satisfying Inada conditions. By choosing, for instance, the five-parameter generalized beta distribution (McDonald and Xu, 1995; McDonald, 2008), such parametrization makes it possible to generate a very large family of utility functions: by varying the five parameters, we show that it is possible to obtain CRRA-, CARA- and HARA-type utility functions as special cases. Secondly, given that we use a more general utility function than the common power utility, explicit solutions are not any more in reach. Thus, here we solve the life-cycle model mainly numerically: we start from the general solution of the problem (available, for instance, in Leung (2007)), and then use a combination of standard and more “recent” tools (such as neural networks) to perform the computations. Thirdly, we illustrate our approach by studying the robustness of the results in Milevsky and Huang (2018) about the magnitude of the AEW and of the related longevity risk pooling coefficient obtained using the CRRA utility function.

The paper is organized as follows. In Section 2, we present the life-cycle model discussed in Milevsky (2018), Milevsky and Huang (2018) and firstly introduced by Yaari (1965). Under a CRRA utility framework and some additional assumptions, Milevsky and Huang (2018) obtain closed-form expressions for the value of longevity risk pooling with fixed life annuities. Our objective is to check the robustness of these results with respect to the choice of the utility function. In Section 3, we propose an alternative numerical approach to solve the life-cycle model when the utility function does not allow to simplify the analysis. Differently to what is typically done in the literature, our approach does not rely on a dynamic programming algorithm, but rather focuses on the general (quasi-explicit) solution (as in Leung (2007)) and an application of neural networks. In Section 4, we present a new way to span the universe of increasing concave utility functions by parametrizing each utility by a (cumulative) distribution function. We show that, by choosing a wide family of distribution functions, it is possible to obtain a general utility function that includes as special or limiting cases most well-known utility functions. Equipped with the methodology presented in Section 3 and with the utility functions defined in Section 4, we challenge the robustness of the results of Milevsky (2018), Milevsky and Huang (2018) and compute the longevity risk pooling for many different utility functions in Section 5. Finally, Section 6 concludes.

2. Problem

Milevsky and Huang (2018) develop a new practical method to quantify, under some specific assumptions, the benefits arising from annuitizing even when a retiree is already endowed with a periodic pension income $\pi > 0$. In a seminal paper, Yaari (1965) showed that a rational consumer with no bequest motives would ideally place all of her wealth into actuarially-fair life annuities instead of standard bonds. Within this setting, purchasing annuities is thus the optimal investment choice from a retiree standpoint. To quantify the benefits of annuitizing, Milevsky and Huang (2018) introduce the value of longevity risk pooling δ as a new metric, expressed as a percentage of the initial retirement wealth, which captures the potential utility gain of annuitizing. δ denotes the extra amount of money an individual would require to compensate for the fact that she does not have access to annuitization. Specifically, δ is the percentage by which a retiree needs to increase her initial wealth in order to reach the same level of expected utility of a full annuitization.

2.1. Life-cycle model

Referring to a classical life-cycle model and choosing to operate within the (Yaari, 1965) framework, Milevsky and Huang (2018) consider an individual at age x , with an initial (non-annuitized) wealth w at time 0 and a fixed periodic pension income π , which is paid until the time of death, and whose goal is to maximize her discounted lifetime expected utility $\mathcal{U}_x(w, \pi)$. The pension income π is a periodic cash-flow that is paid annually until death. In the absence of an annuity market, we assume that the retiree has the opportunity to invest her wealth w in order to receive a risk-free rate r . Note that, at time of retirement, the total endowment of the retiree is $w + PV(\pi)$, where $PV(\pi)$ denotes the present value of the lifetime annual income of π . This present value is related to the annuity factor a_x so that $PV(\pi) = a_x \pi$.

The maximal discounted lifetime utility for an individual at age x can be written as

$$\mathcal{U}_x^*(w, \pi) = \max_{c_t} \int_0^\infty e^{-\rho t} {}_t p_x U(c_t) dt, \quad (1)$$

where c_t is the consumption at time t , $\rho > 0$ is the subjective discount rate and, for the ease of exposition, it is set equal to the constant risk-free interest rate r , so $\rho = r$.² Note that we assume perfect credit markets. Furthermore, ${}_t p_x$ refers to the individual's survival probabilities, and U is an increasing strictly concave utility function. The consumer also faces a dynamic budget constraint that can be expressed with the following differential equation:

$$dW_t = (rW_t + \pi - c_t)dt, \quad W_0 = w, \quad (2)$$

where W_t denotes the wealth at time t . Note that $\mathcal{U}_x^*$ depends on the percentage of preexisting pension income π as it directly affects the optimal consumption c_t^* that maximizes the lifetime expected utility.

As it was already proved in the classical life-cycle model of Yaari (1965), that is frictionless and in which the agent has no bequest motives, it is optimal to annuitize all her initial wealth. Thus the expected utility increases if the individual invests part of her initial wealth in buying life annuities. This could be expressed as follows:

$$\mathcal{U}_x^* \left(w - \varepsilon, \pi + \frac{\varepsilon}{a_x} \right) \geq \mathcal{U}_x^*(w, \pi),$$

where ε is the amount spent on annuities, and a_x is the annuity factor. This inequality is valid for every concave increasing utility function assuming a frictionless annuity market (without loadings (fair pricing), fees, asymmetric information and adverse selection, and assuming that the agent has no bequest motives). The maximum discounted lifetime utility thus increases for every amount ε . Furthermore, in the case $\varepsilon = w$ (meaning that the entire liquid wealth w is annuitized), the *annuity equivalent wealth* (AEW), denoted as $\hat{w}$, is defined as the solution to the following equation:

$$\mathcal{U}_x^*(\hat{w}, \pi) = \mathcal{U}_x^* \left(0, \pi + \frac{w}{a_x} \right).$$

The AEW $\hat{w}$ and the value of longevity risk pooling δ are then directly related by the following relationship: $\delta = \frac{\hat{w}}{w} - 1$, which follows immediately from the definition of δ as the solution to:

$$\mathcal{U}_x^*((1 + \delta)w, \pi) = \mathcal{U}_x^* \left(0, \pi + \frac{w}{a_x} \right),$$

which is the percentage increase of initial wealth required to compensate for the loss due to the absence of an annuity market.

We can now move to the specific assumptions made by Milevsky and Huang (2018) to derive a closed-form expression for δ and draw elegant interpretations.

² This is consistent with what Milevsky and Huang (2018) require but it does not represent additional difficulty for our numerical approach.

2.2. Mortality model

Milevsky and Huang (2018) first consider the case of an exponential remaining lifetime. They then consider a more specific and realistic framework, by adopting the Gompertz mortality model. This choice changes the computation of the survival probabilities ${}_t p_x$, as the individual's hazard rate at age x is no longer constant, but given by:

$$\lambda_x = \frac{1}{b} e^{\frac{x-m}{b}},$$

where m is the modal value of life and b is the dispersion coefficient, both expressed in years. Despite its simplicity, the Gompertz mortality model is still widely used in many insurance economic works and in the actuarial modeling as well.

2.3. Longevity pooling with a CRRA utility function

The Annuity Equivalent Wealth (AEW) measure, and thus the longevity risk pooling coefficient δ , depends on many factors, such as the specific law of mortality, the consumer's utility function, the interest rate, the amount of preexisting pension income π and initial wealth w .

Proposition 1. *In the case of a CRRA utility function $U(y) = \frac{y^{1-\gamma}}{1-\gamma}$, where γ is the Arrow–Pratt coefficient of relative risk aversion, which also captures the retiree's attitude to longevity risk, and assuming a Gompertz mortality model, Milevsky and Huang (2018) show that the optimal consumption, when $\pi = 0$, is*

$$c_t^* = \left(\frac{w}{\int_0^\infty e^{-rt} ({}_t p_x)^{\frac{1}{\gamma}} dt} \right) ({}_t p_x)^{\frac{1}{\gamma}},$$

and the value of longevity risk pooling δ is

$$\delta_0 = \left(\frac{a_x}{a_x^*} \right)^{\frac{\gamma}{1-\gamma}} - 1.$$

It is obtained as a power of the ratio of two annuity factors, with a_x being the classical annuity factor at age x and a_x^* being the annuity factor for a modified risk-adjusted age $x^* = x - b \log(\gamma)$, where b is the Gompertz parameter and γ is the risk aversion coefficient.

Continuing to assume CRRA preferences for the retiree and a Gompertz mortality model, for the case in which $\pi > 0$ the computation of value of δ becomes more involved and needs to be done numerically. Indeed, for replicability and transparency reasons, the SSRN working paper version of Milevsky and Huang (2018) provides a very valuable step-by-step R code. The remainder of the paper deals with testing the robustness of these results, specifically of Proposition 1 to using other utility functions.

3. Optimal consumption and AEW

3.1. Description of the approach

In this section, we will explain our approach for solving the optimal consumption and the AEW problem in the case of a general utility function.

Firstly, let us recall the generic expression of the solution of the life-cycle model in (1). This can be found, for instance, in Leung (2007) or Charupat et al. (2012). Denote as U an increasing and strictly concave utility function, ${}_t p_x$ the survival probability at time t of an agent aged x , α_t the subjective discount function, π_t the periodic pension income, r_t the interest rate, $\bar{T}$ the maximum

lifetime. Then, the optimal consumption path for the problem in (1) is given by

$$c_t^* = \begin{cases} (U')^{-1} \left(\frac{{}_t p_x \alpha_t U'(\pi_t) e^{\int_t^{\bar{T}} r_s ds}}{{}_t p_x \alpha_t} \right) & \text{for } t \in [0, \tau], \\ \pi_t & \text{for } t \in [\tau, \bar{T}], \end{cases} \quad (3)$$

where $(U')^{-1}$ denotes the inverse function of U' and τ is the wealth depletion time (which, by definition, is always less than or equal to $\bar{T}$). When $\pi_t = \pi$, $r_t = r$, $\alpha_t = e^{-\rho t}$ and $\rho = r$, the expression in (3) boils down to

$$c_t^* = \begin{cases} (U')^{-1} \left(\frac{{}_t p_x U'(\pi)}{\pi} \right) & \text{for } t \in [0, \tau], \\ \pi & \text{for } t \in [\tau, \bar{T}]. \end{cases} \quad (4)$$

Furthermore, when U is a CRRA function, we obtain

$$c_t^* = \begin{cases} \pi \left(\frac{{}_t p_x}{\pi} \right)^{\frac{1}{\gamma}} & \text{for } t \in [0, \tau], \\ \pi & \text{for } t \in [\tau, \bar{T}], \end{cases} \quad (5)$$

which coincides with Eq. (11) of Milevsky and Huang (2018). As our aim is to use a more general utility function than a CRRA, while keeping all other assumptions equal, we will focus on the formulation in (4).

Before discussing the longevity risk pooling problem, an essential preliminary step is being able to compute the wealth depletion time τ . By slight adaptation of the argument in Leung (2007), this is equivalent to locating the root t^* of the following expression

$$\phi(t) = w - \int_0^t e^{-rz} \left((U')^{-1} \left(\frac{{}_z p_x U'(\pi)}{{}_z p_x} \right) - \pi \right) dz = 0, \quad (6)$$

where w is the initial wealth of the agent. For standard, invertible utility functions, solving this problem numerically is not particularly hard. However, in our case the difficulty stems from the inverse of the first derivative of the utility function, which could not be available analytically when the utility is of a more general form. To overcome this issue, we develop a simple application of neural networks to approximate the inverse of a function over a finite interval. In Appendix, we provide the details of the algorithm and some examples. This approach would be needed, for instance, in the case of a SAHARA utility function or in the setting analyzed in Section 5.2.


We can now move to the problem of computing the longevity risk pooling coefficient δ . Following (Milevsky and Huang, 2018), this problem consists in finding the coefficient δ that solves the following equality,

$$\begin{aligned} & \int_0^\tau e^{-rt} {}_t p_x U(c_t^{**}) dt + U(\pi) \int_\tau^{\bar{T}-x} e^{-rt} {}_t p_x dt \\ &= U \left(\frac{w}{(1+\kappa) \text{ILA}(x)} + \pi \right) \text{ILA}(x), \end{aligned} \quad (7)$$


where the optimal consumption c_t^{**} on the left hand side is the solution of the problem in (1) for an initial budget $\hat{w} = w(1+\delta)$. Also, the upper bound $\bar{T}-x = b \log \left(1 + 10 \log(10) e^{\frac{m-x}{b}} \right)$ is such that, under a Gompertz law of mortality (for some parameters b and m), the conditional survival probability from age x to $\bar{T}$ is ${}_{\bar{T}-x} p_x = 10^{-10}$; κ denotes the insurance loading and the immediate life annuity (ILA) is defined as follows,

$$\text{ILA}(x) := \int_0^{\bar{T}-x} e^{-rt + e^{\frac{x-m}{b}}} \left(1 - e^{\frac{t}{b}} \right) dt.$$

Intuitively, this problem corresponds to finding the additional amount of wealth (here expressed as a percentage δ) that an



(a) Fraction of balance sheet annuitized $\psi_0 = 0.5$, CRRA $\gamma \in [0.5, 5]$, insurance loading $\kappa \in [0, 0.3]$.



(b) Insurance loading $\kappa = 0$, fraction of balance sheet annuitized $\psi_0 \in [0.5, 0.9]$, risk aversion coefficient $\gamma \in [0.5, 5]$.

Fig. 1. Longevity risk pooling δ for a 65-year-old CRRA agent with initial wealth $w = 10$.

Table 1

CRRA $\gamma \in [0.5, 5]$, $\kappa \in [0, 0.3]$, $\psi_0 \in [0.5, 0.9]$.

ψ_0, κ	γ					
	0.5	1	2	3	4	5
$\psi_0 = 0.5, \kappa = 0.0$	0.24	0.33	0.44	0.53	0.59	0.64
$\psi_0 = 0.5, \kappa = 0.1$	0.12	0.20	0.31	0.38	0.44	0.49
$\psi_0 = 0.5, \kappa = 0.2$	0.02	0.1	0.19	0.26	0.31	0.35
$\psi_0 = 0.5, \kappa = 0.3$	-0.06	0.01	0.09	0.15	0.20	0.23
$\psi_0 = 0.7, \kappa = 0.0$	0.18	0.25	0.35	0.42	0.47	0.52
$\psi_0 = 0.9, \kappa = 0.0$	0.11	0.15	0.21	0.25	0.29	0.32

agent would need in order to achieve the same lifetime expected utility level that he could achieve by purchasing a nominal annuity.

The way we solve the equality (7) with respect to δ is essentially similar to the approach proposed by Milevsky and Huang (2018). Firstly, we compute the difference between the left hand side (LHS) and the right hand side (RHS) for a rough grid of values for δ (say, equally spaced with step 0.1). Then, once we find the interval where δ must be located (as the difference LHS–RHS has switched in sign), we discretize the grid more finely with step 0.001. We compute again LHS–RHS for increasing values of δ and pick the first one that yields the change of sign.

In Fig. 1 we compute the values of δ for a 65-year-old ($x = 65$) CRRA agent with initial wealth $w = 10$. The parameters of the Gompertz model, m and b , are set as 81 and 11.5, respectively, and $\rho = r = 0.025$. In Fig. 1(a) we fix the fraction of personal balance sheet annuitized $\psi_0 := \frac{\pi \text{ILA}(x)}{w + \pi \text{ILA}(x)} = 0.5$ and set the risk aversion $\gamma \in [0.5, 5]$ and the insurance loading $\kappa \in [0, 0.3]$. In Fig. 1(b), we fix $\kappa = 0$ and choose $\gamma \in [0.5, 5]$ and $\psi_0 \in [0.5, 0.9]$. In Table 1, we provide a few additional reference values in tabulated format.

As it is intuitive and already noted by Milevsky and Huang (2018), we observe in Fig. 1 and Table 1 that the value of longevity risk pooling increases with the risk aversion γ and with the fraction of personal balance sheet annuitized ψ_0 , but obviously decreases with the insurance loading κ (which measures the costs associated to the purchase of annuities and thus the departure to the “fair” pricing). Also, consistently with Yaari (1965), when $\kappa = 0$, we note that δ is always positive.

4. A general concave utility function

In this section, we construct a family of utility functions that encompasses all standard utility functions commonly used in

the literature, such as the CRRA, logarithmic, CARA, HARA and SAHARA utility functions. We parametrize each utility function by a cumulative distribution function (CDF) and show that, the more general the CDF, the larger the class of utility functions that we can possibly span.

This approach is inspired by the optimal portfolio choice problem studied in Bernard et al. (2015). In particular, they show how the expected utility paradigm can rationalize all optimal investment choices made by investors with preferences that are consistent with first-order stochastic dominance. They derive an explicit expression of a utility function that can be used to explain the investor’s demand for any given increasing continuous distribution F of final wealth. This utility function depends on the pricing kernel in the financial market. In the present context, we do not need to model a financial market and thus we do not use their general expression, rather get some inspiration on a new way to describe all suitable increasing concave utility functions.

4.1. Construction of the utility function

Let Φ denote the CDF of a standard normal distribution. We define a utility function over a (possibly infinite) interval (y_1, y_2) as follows,

$$U(y) = \int_c^y e^{-\Phi^{-1}(F(s))} ds, \quad (8)$$

where F is a continuous CDF with support (y_1, y_2) and $c \in (y_1, y_2)$ is chosen arbitrarily.³ Also, note that the utility function in (8) is strictly increasing over (y_1, y_2) , concave for any choice of the continuous distribution F with support (y_1, y_2) , and twice differentiable, with first and second derivative given as

$$U'(y) = e^{-\Phi^{-1}(F(y))},$$

$$U''(y) = \frac{-f(y)e^{-\Phi^{-1}(F(y))}}{\phi(\Phi^{-1}(F(y)))},$$

where ϕ denotes the probability distribution function (PDF) of a standard normal distribution and f denotes the derivative of F , which is of course the associated PDF. Moreover, the Inada conditions are satisfied, that is, $U'(y_1) = +\infty$ and $U'(y_2) = 0^+$.

This expression of the utility function covers all concave increasing twice differentiable utility functions. Indeed, in the following proposition we show that there exists a one-to-one

³ Since utility functions can be defined up to a constant term, the choice of c in (8) does not influence the modeling of agents’ preferences.

correspondence between the set of utility functions and the set of all continuous distribution functions.

Proposition 2. Any continuously differentiable utility function U that is increasing, concave over (y_1, y_2) , and is such that Inada conditions are satisfied ($U'(y_1) = +\infty$ and $U'(y_2) = 0^+$), can be written as in (8), wherein F denotes a continuous CDF with support (y_1, y_2) .

Proof. Consider a continuously differentiable utility function U that is increasing, concave over (y_1, y_2) , and is such that Inada conditions are satisfied. Then $U'(y)$ is well-defined and strictly positive at any point of (y_1, y_2) . Define

$$F(y) := \Phi(-\log(U'(y))). \quad (9)$$

From the concavity of U , U' is decreasing and thus F as defined in (9) is increasing over (y_1, y_2) . Moreover, due to the Inada conditions, it is such that the limit of F at y_1 is 0 and the limit at y_2 is 1. Finally, as U' is continuous, F is a cumulative distribution function of a continuous distribution over a support (y_1, y_2) , being strictly increasing from y_1 to y_2 . $\square$

4.2. Parametrization of the utility function

In order to parametrize our utility function, we propose to use the five-parameter generalized beta (GB) distribution function, which was introduced by McDonald and Xu (1995). This can be characterized by the following probability density function,

$$f_{GB}(y; a, b, \zeta, p, q) = \frac{|a| y^{ap-1} \left(1 - (1 - \zeta) \left(\frac{y}{b}\right)^a\right)^{q-1}}{b^{ap} B(p, q) \left(1 + \zeta \left(\frac{y}{b}\right)^a\right)^{p+q}}, \quad (10)$$

defined for $0 < y^a < \frac{b^a}{1-\zeta}$, $0 \leq \zeta \leq 1$ and b, p, q positive. Also, $B(p, q)$ denotes the beta function. As shown in McDonald (2008), the GB family includes most of the well-known univariate distributions as special or limiting cases: lognormal, half-normal, χ^2 , exponential, power function and many other common distributions. Its range of applications has been discussed in several papers, such as Eugene et al. (2002), Ye et al. (2012), the more recent (Núñez et al., 2019) and references therein.

Of course, other choices could have been made. For instance, we could have similarly chosen the closely related exponential generalized beta (EGB) distribution, which is given by the distribution of the random variable $Z = \log(Y)$, for $Y \sim f_{GB}(y; a, b, \zeta, p, q)$. For a yet larger family of distributions, we could have also considered the generalized beta-generated (GBG) distributions, introduced by Alexander et al. (2012). Given a parent distribution $F(y; \tau)$ with parameter vector τ , support $\mathcal{S}$ and density $f(y; \tau)$, a GBG is characterized by the following density

$$f_{GBG}(y; \tau, a, b, \zeta) = \zeta B(a, b)^{-1} f(y; \tau) F(y; \tau)^{a\zeta-1} (1 - F(y; \tau)^\zeta)^{b-1},$$

defined for $y \in \mathcal{S}$, with positive parameters a, b, ζ . As an example, when $F(y; \tau)$ is taken to be the uniform, we obtain the generalized beta distribution of the first kind (GB1), which corresponds to $f_{GB}(y; a, b, \zeta = 0, p, q)$.

For a review of different methods for generating families of univariate continuous distributions, see Lee et al. (2013).

4.3. Examples and utility functions family tree

In this section, we show that, by choosing an appropriate class of distribution functions and varying the parameters, from our suggested construction in ((8)) we are able to retrieve all the

special cases of utility functions commonly used in the literature.⁴ By means of the formula in (9), we will compute the CDFs that are associated with each of the utility functions that we will consider. All such CDFs will turn out to be related to the normal or lognormal distribution, meaning that, by using the generalized beta distribution in (10), possibly modulo a transformation, we are able to generate a very large family of utility functions.

CRRA utility. Firstly, let us start with a classic *constant relative risk aversion* (CRRA) power utility with parameter $\gamma > 0$:

$$U(y) = \frac{y^{1-\gamma} - 1}{1 - \gamma}. \quad (11)$$

Given $U'(y) = y^{-\gamma}$, from (9) we obtain that

$$F_{CRRA}(y) = \Phi(-\log(y^{-\gamma})) = \Phi(\gamma \log(y)),$$

which is the CDF of a Lognormal $\left(0, \frac{1}{\gamma}\right)$ (i.e., of a random variable that can also be written as $\exp\left(\frac{Z}{\gamma}\right)$, where Z is a standard normal distribution). Trivially, when $\gamma \rightarrow 1$, the CRRA utility converges to a simple case of logarithmic utility $U(y) = \log(y)$, whose more general version is discussed next.

Logarithmic utility. We consider then a logarithmic utility function with parameters $a > 0$ and d :

$$U(y) = \log(ay + d), \quad (12)$$

defined for $y > -\frac{d}{a}$. Thus, given $U'(y) = \frac{a}{ay + d}$, in this case the corresponding CDF is

$$F_{log}(y) = \Phi\left(-\log\left(\frac{a}{ay + d}\right)\right) = \Phi\left(\log\left(y + \frac{d}{a}\right)\right), \quad (13)$$

which is the CDF of a shifted Lognormal $(0, 1)$ with shift $-\frac{d}{a}$.

CARA utility. Now we consider a *constant absolute risk aversion* (CARA) exponential utility with parameter $\gamma > 0$:

$$U(y) = -e^{-\gamma y}. \quad (14)$$

As $U'(y) = \gamma e^{-\gamma y}$, we have that

$$F_{CARA}(y) = \Phi(-\log(e^{-\gamma y})) = \Phi(\gamma y),$$

which is the CDF of a Normal $\left(0, \frac{1}{\gamma}\right)$.

HARA utility. For constants $a > 0$, $\gamma > 0$ and $d \in \mathbb{R}$, a *hyperbolic absolute risk aversion* (HARA) utility function is given as follows:


$$U(y) = \frac{\gamma}{1 - \gamma} \left(\frac{ay}{\gamma} + d\right)^{1-\gamma}, \quad (15)$$

defined for $y > -\frac{d\gamma}{a}$. As $U'(y) = a \left(\frac{ay}{\gamma} + d\right)^{-\gamma}$, we get


$$\begin{aligned} F_{HARA}(y) &= \Phi\left(-\log\left(a \left(\frac{ay}{\gamma} + d\right)^{-\gamma}\right)\right) \\ &= \Phi\left(\gamma \log\left(a^{-\frac{1}{\gamma}} \left(\frac{ay}{\gamma} + d\right)\right)\right), \end{aligned}$$

which is the CDF of a shifted Lognormal $\left(\log\left(\gamma a^{\frac{1-\gamma}{\gamma}}\right), \frac{1}{\gamma}\right)$ with shift $-\frac{d\gamma}{a}$.

⁴ For an alternative method for constructing family of utility functions, see Brockett and Golden (1987).



(a) Fraction of balance sheet annuitized $\psi_0 = 0.5$, CARA $\gamma \in [0.1, 0.6]$, insurance loading $\kappa \in [0, 0.3]$.



(b) Insurance loading $\kappa = 0$, fraction of balance sheet annuitized $\psi_0 \in [0.5, 0.9]$, CARA $\gamma \in [0.1, 0.6]$.

Fig. 2. Longevity risk pooling δ for a 65-year-old CARA agent with initial budget $w = 10$.

Notice that this class of utility functions includes the previous ones as special or limiting cases:

- a CRRA utility with risk aversion parameter γ is obtained in the case when $\gamma > 0$, $d = 0$ and $a = \gamma^{-\frac{\gamma}{1-\gamma}}$;
- an exponential utility with risk aversion parameter a is obtained when $d = 1$, $\gamma \rightarrow \infty$;
- a logarithmic utility with parameters a and d (defined similarly as in (12)) is obtained when $\gamma = 1$.

SAHARA utility. Finally, we consider the case of the *symmetric asymptotic hyperbolic absolute risk aversion* (SAHARA) utility function introduced by Chen et al. (2011). For positive parameters α , β , a SAHARA utility function is given by:

$$U(y) = \begin{cases} -\frac{1}{\alpha^2 - 1} (y + \sqrt{\beta^2 + y^2})^{-\alpha} (y + \alpha\sqrt{\beta^2 + y^2}) & \text{if } \alpha \neq 1, \\ \frac{1}{2} \log(y + \sqrt{\beta^2 + y^2}) + \frac{1}{2\beta^2} y (\sqrt{\beta^2 + y^2} - y) & \text{if } \alpha = 1. \end{cases} \quad (16)$$

The derivative can be computed as $U'(y) = (y + \sqrt{\beta^2 + y^2})^{-\alpha}$, equally for $\alpha \neq 1$ and $\alpha = 1$. By plugging this expression into (9), we get

$$F_{\text{SAHARA}}(y) = \Phi \left(\alpha \log(\beta) + \alpha \sinh^{-1} \left(\frac{y}{\beta} \right) \right),$$

which is the CDF of a Johnson S_U -distribution (Johnson, 1949), itself a transformation of the normal distribution.

5. Robustness to the choice of utility function

Equipped with the approach discussed in Section 3, here we show that the results of Milevsky (2018) and Milevsky and Huang (2018) on the optimal demand for annuities, proved in the case of a CRRA and logarithmic utility maximizer, hold more generally.

5.1. Special cases of known utility functions

In what follows, we illustrate the sensitivity of δ as a function of the insurance loading (fee) κ applied by the insurance company, the risk aversion of the agent and her fraction of balance sheet annuitized ψ_0 . Throughout the section, we keep the same assumptions of Section 3: in particular, $\rho = r = 0.025$, the agent's age $x = 65$, the Gompertz parameters $m = 81$ and

Table 2

CARA $\gamma \in [0.1, 0.6]$, $\kappa \in [0, 0.3]$, $\psi_0 \in [0.5, 0.9]$.

ψ_0, κ	γ					
	0.1	0.2	0.3	0.4	0.5	0.6
$\psi_0 = 0.5, \kappa = 0.0$	0.16	0.21	0.25	0.28	0.3	0.33
$\psi_0 = 0.5, \kappa = 0.1$	0.05	0.09	0.12	0.15	0.17	0.19
$\psi_0 = 0.5, \kappa = 0.2$	-0.04	0.0	0.02	0.05	0.07	0.09
$\psi_0 = 0.5, \kappa = 0.3$	-0.12	-0.09	-0.06	-0.04	-0.02	-0.01
$\psi_0 = 0.7, \kappa = 0.0$	0.16	0.21	0.25	0.28	0.3	0.33
$\psi_0 = 0.9, \kappa = 0.0$	0.16	0.21	0.25	0.28	0.3	0.33

$b = 11.5$, $\bar{T} - x = b \log \left(1 + 10 \log(10) e^{\frac{m-x}{b}} \right)$. Also, we set the initial budget to $w = 10$.

CARA utility. Following the definition in (14), in Fig. 2 we compute the longevity risk pooling δ for a CARA agent. In Fig. 2(a), we fix the fraction of balance sheet annuitized ψ_0 at 0.5, and set $\gamma \in [0.1, 0.6]$ and the insurance loading $\kappa \in [0, 0.3]$. In Fig. 2(b), instead, we fix κ at 0 and set $\gamma \in [0.1, 0.6]$ and $\psi_0 \in [0.5, 0.9]$. A few reference values are reported in tabulated format in Table 2.

The qualitative interpretation of the results obtained in Fig. 2(a) is unchanged with respect to the CRRA setting. Namely, the insurance loading decreases significantly the value of risk pooling while the level of risk aversion, measured by γ , increases this value. On the other hand, from Fig. 2(b), we can observe that in the case of a CARA utility, the longevity risk pooling δ does not depend on the fraction of balance sheet annuitized ψ_0 . This is due to the fact that, when using a CARA exponential utility, the equality in (7) can be simplified and shown to be independent of the pension income π .⁵ Thus, this entails that, for a CARA agent, the benefits of pooling do not change based on the preexisting amount of lifetime annual income. This is relatively intuitive as the risk aversion of a CARA utility function does not depend on the wealth level.

HARA utility. Following the definition in (15), in Fig. 3 we compute the longevity risk pooling δ for a HARA agent. In Fig. 3(a), we fix the fraction of balance sheet annuitized ψ_0 at 0.5 and set the risk aversion coefficient $\gamma \in [0.2, 1.4]$ and the insurance loading

⁵ We recall that ψ_0 is a function of the pension income π , the budget w and the immediate life annuity $ILA(x)$. However, in our experiments w and all other parameters of the model are kept constant, so we only let π change in order to determine the values of ψ_0 .

Table 3HARA $\gamma \in [0.2, 1.4]$, $\kappa \in [0, 0.3]$, $\psi_0 \in [0.5, 0.9]$, $a = 1$ or $a = 2$, $d = 1$.

ψ_0, κ		γ						
		0.2	0.4	0.6	0.8	1.0	1.2	1.4
$a = 1$	$\psi_0 = 0.5, \kappa = 0.0$	0.14	0.19	0.22	0.24	0.26	0.27	0.28
	$\psi_0 = 0.5, \kappa = 0.1$	0.03	0.08	0.10	0.12	0.14	0.15	0.16
	$\psi_0 = 0.5, \kappa = 0.2$	−0.06	−0.02	0.01	0.02	0.04	0.05	0.06
	$\psi_0 = 0.5, \kappa = 0.3$	−0.13	−0.1	−0.08	−0.06	−0.05	−0.04	−0.03
	$\psi_0 = 0.7, \kappa = 0.0$	0.11	0.15	0.18	0.20	0.22	0.23	0.24
	$\psi_0 = 0.9, \kappa = 0.0$	0.07	0.09	0.11	0.13	0.14	0.15	0.16
$a = 2$	$\psi_0 = 0.5, \kappa = 0.0$	0.14	0.2	0.24	0.27	0.29	0.31	0.32
	$\psi_0 = 0.5, \kappa = 0.1$	0.04	0.09	0.12	0.14	0.16	0.18	0.19
	$\psi_0 = 0.5, \kappa = 0.2$	−0.05	0.0	0.02	0.04	0.06	0.07	0.09
	$\psi_0 = 0.5, \kappa = 0.3$	−0.13	−0.09	−0.06	−0.04	−0.03	−0.01	0.0
	$\psi_0 = 0.7, \kappa = 0.0$	0.11	0.16	0.19	0.21	0.23	0.25	0.27
	$\psi_0 = 0.9, \kappa = 0.0$	0.07	0.1	0.12	0.13	0.15	0.16	0.17

Table 4SAHARA $\alpha \in [0.5, 5]$, $\beta = 1$ or $\beta = 2$, $\kappa \in [0, 0.3]$, $\psi_0 \in [0.5, 0.9]$.

ψ_0, κ		α				
		0.5	1	2	3	4
$\beta = 1$	$\psi_0 = 0.5, \kappa = 0.0$	0.22	0.30	0.41	0.49	0.57
	$\psi_0 = 0.5, \kappa = 0.1$	0.1	0.18	0.27	0.34	0.41
	$\psi_0 = 0.5, \kappa = 0.2$	0.0	0.07	0.16	0.22	0.29
	$\psi_0 = 0.5, \kappa = 0.3$	−0.08	−0.02	0.06	0.17	0.18
	$\psi_0 = 0.7, \kappa = 0.0$	0.18	0.25	0.34	0.41	0.46
	$\psi_0 = 0.9, \kappa = 0.0$	0.11	0.15	0.21	0.25	0.29
$\beta = 2$	$\psi_0 = 0.5, \kappa = 0.0$	0.19	0.26	0.36	0.42	0.49
	$\psi_0 = 0.5, \kappa = 0.1$	0.07	0.14	0.22	0.28	0.39
	$\psi_0 = 0.5, \kappa = 0.2$	−0.02	0.04	0.11	0.16	0.22
	$\psi_0 = 0.5, \kappa = 0.3$	−0.1	−0.1	0.02	0.07	0.12
	$\psi_0 = 0.7, \kappa = 0.0$	0.16	0.23	0.32	0.38	0.43
	$\psi_0 = 0.9, \kappa = 0.0$	0.11	0.15	0.21	0.25	0.28

$\kappa \in [0, 0.3]$. In Fig. 3(b), we fix $\kappa = 0$ and choose $\gamma \in [0.2, 1.4]$ and $\psi_0 \in [0.5, 0.9]$. Also, the coefficients a and d in (15) are set equal to 1. A few additional values are reported in tabulated format in Table 3, where we also consider the case when $a = 2$.

In the case of the HARA utility, the qualitative results are also consistent with those reported for the CRRA utility function.

SAHARA utility. Finally, following the definition in (16), in Fig. 4 we compute the longevity risk pooling δ for a SAHARA agent.

In Fig. 4(a), we fix the fraction of balance sheet annuitized ψ_0 at 0.5 and set $\alpha \in [0.5, 4]$ and the insurance loading $\kappa \in [0, 0.3]$. In Fig. 4(b), we fix the insurance loading to κ at 0 and choose $\alpha \in [0.5, 4]$ and $\psi_0 \in [0.5, 0.9]$. Furthermore, the coefficient β is set as equal to 1. A few additional numerical experiments in the case of a SAHARA utility function are reported in tabulated format in Table 4, where we also consider the case when $\beta = 2$.

5.2. Sensitivity to the choice of utility function

In this last section, we illustrate how we can use our parametrization of the utility function via a probability distribution in order to study the sensitivity of the risk pooling coefficient to the choice of the utility function.

For instance, let us consider a family of utility functions parametrized by the following three-parameter generalized gamma distribution (GG) that has a density usually defined for $\alpha > 0$, $\beta > 0$, $\theta > 0$ as

$$f_{GG}(y; \alpha, \beta, \theta) = \frac{\beta}{\Gamma(\alpha) \cdot \theta} \left(\frac{y}{\theta}\right)^{\alpha\beta-1} e^{-\left(\frac{y}{\theta}\right)^\beta}, \quad (17)$$

with $y > 0$. For reasons of numerical stability, however, in our experiments we use an alternative parametrization (see Lawless

Table 5General utility: $\mu \in [0, 1]$, $\sigma \in [0.5, 0.75, 1]$, $\lambda \in [-0.4, 0.4]$.

μ, σ	λ				
	−0.4	−0.2	0.0	0.2	0.4
$\mu = 0.0, \sigma = 1.0$	0.32	0.32	0.33	0.33	0.33
$\mu = 0.0, \sigma = 0.75$	0.36	0.37	0.37	0.38	0.38
$\mu = 0.0, \sigma = 0.5$	0.42	0.43	0.44	0.45	0.46
$\mu = 1.0, \sigma = 1.0$	0.34	0.33	0.33	0.32	0.32
$\mu = 1.0, \sigma = 0.75$	0.39	0.38	0.37	0.36	0.35
$\mu = 1.0, \sigma = 0.5$	0.47	0.46	0.44	0.43	0.41

(2011)) with parameters $-\infty < \mu < \infty$, $\sigma > 0$ and λ (arbitrary):

$$f_{GG}(y; \mu, \sigma, \lambda) = \begin{cases} \frac{|\lambda|}{\sigma y \cdot \Gamma\left(\frac{1}{\lambda^2}\right)} \exp\left(\frac{\lambda \frac{\log(y)-\mu}{\sigma} + \log\left(\frac{1}{\lambda^2}\right) - e^{\lambda \frac{\log(y)-\mu}{\sigma}}}{\lambda^2}\right) & \text{if } \lambda \neq 0, \\ \frac{1}{\sigma y \sqrt{2\pi}} \exp\left(-\frac{1}{2} \left(\frac{\log(y)-\mu}{\sigma}\right)^2\right) & \text{if } \lambda = 0. \end{cases} \quad (18)$$

Notice that the generalized gamma is a limiting case of the generalized beta introduced in Section 4. More precisely,


$$f_{GG}(y; \alpha, \beta, \theta) = \lim_{q \rightarrow \infty} f_{GB}(y; \alpha, b = \theta \beta q^{\frac{1}{\alpha}}, \zeta = 0, \beta, q).$$

Also, it includes as special cases the gamma distribution, the Weibull distribution and the lognormal distribution.⁶


In Fig. 5, we compute the longevity risk pooling δ with respect to μ, σ and λ . As already pointed out, the generalized gamma distribution converges to a lognormal distribution when $\lambda = 0$. If, in addition, $\mu = 0$, our parametrized utility function boils down to a CRRA utility with $\gamma = \frac{1}{\sigma}$, and logarithmic utility for $\sigma = 1$. Figs. 5(a) and 5(b) both illustrate how to deviate away from the CRRA utility function by varying a parameter, here λ . This makes it possible to check the robustness with respect to the choice of the utility function, not only by considering a few examples of utility functions that are used in the literature but by moving away from the CRRA utility smoothly.

From Table 5, we can see that, using our general approach, when $\mu = 0$, $\sigma = 0.5$, $\lambda = 0$, corresponding to a CRRA utility with $\gamma = 2$, we obtain $\delta = 0.44$. This compares favorably with the “direct” method, which does not rely on the approximation

⁶ From the definition in (18), it is easy to see that, when $\lambda = 0$, the generalized gamma converges to a Lognormal(μ, σ).




(a) Fraction of balance sheet annuitized $\psi_0 = 0.5$, HARA $\gamma \in [0.2, 1.4]$, insurance loading $\kappa \in [0, 0.3]$.




(b) Insurance loading $\kappa = 0$, fraction of balance sheet annuitized $\psi_0 \in [0.5, 0.9]$, HARA $\gamma \in [0.2, 1.4]$.

Fig. 3. Longevity risk pooling δ for a 65-year-old HARA agent with initial budget $w = 10$. Coefficients a and d in (15) are set as 1.




(a) Fraction of balance sheet annuitized $\psi_0 = 0.5$, SAHARA $\alpha \in [0.5, 4]$, insurance loading $\kappa \in [0, 0.3]$.




(b) Insurance loading $\kappa = 0$, fraction of balance sheet annuitized $\psi_0 \in [0.5, 0.9]$, SAHARA $\alpha \in [0.5, 4]$.

Fig. 4. Longevity risk pooling δ for a 65-year-old SAHARA agent with initial budget $w = 10$. The coefficient β in (16) is set as equal to 1.



(a)



(b)

Fig. 5. Longevity risk pooling δ for a 65-year-old agent with initial budget $w = 10$ and utility function parametrized by a generalized gamma distribution with parameters μ , σ , λ . Fraction of balance sheet annuitized $\psi_0 = 0.5$, insurance loading $\kappa = 0$ and $\lambda \in [-0.4, 0.4]$.

of the inverse of the first derivative of the utility function (but rather employs its explicit form, when known) and also yields $\delta = 0.44$.⁷ In addition, when $\mu = 0$, $\sigma = 1$, $\lambda = 0$, corresponding to a logarithmic utility, we obtain $\delta = 0.33$. In Milevsky and Huang (2018), by using a CRRA utility with $\gamma = 1.01$ (a close approximation of the logarithmic utility) and budget $w = 10\,000$, the authors obtain $\delta = 0.326$.

Fig. 5 illustrates that the results are mildly sensitive to the choice of the utility function. Throughout our numerical study, besides few remarks concerning specifically the case of a CARA exponential utility, we did not find any significant (qualitative) behavior change by using a more general utility function than the CRRA utility function used by Milevsky and Huang (2018). Thus, we can conclude that their results related to the optimal demand for annuities are overall robust to the assumption of the CRRA utility function.

6. Conclusions

This paper makes several contributions. A first contribution is to provide new families of concave increasing utility functions that are of interest by themselves and can be used to generalize results proven in the literature for specific utility functions. A second contribution is to show that the results of Milevsky and Huang (2018) about the welfare benefits of purchasing annuities are robust to the assumption of a CRRA utility function, which is a key assumption used in the derivation of their closed-form expressions for the longevity risk pooling coefficient. Of course, specific choices of utility functions may lead to specific sensitivities, e.g., in the case of a CARA utility function, the longevity pooling coefficient does not depend on the percentage of annuitized wealth due to the absence of sensitivity of the risk aversion to the wealth level.

Our method is flexible and we have presented the simplest case in which the only investment opportunity for the retiree is the risk-free asset. Generalizing the results to the case of the optimal demand for annuities in the presence of a financial market, thus including additional investment opportunities, is a natural extension that we leave for further research.

Another natural extension is to go beyond the expected utility theory and study the annuity demand in a non-expected utility setting. On the one hand, we do not expect results to be necessarily different if essentially the behavioral theory is consistent with first-order stochastic dominance. In this case, Bernard et al. (2015) have indeed proved that the optimal portfolio of a non-expected utility investor always coincides with the one of an expected utility investor with concave increasing utility. However, their setting is based on a one-period investment, which is still consistent with time-additive preferences in which the agent's attitude to risk are described similarly to attitude to intertemporal substitution. An interesting but challenging avenue for future research is to investigate the demand for annuities assuming more general models of preferences such as Davidoff et al. (2005) and Cannon and Tonks (2008) (e.g., Epstein–Zin preferences), but this is beyond the scope of our paper.

Appendix. Inverse functions via neural networks

In this appendix, we show how to approximate the inverse of a function over a finite range by means of neural networks.

Let us consider a sufficiently smooth function f defined on an interval $[y_1, y_2]$. In a nutshell, the algorithm works simply as follows. First, start by sampling $\mathbf{x} := \{x_i\}_{i=1}^p$ from a $\text{Uniform}(y_1, y_2)$




Fig. 6. Inverse of the function $f(x) = -\log(x)$ on the interval $[0, 4]$.

and evaluate $f(\mathbf{x})$. Then, approximate $f^{-1}(f(\mathbf{x}))$ by $h(f(\mathbf{x}); \mathbf{w})$, where h is a neural network with weights $\mathbf{w}$, and minimize the L^2 -norm $\|h(f(\mathbf{x}); \mathbf{w}) - \mathbf{x}\|_2$ by using stochastic gradient descent (more specifically, we use the Adam optimizer (Kingma and Ba, 2015), which is an extension to stochastic gradient descent). then repeat the experiment until convergence.

Below, we provide a pseudo-code for the algorithm. Our implementation was done in Python with TensorFlow library.

Algorithm 1: Inverse function via neural nets.

Inputs: target function f defined on a finite set $[y_1, y_2]$; batch size p ; penalty function β ; hyper-parameters for the neural networks architecture Θ_m .

Require: random initialization of weights $\mathbf{w}$.

while not converged **do**

sample $\{x_i\}_{i=1}^p = \mathbf{x} \sim U(y_1, y_2)$;

evaluate $f(\mathbf{x})$;

approximate $f^{-1}(f(\mathbf{x}))$ by $h(f(\mathbf{x}); \mathbf{w})$ via a neural network;

evaluate $\phi(f; \mathbf{w}) = \|h(f(\mathbf{x}); \mathbf{w}) - \mathbf{x}\|_2$;

$\mathbf{w} \leftarrow \text{Adam}(\phi(f; \mathbf{w}))$;

end while

Now we show a couple of examples to verify the accuracy of the algorithm. Firstly, we consider the quite simple case of $f(x) = -\log(x)$ on the interval $[0, 4]$. Of course, the exact inverse is given by $f^{-1}(x) = e^{-x}$. We train the algorithm for 20000 iterations⁸ with a batch size of 2^7 . As network architecture, we use a simple feedforward network with one hidden layer and 64 hidden neurons. As we can see in Fig. 6, the two curves are almost indistinguishable. In this case, we achieved a training mean squared error (MSE) of 2.1×10^{-7} and a test MSE of 3×10^{-6} .

Secondly, we consider the more complicated case of $f(x) = \sin(x)$ on the interval $[-\frac{\pi}{2}, \frac{\pi}{2}]$. The inverse here is given by $f^{-1}(x) = \arcsin(x)$. As we can see from Fig. 7, the algorithm still achieves quite a good level of accuracy. The training MSE is 2.4×10^{-7} , while the test MSE is 7×10^{-4} .

As a final note, we remark that the objective of the algorithm is not to approximate the (known) inverse function f^{-1} per se.

⁷ Note that this method can only be used when $\mu = 0$, $\lambda = 0$, as then we retrieve the special case of the CRRA.

⁸ Since we are sampling uniformly in the interval $[y_1, y_2]$, the larger such interval the higher the number of iterations we would need in order for the algorithm to be trained adequately everywhere. Also, notice that, although the function is approximated only on the interval $[y_1, y_2]$, there is no a priori impediment to evaluating it outside of this range (by extrapolation), if necessary.




Fig. 7. Inverse of the function $f(x) = \sin(x)$ on the interval $[-\frac{\pi}{2}, \frac{\pi}{2}]$.

Of course, while training we do not assume any knowledge on f^{-1} . Rather, the algorithm looks for the function (i.e., the neural network) h such that $h(f(x)) = x$.

References

- Alexander, C., Cordeiro, G.M., Ortega, E.M., Sarabia, J.M., 2012. Generalized beta-generated distributions. *Comput. Statist. Data Anal.* 56 (6), 1880–1897.
- Bernard, C., Chen, J.S., Vanduffel, S., 2015. Rationalizing investors' choices. *J. Math. Econom.* 59, 10–23.
- Brockett, P.L., Golden, L.L., 1987. A class of utility functions containing all the common utility functions. *Manage. Sci.* 33 (8), 955–964.
- Cannon, E., Tonks, I., 2008. *Annuity Markets*. Oxford University Press.
- Charupat, N., Huang, H., Milevsky, M.A., 2012. *Strategic Financial Planning Over the Lifecycle: A Conceptual Approach to Personal Risk Management*. Cambridge University Press.
- Chen, A., Pelsser, A., Vellekoop, M., 2011. Modeling non-monotone risk aversion using SAHARA utility functions. *J. Econom. Theory* 146 (5), 2075–2092.
- Davidoff, T., Brown, J.R., Diamond, P.A., 2005. Annuities and individual welfare. *Amer. Econ. Rev.* 95 (5), 1573–1590.
- Eugene, N., Lee, C., Famoye, F., 2002. Beta-normal distribution and its applications. *Comm. Statist. Theory methods* 31 (4), 497–512.
- Hanewald, K., Piggott, J., Sherris, M., 2013. Individual post-retirement longevity risk management under systematic mortality risk. *Insurance Math. Econom.* 52 (1), 87–97.
- Johnson, N.L., 1949. Systems of frequency curves generated by methods of translation. *Biometrika* 36 (1/2), 149–176.
- Kingma, D.P., Ba, J., 2015. Adam: A method for stochastic optimization. In: *International Conference on Learning Representations, ICLR 2015*.
- Lawless, J.F., 2011. *Statistical Models and Methods for Lifetime Data*, vol. 362. John Wiley & Sons.
- Lee, C., Famoye, F., Alzaatreh, A.Y., 2013. Methods for generating families of univariate continuous distributions in the recent decades. *Wiley Interdiscip. Rev. Comput. Stat.* 5 (3), 219–238.
- Leung, S.F., 2007. The existence, uniqueness, and optimality of the terminal wealth depletion time in life-cycle models of saving under uncertain lifetime and borrowing constraint. *J. Econom. Theory* 134 (1), 470–493.
- McDonald, J.B., 2008. Some generalized functions for the size distribution of income. In: *Modeling Income Distributions and Lorenz Curves*. Springer, pp. 37–55.
- McDonald, J.B., Xu, Y.J., 1995. A generalization of the beta distribution with applications. *J. Econometrics* 66 (1–2), 133–152.
- Milevsky, M.A., 2018. Swimming with wealthy sharks: longevity, volatility and the value of risk pooling. *J. Pension Econ. Finance* 1–30.
- Milevsky, M.A., Huang, H., 2018. The utility value of longevity risk pooling: Analytic insights. *N. Am. Actuar. J.* 22 (4), 574–590.
- Mitchell, O.S., 2002. Developments in decumulation: The role of annuity products in financing retirement. In: *Ageing, Financial Markets and Monetary Policy*. Springer, pp. 97–125.
- Níguez, T.-M., Paya, I., Peel, D., Perote, J., 2019. Flexible distribution functions, higher-order preferences and optimal portfolio allocation. *Quant. Finance* 19 (4), 699–703.
- Yaari, M.E., 1965. Uncertain lifetime, life insurance, and the theory of the consumer. *Rev. Econom. Stud.* 32 (2), 137–150.
- Ye, Y., Oluyede, B.O., Pararai, M., 2012. Weighted generalized beta distribution of the second kind and related distributions. *J. Stat. Econom. Methods* 1 (1), 13–31.