NBER WORKING PAPER SERIES

EVALUATION OF MEDICAL TECHNOLOGIES WITH UNCERTAIN BENEFITS

Darius N. Lakdawalla
Charles E. Phelps

Working Paper 26058
http://www.nber.org/papers/w26058

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2019

The authors thank Mike Drummond, Alan Garber, Lou Garrison, Jon Gruber, Emmett Keeler, Mark Pauly, Julian Reif, and seminar participants at the University of Washington for helpful comments, and we thank Hanh Nguyen for exceptional assistance with the manuscript. All errors are our own. DNL discloses that he owns equity in Precision Medicine Group, which provides consulting services to firms in the healthcare and pharmaceutical industries. DNL is grateful for funding from the National Institute on Aging (1R01AG062277-01). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w26058.ack

Evaluation of Medical Technologies with Uncertain Benefits
Darius N. Lakdawalla and Charles E. Phelps
NBER Working Paper No. 26058
July 2019, Revised May 2020
JEL No. I1,I13,I18

ABSTRACT

Cost-effectiveness analysis (CEA), despite its known limitations, continues as the primary method used for health technology assessment (HTA) both officially (UK, Australia and Canada) and less formally elsewhere. Standard CEA models compare incremental cost increases to incremental average gains in health, commonly expressed in Quality-Adjusted Life Years (QALYs). Our research generalizes earlier CEA models in several ways. First, we introduce risk aversion in Quality of Life (QoL), which affects willingness to pay (WTP) for health care, leading to WTP thresholds that rise with illness severity. Ignoring risk aversion in QoL over-values treatments for minor illnesses and under-values treatments for highly severe illnesses, perhaps by an order of magnitude. We call our generalized WTP threshold the Risk-Aversion and Severity-Adjusted WTP (RASA-WTP). Unlike traditional CEA analyses, which discriminate against persons with disabilities, our analysis implies that the marginal value of improving QoL rises for disabled individuals. Our model can also value the uncertain benefits of medical interventions by employing well-established analytic methods from finance. We develop a certainty-equivalent quality of life measure that we call the Risk-Adjusted QALY (RA-QALY), which accounts for consumer preferences over risky health outcomes. Finally, we show that traditional QALYs no longer serve as a single index of health, when consumers are risk-averse. To address this problem, we derive a generalized single-index of health outcomes—the Generalized Risk-Adjusted QALY (GRA-QALY). The GRA-QALY reinstates the equivalence between health gains from quality of life and gains from life extension, even in the presence of risk-aversion and treatment outcome uncertainty. Earlier models of CEA that abstract from risk-aversion nest as special cases of our more general model. We discuss new data necessary to implement our model and standard analytic methods by which the necessary parameters can be obtained.

Darius N. Lakdawalla
University of Southern California
635 Downey Way, VPD 414-K
Schaeffer Center for Health Policy and Economics
Los Angeles, CA 90089-7273
and NBER
dlakdawa@usc.edu

Charles E. Phelps
University of Rochester
30250 S Highway One
Gualala, CA 95445
charles.phelps@rochester.edu

Abstract (287 words)

Cost-effectiveness analysis (CEA), despite its known limitations, continues as the primary method used for health technology assessment (HTA) both officially (UK, Australia and Canada) and less formally elsewhere. Standard CEA models compare incremental cost increases to incremental average gains in health, commonly expressed in Quality-Adjusted Life Years (QALYs). Our research generalizes earlier CEA models in several ways. First, we introduce risk aversion in Quality of Life (QoL), which affects willingness to pay (WTP) for health care, leading to WTP thresholds that rise with illness severity. Ignoring risk aversion in QoL over-values treatments for minor illnesses and under-values treatments for highly severe illnesses, perhaps by an order of magnitude. We call our generalized WTP threshold the Risk-Aversion and Severity-Adjusted WTP (RASA-WTP). Unlike traditional CEA analyses, which discriminate against persons with disabilities, our analysis implies that the marginal value of improving QoL rises for disabled individuals. Our model can also value the uncertain benefits of medical interventions by employing well-established analytic methods from finance. We develop a certainty-equivalent quality of life measure that we call the Risk-Adjusted QALY (RA-QALY), which accounts for consumer preferences over risky health outcomes. Finally, we show that traditional QALYs no longer serve as a single index of health, when consumers are risk-averse. To address this problem, we derive a generalized single-index of health outcomes—the Generalized Risk-Adjusted QALY (GRA-QALY). The GRA-QALY reinstates the equivalence between health gains from quality of life and gains from life extension, even in the presence of risk-aversion and treatment outcome uncertainty. Earlier models of CEA that abstract from risk-aversion nest as special cases of our more general model. We discuss new data necessary to implement our model and standard analytic methods by which the necessary parameters can be obtained.

1. INTRODUCTION

For decades, economists have studied how consumers trade off mortality risk and money (Murphy and Topel, 2006; Rosen, 1988). These analyses guide policymakers in addressing incomplete markets for life-extension investments, e.g., by determining optimal investments in safer transportation infrastructure or optimal environmental standards. In the healthcare context, economic theories of mortality risk-reduction help determine how much third-party payers and social planners should spend on extending life. However, they do not determine how to allocate resources among competing interventions. The related theory of cost-effectiveness, of roughly the same vintage as the theory of mortality risk-reduction, addresses these gaps (Weinstein and Stason, 1977).

Cost-effectiveness has long been used in Britain, Canada, and Australia to determine coverage and reimbursement of new medical technologies by health insurers and to evaluate medical technologies in the US and elsewhere. The Institute for Clinical and Economic Review (ICER), a US nonprofit organization, routinely conducts and releases cost-effectiveness studies for use by American healthcare payers and providers.[1] Nearly 60% of US payers have consulted cost-effectiveness analyses in their price negotiations or reimbursement decisions (Lising et al., 2016). In 2018, a large Pharmacy Benefit Manager (CVS Caremark) proposed a health plan that limits payment for prescription drugs to a maximum of $100,000 per estimated Quality Adjusted Life Year (QALY) gained.[2] Growing reliance on and interest in cost-effectiveness analysis raises the importance of assuring that CEA methods are robust.

Innovations in our Model. Standard CEA models assume that consumers are risk-neutral in health (Garber and Phelps, 1997), which simplifies analyses but risks misrepresenting true consumer preferences. We introduce risk aversion in Quality of Life (QoL), a generalization with significant implications for the conduct of CEA. First, this approach shows that a uniform CEA Willingness to Pay (WTP) threshold should be replaced by thresholds that grow as disease severity increases. Risk-averse consumers derive greater value from health improvements when they face bleaker health prospects. Therefore, ignoring risk aversion over-values treatments for minor illnesses and under-values treatments for highly serious ones. Our empirical calibrations

[1] https://icer-review.org

[2] https://www.forbes.com/sites/joshuacohen/2018/09/20/will-cvs-caremark-make-icer-the-american-nice/#1cb23a8b6173 last visited May 20, 2020.

suggest that mild illness treatment might be over-valued by a factor of two or three, while severe illness treatment could be under-valued by up to an order of magnitude. Importantly, this approach also helps mitigate the discrimination against disabled persons that is embedded in current CEA methods. The Affordable Care Act forbids use of CEA methods in Medicare and PCORI that so discriminate.

Second, we incorporate random health endowments in future periods, thereby allowing both *ex ante* perspectives for formulating health insurance coverage decisions and *ex post* analyses for patients with realized medical disorders to select specific therapies.

Third, our approach expands the valuation of medical care to account for the variability in treatment outcomes, just as the world of finance has understood the importance of variability for nearly three quarters of a century. We connect measures of uncertainty in treatment outcomes to standard parameters measuring risk attitudes to quantify the importance of uncertainty.

Fourth, introducing risk aversion in QoL breaks the traditional equivalence between health gains in life years (LY) and QoL, central to current models. Without that equivalence, researchers can no longer conclude that "a QALY is a QALY." To rebuild a tractable framework for decision making, we develop a generalized single index of value (the Generalized Risk-Adjusted Quality Adjusted Life Year, or GRA-QALY) that restores the equivalence between health gains in LY and QoL.

Uncertainty is increasingly salient with the rise of personalized and targeted therapies and the value of "companion diagnostics" that help predict when a given technology will or will not assist any specific patient. Technologies (including diagnostic tests) that reduce variances of health outcomes generate value to risk-averse consumers, even if average outcomes remain fixed (Lakdawalla et al., 2017). Further, technologies that increase variance may still have incremental value to risk-averse consumers if they sufficiently increase the positive skewness in distributions of variable health outcomes (Eeckhoudt et al., 1995)— e.g., the "value of hope" exhibited by cancer patients (Lakdawalla et al., 2012).

A recent economic task force report called for methods to monetize the values of risk reduction and of hope, and the influence of illness severity in value assessment (Garrison et al., 2020; Lakdawalla et al., 2018). Our approach provides direct solutions to all three of these issues using standard expected-utility maximization methods.

The Relevance of Uncertainty. A wide variety of diseases exhibit variable treatment responses (Alatorre et al., 2011), including highly prevalent diseases like depression (Carter et al., 2012), rheumatoid arthritis (Goetz et al., 2011), diabetes (Cantrell et al., 2010), and cancer (Yu and Cui, 2018). While continued progress in diagnostics might eventually improve the ability to forecast which patients will respond to particular treatments, there remains a considerable amount of variability in treatment response that is unknown before treatments begin. Until accurate forecasting of treatment success is ubiquitous, value assessment methods should incorporate practical strategies for quantifying effects of uncertainty.

Multi-dimensional Value. Health outcomes may have multiple dimensions of value. For example, the widely used European Quality of Life (EQ-5) measure has five dimensions. The Health Utility Index 3 (HUI3) captures 8 dimensions: vision, hearing, speech, ambulation, dexterity, emotion, cognition, and pain (Maxwell et al., 2016). Judging from the frequency of Direct to Consumer advertisements on TV and elsewhere, one might readily add dimensions such as "physical appearance" and "sexual function." Whatever the total number of elements, our model presumes a scalar measure of QoL, which could (for example) use multi-criteria decision analysis (MCDA) methods to combine vectors of quality measures into scalars, along with associated variances, skewness, and other higher-order moments from the distributions of these combined measures.

Limitations. Our approach assesses the effects of risk aversion in QoL and uncertainty in treatment outcomes for a representative utility-maximizing consumer. Since our work generalizes traditional CEA, it (like traditional CEA) cannot resolve issues arising in some public policy decisions that involve inter-personal comparisons (equity and fairness) (Asaria et al., 2015), nor can it address issues such as externalities arising from scientific spillovers, contagious disease spread or herd immunity against infectious diseases. Nevertheless, we believe that a full understanding of the effect of risk aversion in QoL and uncertainty in the benefits of medical interventions can improve decision-making at all levels, just as traditional CEA has informed numerous private and public policy decisions.

2. THEORETICAL DEVELOPMENT

2.1. Reviewing the Foundations of Cost-Effectiveness Analysis

We begin with the familiar theoretical framework for cost-effectiveness analysis of Garber and Phelps (1997), hereafter GP. Consider a representative consumer choosing consumption and medical spending over two periods.[3] Define Y as exogenous income that arrives each period. Define M_i as composite medical spending in period i, and $C_i \equiv Y - M_i$ as non-medical consumption in period i. Define $H_i \in [0,1]$ as QoL in period i. Period i utility is given by $U(C_i)H_i$, and the probability of survival to period 1 is p_1. Expected utility is $U(C_0)H_0 + p_1 U(C_1)H_1$. Note the key assumption that utility is linear in QoL.

In GP, the optimal cost-effectiveness threshold is given by:[4]

$$K \equiv \frac{U(C_1)}{U'(C_0)}\left[\frac{1}{H_0}\right] \quad (1)$$

K reflects the value of future QoL gains paid for today, measured as consumption willingly foregone in exchange for one unit of QoL. GP sets baseline QoL to unity, in the sense that $H_0 \equiv 1$. Cost-effectiveness is measured from the perspective of a consumer with "perfect" or "excellent" baseline health. We start by following this convention and generalize it later (in Section 3.3.2) to allow for the possibility that $H_0 < 1$. As things stand, K is the WTP for one QALY.

H_1 reflects the health status of patients when they need treatment. We assume that H_1 incorporates the effects of the illness in question and, possibly, other unrelated disabilities or illnesses. For example, physically disabled patients with diabetes would have lower values of H_1 than patients with identical diabetes status but no physical disabilities.

If individuals choose medical spending optimally, then K represents the cost-effectiveness ratio that leads to the first-best allocation of medical consumption. If instead the individual faces an arbitrarily fixed budget – e.g., fixed by a government payer – then the resulting K produces suboptimal allocations, but they are second-best optimal conditional on the chosen budgetary level (Phelps, 2019b).

---

[3] GP use three-periods, but in our setting, the final period adds no insight.

[4] GP omit H_0 in their formulation because they set $H_0 = 1$. They also use $U(Y)$ instead of $U(C_1)$ in the numerator. When $U(Y) = \ln(Y)$, K differs proportionately from our definition by the budget share of medical care.

For technologies that add Δp_1 in survival probability and ΔH_1 in QoL, QALY gains are defined as $\Delta QALY \equiv \Delta p_1 H_1 + p_1 \Delta H_1$. GP prove the legitimacy of guiding decisions by comparing costs per quality-adjusted life-years (QALYs) to K. In addition, their framework implies that the value of a technology equals the value of QALYs gained through it. This is a non-trivial result, because health has two dimensions—QoL and LY – and they distill its value into one dimension, namely the QALY improvement.

The QALY as a single index of health improvement rests crucially on the assumption that marginal utility equals average utility in both QoL and LY. Elaborating, consider a representative consumer's willingness to increase annual health insurance premiums by the expected cost of a new medical technology. The marginal utility of QoL is $U(C_1)$, and WTP for this QoL gain in terms of baseline consumption (C_0) units is $K = \frac{U(C_1)}{U'(C_0)}$. Therefore, the period zero WTP for an expected QoL gain of $p_1 \Delta H_1$ is $K p_1 \Delta H_1$. Since $p_1 \Delta H_1$ is the QALY gain associated with ΔH_1 additional QoL units, it follows that K yields the value per QALY gained via QoL improvement. At the same time, the period zero WTP for a marginal gain in survival probability, Δp_1, is $\frac{U(C_1)H_1}{U'(C_0)} \Delta p_1 = K \Delta p_1 H_1$. Since $\Delta p_1 H_1$ is the QALY gain associated with this gain in LY, it follows that K measures value per QALY gained *regardless* of whether it is gained via life-extension or QoL improvement.

This implication holds because, under utility that is linear in QoL, the marginal rate of substitution between QoL and survival gains happens to be $\frac{U(C_1)H_1}{U(C_1)} = H_1$. To appreciate the importance of this condition, recall that the QALY gain is defined as $p_1 \Delta_H + \Delta p_1 H_1$. The two terms in this QALY formula compactly express the value of a two-dimensional gain in health in the common units of QoL improvement. The first term is literally the expected gain in QoL units. The second is the gain in survival multiplied by the marginal rate of substitution; this product equals the expected QoL units the consumer will give up in return for Δp_1 more survival probability.

This result produces strong decision-analytic implications. Consider a technology that increases survival probability by Δp_1, QoL by ΔH_1, and incremental costs by ΔC_0. Based on the arguments above, the net value of such a technology is $K \Delta p_1 H_1 + K p_1 \Delta H_1 - \Delta C_0$. It improves consumer welfare if and only if:

$$\frac{\Delta C_0}{\Delta p_1 H_1 + p_1 \Delta H_1} \leq K \tag{2}$$

This proves—in the risk-neutral world of GP—that the cost per QALY gained, regardless of how the QALY was gained, represents a sufficient index of welfare when compared to the threshold K.

This "QALY is a QALY" conclusion serves as a fundamental result in CEA methodology. However, it rests somewhat precariously on the assumption that expected utility is linear—not concave—in QoL, or equivalently that consumers are risk-neutral in QoL. While linearity of expected utility in QoL provides strong and helpful implications, it also places the cost-effectiveness literature out of step with the foundational literature on health production and health human capital within health economics (Ehrlich and Chuma, 1990; Galama and Kapteyn, 2011; Grossman, 1972; Muurinen, 1982). In this respect, the cost-effectiveness literature on valuing health investment is incompatible with the economics literature that explains those health investment decisions. Other well-known studies assuming concavity in health-related QoL include Arrow's (1976) work on health insurance benefit design, Cutler et al's (1998) study of medical care price indices, and Hall and Jones (2007) on the macroeconomic determinants of medical spending. From an empirical standpoint, linearity of utility in QoL is also hard to reconcile with survey-based studies finding preferences for treating more severe diseases, holding cost-effectiveness fixed (Green and Gerard, 2009; Linley and Hughes, 2013; Nord et al., 1995).

Figure 1: Indifference curves for constant and diminishing marginal utility of QoL.

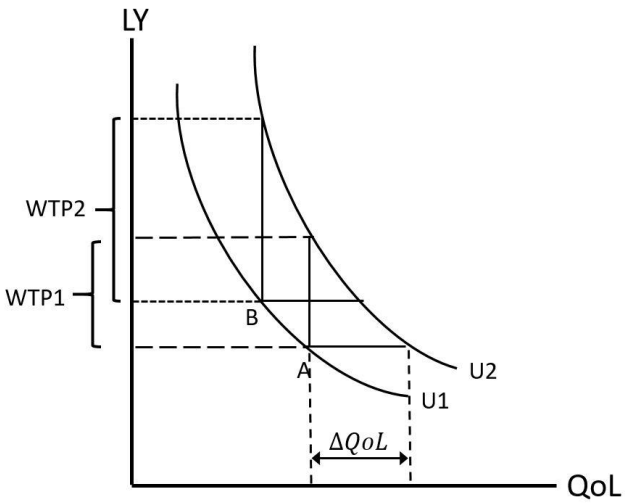

Figure 1a

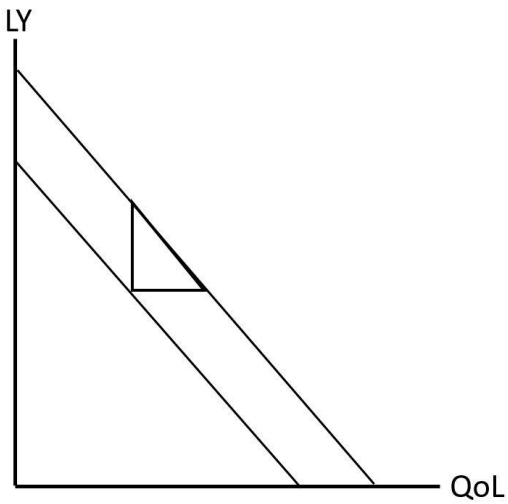

Figure 1b

Figure 1 illustrates the challenges that linear QoL utility poses to intuition. In Figure 1a, utility is concave in QoL, while in Figure 1b, it is linear. Starting in Figure 1a, from point A on indifference curve U_1, a fixed health gain of ΔQoL puts the consumer on indifference curve U_2. The vertical distance WTP_1— the compensating variation—returns the consumer to U_1 and measures (in LY) the value of this QoL gain. The same exercise conducted at point B (starting at a lower QoL) gives WTP_2. Because of diminishing marginal utility of QoL, $WTP_2 > WTP_1$ for the same ΔQoL.

In contrast, Figure 1b illustrates preferences when utility is linear in both LY and QoL, as GP assume. In this case, LY and QoL are perfect substitutes (hence the conclusion that "a QALY is a QALY"), and thus indifference curves are straight lines. Thus, WTP for improvement in QoL is the same for all values of QoL.

The familiar curvature in Figure 1a reflects the preference for variety that microeconomists typically assume. Equating apples and oranges, consumers endowed with 10 oranges and 1 apple are more eager to trade an orange for an apple, than when endowed with a balanced bundle of fruit.

The preference for variety in health bundles is similarly intuitive. Imagine Consumer A, with nearly perfect QoL, but with only 6 months left to live. Consumer B has 5 years to live, but at QoL level 0.1. The linear indifference curves of GP imply that these two consumers have identical willingness to trade away longevity in exchange for QoL. However, introspection

suggests that Consumer B—blessed with significantly greater longevity but hampered by very low QoL—will have much greater interest in making such a trade. The survey-based evidence cited earlier suggests that most real-world respondents agree, in the sense that health improvements are viewed as more valuable for the severely ill (Green and Gerard, 2009; Linley and Hughes, 2013; Nord et al., 1995).

In light of our discussion about concavity in QoL, readers might wonder why we still maintain that expected utility is linear in survival probabilities. Our reasoning corresponds to that of the QoL case: The vast majority of the theoretical literature in health economics, and indeed economics more generally, assumes that expected utility is additively separable across periods of time and thus linear in survival probability. This applies to the health human capital literature spawned by Grossman (1972), the CEA analysis of GP (Garber and Phelps, 1997), theoretical cost-effectiveness papers by Meltzer (Meltzer, 1997; Meltzer and Smith, 2011), and the value of statistical life literature following Rosen (1981, 1988).

Formal justification for additive time-separability in economics generally traces back to Koopmans (1972), who shows how additive time separability follows from two relatively intuitive axioms.[5] While the intellectual history is useful and interesting, it is not our goal to mount a renewed defense of additively separable utility over time. We merely follow the bulk of the economics literature in this respect.[6] In any event, our key observations about the fragility of QALYs would carry through even if utility were concave in survival, as long as the concavity in QoL differed meaningfully from concavity for survival.

Concavity of utility in QoL undermines the traditional equivalence between survival gains and QoL gains, and it simultaneously undermines the rationale for use of the QALY as a single index of health improvements. We will propose a solution and use it to build a framework that properly values health gains when individuals are risk averse in QoL, and when treatment benefits are uncertain.

---

[5] Bleichrodt et al. (2008) provide intuitive explanations for the two Koopmans axioms: 1) the marginal rate of substitution between today and tomorrow is unaffected by consumption levels on the day after tomorrow and beyond; and 2) if consumers prefer the consumption sequence x to the sequence y, they will harbor the same preferences if those sequences started tomorrow and we hold their consumption level today fixed.

[6] For analysis departing from this tradition, see Córdoba and Ripoll (2016).

2.2. Incorporating Risk-Aversion and Uncertainty

GP abstracted from health risk in three ways. We relax all three of these restrictions and explore their implications for value assessment.

First, we first define utility as $V(C,H) = U(C)W(H)$ and we allow W to exhibit concavity and thus risk-aversion. While we could use an even more general form of utility, additional interaction terms between C and H bring little insight into our central issue and further complicate presentation of our model.

Second, to introduce uncertainty over untreated health state, we assume that consumers are either sick (with probability ϕ) or well (with probability $(1-\phi)$) in period 1, with corresponding QoL levels of H_{1s} and H_{1w}, respectively. This allows consideration of value from an *ex ante* perspective, as when considering health insurance benefit structure. One can return to the deterministic valuation framework of traditional CEA modeling simply by setting $\phi = 1$, and a fully deterministic framework by setting $p_1 = 1$.

Third, we introduce uncertainty in outcomes of medical treatments and consider the expected utility of QoL, $E[W(H)]$. This is analogous to introducing uncertainty in financial transactions, a common practice in economic evaluations involving risky assets or income streams. We assume non-medical consumption is equal across the healthy and sick states in period 1, to avoid randomness in $U(C_1)$.[7]

As we will show, introduction of risk aversion in QoL has two distinct effects on Health Technology Assessment (HTA). First, it affects the WTP for health gains, both by introducing diminishing marginal utility in health and by highlighting that marginal values of health gains rise with illness severity and/or permanent disability. Second, it provides a way to incorporate variability in treatment outcomes that has been ignored in standard HTA practices.

Valuation of risk-reduction in economics and finance began with Markowitz (1952), Pratt (1964) and Arrow (1965). These early studies led to the now-classic "mean-variance" tradeoff in

[7] In other words, we assume consumers can purchase insurance against healthcare expenditures, but cannot purchase pure health indemnity insurance that enables transfer of consumption across health states. Such insurance would require perfect identification of health states to avoid fraudulent claims against the insurer, a technology that does not currently exist except in the imaginings of Star Trek "tricorders." Therefore, it sacrifices little salience, but conserves notation, if we eliminate markets for pure consumption insurance. Some firms do offer modest quantities of "critical illness insurance," which provides income transfers for a set of precisely defined diagnoses, including cancers, accidents, and heart attacks, but this set omits major portions of the potential illness spectrum. Indeed, the fact that consumers still seek to prevent illness provides prima facie evidence that financial markets are not completely insuring them against illness risk (Lakdawalla et al., 2017).

financial economics, extended by Kimball (1990) to include higher-order risk attitudes. From this background, we develop a framework that incorporates mean, variance, skewness, and (if desired) kurtosis into the valuation of stochastic health improvements. While we focus on cost-effectiveness analysis in healthcare, these methods could potentially apply to other human capital investments, e.g., vehicle safety and environmental quality.

Define p_{1T} and p_{1NT} as the probability of survival with and without treatment. Since expected utility is additively separable over time, it is also linear in survival. As a result, there is no need to consider the distribution of survival improvements; average improvements are sufficient.[8] Therefore, expected utility with treatment (T) is:

$$E[V(C,H)_T] = U(C_0)W(H_0) + p_{1\mathrm{T}}U(C_1)\{\phi E[W(H_{1s}+B)] + (1-\phi)W(H_{1w})\} \quad (3)$$

Similarly, expected utility with no treatment (NT) is:

$$E[V(C,H)_{NT}] = U(C_0)W(H_0) + p_{1\mathrm{NT}}U(C_1)\{\phi E[W(H_{1s})] + (1-\phi)W(H_{1w})\} \quad (4)$$

The difference between these two measures gives the expected incremental value of treatment T. As previously noted, traditional *ex post* evaluation perspectives coincide with the special case of $\phi = 1$.

Before proceeding to value assessment, it helps to define and discuss consumer risk preferences, which play substantial roles in valuation of risky technologies. Define the coefficient of relative risk-aversion in QoL as $r_H^* \equiv -\frac{W''(\mu_H)}{W'(\mu_H)}\mu_H$. To the best of our knowledge, there are no extant empirical estimates of relative risk-aversion over QoL, a gap in the literature that could fruitfully be addressed.[9] Lacking direct estimates of health-related risk preferences, we suggest the interim approach of assuming that relative risk parameters are similar for both QoL and consumption, and conducting sensitivity analyses around this baseline.

Following Kimball (1990), we also define "relative prudence" as $\pi_H^* \equiv -\frac{W'''(\mu_H)}{W''(\mu_H)}\mu_H$. Consumers with greater prudence are more likely to invest additional resources in QoL today in

[8] We leave to future research the questions of whether and how health technology assessment ought to depart from the conventional assumption of additively time-separable utility and the implied risk-neutrality over survival.

[9] Córdoba and Ripoll (2016) estimate a lifetime expected utility model that separates intertemporal substitution (the inverse of the coefficient of relative risk-aversion) and a "mortality aversion," which differs from risk aversion over QoL, as our model requires.

anticipation of future QoL risks (e.g., more vaccines), or changes in lifestyle that affect health risks (e.g., smoking cessation). Box A further discusses these higher-order risk terms.

[BOX A GOES HERE]

Consistent with much of the applied literature in public economics, we treat r_H^* and π_H^* as roughly constant empirically (Chetty, 2006; Noussair et al., 2013). The evidence supports this assumption directly in the case of risk-aversion over consumption, but future research must quantify relative risk preferences over QoL empirically. Our task is to explain how to value medical technology for consumers who are risk averse in QoL.

2.2.1. Quality of Life (QoL) Gains

We begin by focusing on the QoL gains produced by the technology, T. Therefore, assume that $p_{1T} = p_{1NT} = p_1$ for the moment, and compute the difference in utility between (3) and (4). To monetize this utility difference, we normalize by the marginal utility of baseline period consumption, $U'(C_0) * W(H_0)$.

Suppose that H_{1s} is a random variable measuring the untreated QoL level of a sick patient in period 1. Next, consider a treatment that produces a random QoL benefit (B) for patients in the sick state. H_{1s} and B have means μ_H and μ_B, variances σ_H^2 and σ_B^2, and Pearson skewness coefficients $\gamma_H \equiv \frac{E(H-\mu_H)^3}{\sigma_H^3}$ and $\gamma_B \equiv \frac{E(B-\mu_B)^3}{\sigma_B^3}$. Finally, define σ_{H+B}^2 and γ_{S+B} as variance and skewness coefficients of $H + B$, the QoL in the treated state. For easy reference, Box B summarizes the relevant parameter definitions.

Box B about here.

We now have:

$$EV(B) = \left[\frac{U(C_1)}{U'(C_0)}\right] \frac{p_1 \phi E[W(H_{1s}+B)-W(H_{1s})]}{W(H_0)} = K p_1 \phi \frac{[Expected\ utility\ gain\ from\ T]}{W(H_0)} \qquad (5)$$

Using Taylor series methods[10] to expand $E[W(H_{1s} + B)]$ and $E[W(H_{1s})]$ around μ_H and taking their difference gives:[11]

[10] For the salient case of utility functions in the Hyperbolic Absolute Risk-Aversion (HARA) family, Appendix 7.1 proves that this Taylor Series expansion converges.

[11] The fourth derivative, $W''''(H)$, could readily be added here, introducing kurtosis and, following Kimball (1990, 1992, 1993), relative "temperance" $= \tau_H^* = \mu_H \frac{W''''(H)}{W'''(H)}$. For compactness, we omit these terms.

$$EV(B) = \frac{Kp_1\phi}{W(H_0)}\Big\{W'(\mu_H)\mu_B + \frac{1}{2}[W''(\mu_H)[\sigma^2_{H+B} - W''(\mu_H)\sigma^2_H] +$$

$$\frac{1}{6}\left[W'''(\mu_H)\gamma_{1(H+B)}\sigma^3_{H+B} - W'''(\mu_H)\gamma_{1H}\sigma^3_H\right] + \cdots\Big\} \quad (6)$$

Collecting terms and recognizing that $r^*_H = -\mu_H \frac{W''(\mu_H)}{W'(\mu_H)}$ and $r^*_H\pi^*_H = \mu^2_H \frac{W'''(\mu_H)}{W'(\mu_H)}$ gives:[12]

$$EV(B) \approx K\frac{W'(\mu_H)}{W(H_0)}p_1\phi\mu_B\left\{1 + \left[\frac{1}{\mu_B}\right]\left[-\frac{1}{2}\left(\frac{1}{\mu_H}\right)r^*_H\,\Delta\,\sigma^2 + \frac{1}{6}\pi^*_H r^*_H\left(\frac{1}{\mu_H}\right)^2\,\Delta\,(\gamma\sigma^3) + \cdots\right]\right\} \quad (7)$$

Note that the term outside the curly braces in Equation (7) has $W'(\mu_H)$ in the numerator and $W(H_0)$ in the denominator—evaluated at different levels of QoL. We return to this issue in Section 3.1.1.

Equation (7) contains two new expressions: $\Delta\,\sigma^2 \equiv (\sigma^2_{H+B} - \sigma^2_H)$ and $\Delta\,(\gamma\sigma^3) \equiv (\gamma_{H+B}\sigma^3_{H+B} - \gamma_H\sigma^3_H)$. These expressions (respectively) reflect the difference in variances between treated and control populations, and the difference in the comparable values of Pearson's skewness (γ) times σ^3 in treated and control populations.[13] As we discuss later, these parameters are readily estimated in standard Randomized Controlled Trials (RCT) and in comparable non-randomized comparisons of treatments and control therapies (either "no treatment" or "standard of care" as appropriate).

Several things become apparent from this formulation. First, when treatment effects (μ_B) are large, the stochastic terms become relatively less important in overall incremental valuation of medical interventions. Conversely, as average treatment gains become smaller, the stochastic terms rise in importance. Over time, we expect stochastic terms to become steadily more important, since the general progress of research will typically have already found any easy and large gains in average treatment efficacy (Jones, 2009). Low-hanging fruit is usually picked

[12] Appendix 7.2 discusses convergence properties of these Taylor Series estimates. In general, convergence should be relatively rapid for $\mu_H > .2$ to .3. Even for most people with fatal disease diagnoses and only a few months to live, μ_H will generally exceed these levels.

[13] With no covariance, $\sigma^2_{H+B} - \sigma^2_H = \sigma^2_B$, and with perfect negative covariance (treatments fully restore all treated individuals to baseline health, no matter how large the acute illness severity), then $\sigma^2_{H+B} = 0$.

first. In parallel, from equation (7), given our standard assumption that relative risk parameters are approximately constant, we can see that the stochastic terms grow in importance as μ_H falls.[14]

Our framework incorporates two kinds of generalizations: 1) Diminishing marginal utility to QoL gains; and 2) stochastic gains in health. We next explore the independent effects of these phenomena on the value of QoL improvement.

The Consequences of Diminishing Marginal Utility of W

To focus on the implications of diminishing marginal utility alone, assume for the moment that QoL levels and improvements are entirely non-stochastic. Conditional on survival, the individual becomes sick with probability $\phi = 1$. Treatment benefits are also known, so that $\sigma = 0$. Equation (7) then implies that the value of non-stochastic QoL improvements is $V(B) \approx K\frac{W'(\mu_H)}{W(H_0)}p_1\mu_B$, not simply $Kp_1\mu_B$. The difference is due entirely to concavity in W, since it persists even when health is deterministic.

To explore the difference further, define two terms, first the elasticity of W with respect to H, $\omega_H \equiv \frac{W'(H_0)H_0}{W(H_0)}$, which equals $\frac{W'(H_0)}{W(H_0)}$ when $H_0 = 1$. Next, define $R \equiv \frac{W'(\mu_H)}{W'(H_0)}$, the "severity ratio," reflecting how severity of illness alters the marginal utility of QoL relative to baseline health H_0. Given these definitions, $V(B) = K\omega_H R p_1\mu_B$. In GP, the marginal value per unit of expected QoL improvement equals K. With concavity in $W(H)$, this generalizes to $K\omega_H R$. Thus—as a result of introducing risk aversion over QoL—the WTP threshold K of GP is adjusted by two multiplicative factors, ω_H and R. Diminishing marginal utility of H implies $0 < \omega_H < 1$ and $R > 1$, where R grows for more severe illnesses. In GP, on the other hand, both ω_H and R equal 1.0 as a result of utility's linearity in QoL. In subsequent sections, we calibrate ω_H and R, and we discuss their operational significance for HTA.

Our WTP threshold can be expressed as:

$$K\omega_H R = \frac{U(C_0)}{U'(C_0)}\frac{W'(\mu_H)}{W(H_0)} \quad (8)$$

[14] Alternatively, if we were to assume that absolute (instead of relative) risk parameters were constant, this $\frac{1}{\mu_H}$ effect would vanish, but the $\frac{1}{\mu_B}$ effect remains. While we currently have no evidence on how risk parameters in health behave over different levels of H, the vast bulk of data regarding consumption (C) suggest that constant relative risk parameters are much closer to reality than constant absolute risk parameters. For example, for CRRA, $\pi^* = r^* + 1$, whereas with CARA, $\pi^* = r^*$. Estimates generally show that $\pi^* \approx r_H^* + 1$, suggesting that CRRA approximately applies (Noussair et al., 2013).

For more severe illnesses, as μ_H falls, $W'(\mu_H)$ rises, and thus $K\omega_H R$ rises. Therefore, WTP for QoL improvements is higher for patients with disabilities and/or more severe illnesses. This finding resolves a long-standing question in CEA, posed most recently by the ISPOR Task Force on economic evaluation of medical technologies (Lakdawalla et al., 2018): "Does untreated health status affect the value of health gains?" By introducing risk aversion over QoL, we demonstrate that the answer is unequivocally "yes." R and hence $K\omega_H R$ rise as untreated QoL worsens. Later, we will show that $K\omega_H R$ functions as the cost-effectiveness threshold, analogous to the role of K in GP. Therefore, R also captures the effect of illness severity on the proper cost-effectiveness threshold.

The Consequences of Uncertain Treatment Outcomes

To address how stochastic QoL modifies values of QoL improvements, we turn to the portion of equation (7) in curly braces and define:

$$\epsilon \equiv \left\{1 + [\tfrac{1}{\mu_B}][-\tfrac{1}{2} r_H^* \left(\tfrac{1}{\mu_H}\right) \Delta\sigma^2 + \tfrac{1}{6}\pi_H^* r_H^* \left(\tfrac{1}{\mu_H}\right)^2 \Delta(\gamma\sigma^3) + \cdots]\right\} \tag{9}$$

In Equation (7), expected QoL gains are $p_1\phi\mu_B$. Therefore:

$$E\big(V(B)\big) = [K\omega_H R][p_1\phi\mu_B\epsilon] \tag{10}$$

We describe $[K\omega_H R]$ as the Risk-Aversion and Severity Adjusted Willingness to Pay (RASA-WTP), and $[p_1\phi\mu_B\epsilon]$ as the Risk-Adjusted QALY (RA-QALY). The former affects the value per unit of health gain, and the latter represents the "certainty-equivalent" amount of QoL gained from a medical treatment. Thus

$$E\big(V(B)\big) = \text{RASA-WTP} * \text{RA-QALY} \tag{11}$$

This generalizes the traditional value measure, which states that $E\big(V(B)\big) = K * QALY$.

Note that $\mu_B\epsilon$ yields the "certainty-equivalent" gain in QoL associated with the medical technology. That is, a technology that produces $\mu_B\epsilon$ units of QoL with certainty is equal in value to the technology producing the stochastic gain, B, with average gain, μ_B. The term, ϵ, measures the certainty-equivalent QoL units the consumer requires in exchange for each additional unit of average QoL gain. For this reason, we call ϵ the "certainty-equivalence ratio."

Among its other functions, ϵ quantifies the importance of stochastic terms in value assessment. With entirely deterministic QoL, $\epsilon = 1$. Analogously, when $\epsilon \approx 1$, the stochastic

terms are easily ignored, and average QoL gains are largely sufficient metrics of benefit. As ϵ departs from unity, in either direction, the stochastic terms become relatively more important.

Variance and skewness have opposing effects on ϵ. Increases in variance reduce expected utility, and vice-versa. If a treatment increases the variance of QoL (i.e., $\sigma_{H+B}^2 > \sigma_H^2$), any given average gain in QoL will be worth less. For sufficiently small μ_B and increase in variance, ϵ might even become negative, so that positive average gains in QoL could in fact be costly, when considering expansions in variance. In contrast, treatments that reduce variance augment the value of the average QoL improvement. The latter gives an example of "insurance value" from medical technology by reducing the uncertainty surrounding health outcomes (Lakdawalla et al., 2017).

Conversely, increases in positive skewness add value. If $\gamma_{1(H+B)}\sigma_{H+B}^3 > \gamma_{1H}\sigma_H^3$, this increase in positive skewness augments average values of QoL improvements. This effect has not previously been formalized in the literature, but it seems related to empirical findings that patients value therapies providing modest chances of a very large health improvement. Prior research has referred to this as the "value of hope" associated with, for instance, immunotherapies treating metastatic cancer (Lakdawalla et al., 2012). Risk-averse consumers dislike variance in treatment outcomes. However, to the extent variance exists in competing treatments, consumers will prefer those with variance that has greater positive skewness. Table 1 summarizes various combinations of changes in variance and skewness and their effects on medical technologies' value.

[INSERT TABLE 1 HERE]

When $\epsilon < 1$, gains in average QALYs are less valuable than corresponding gains in certain QALYs. This will be true if the technology increases the variance of QoL. In contrast, if $\epsilon > 1$, consumers would rather take one unit of average QoL gain instead of a sure one-unit gain in QoL. This could occur for technologies with lower variance and/or highly positively skewed treatment benefits.

2.2.2. Longevity Gains

Next, we generalize our analysis to consider technologies that produce both random QoL benefits B, *and* average survival increases of μ_p. Building upon equation (5), the technology's expected value is given by:

$$EV(\mu_p, B) = K\frac{\mu_p[\phi E(W(H_{1S}+B))+(1-\phi)W(H_{1W})]+p_1\phi E[W(H_{1S}+B)-W(H_{1S})]}{W(H_0)} \quad (12)$$

The first part of the health-related numerator component (involving μ_p) describes gains in expected utility from increased survival, and the second part restates the gains arising from improved QoL, holding survival constant. Section 2.2.1 characterized the second part of this expression by proving that $K\frac{p_1\phi E[W(H_{1S}+B)-W(H_{1S})]}{W(H_0)} = [K\omega_H R][p_1\phi\mu_B\epsilon]$. We now characterize the first part of this expression, which corresponds to the contributions of survival gains, $K\frac{\mu_p[\phi E(W(H_{1S}+B))+(1-\phi)W(H_{1W})]}{W(H_0)}$.

To relate this to our earlier analyses of GP, define δ, the marginal rate of substitution between survival and QoL, (with dimension of H).[15]

$$\delta \equiv \frac{[\phi E(W(H_{1S}+B))+(1-\phi)W(H_{1W})]}{W'(\mu_H)} \quad (13)$$

Here, the numerator describes the marginal expected utility from gains in survival (Δp_1) and the denominator describes the marginal utility of gains in QoL.[16] Recall that this marginal rate of substitution equals H_1 when W is linear, a result that undergirds the equivalence of survival and QoL gains in traditional CEA, per GP.

This definition of δ, coupled with our earlier analysis of QoL improvements, allows rewriting (12) as:

$$EV(\mu_p, B) = K\omega_H R\{\mu_p\delta + p_1\phi\mu_B\epsilon\} \quad (14)$$

Recall that $K\omega_H R$ measures the value of gains in QoL, either through survival improvement or quality-adjusted life improvement. Define the term in curly braces as "generalized risk-adjusted-quality-adjusted life-year" (GRA-QALY):

$$GRA\text{-}QALY(\mu_p, B) = \{\mu_p\delta + p_1\phi\mu_B\epsilon\} \quad (15)$$

The GRA-QALY can be used in place of the QALY in decision analysis and value assessment. To see why, observe that increases in the GRA-QALY are always worth $K\omega_H R$ on the margin, regardless of whether generated by survival gains, QoL gains, or some combination

[15] A later section discusses methods to estimate δ.

[16] This is the "*ex post*" rate of trade, just as in GP, so $p_1 = 1$..

of the two. It functions as a "single index of value," just like the QALY in GP's framework. The GRA-QALY generalizes the RA-QALY further by incorporating longevity-extending benefits. In addition, it transforms the average QoL gain, $p_1\phi\mu_B$, into the certainty-equivalent gain in QoL, $p_1\phi\mu_B\epsilon$.

2.2.3. Implications for Cost-Effectiveness Analysis

We now turn to cost-effectiveness analysis for medical technologies with uncertain benefits. Building upon equation (2), the technology is welfare-improving if

$$K\omega_H R\{\mu_p\delta + p_1\phi\mu_B\epsilon\} \geq \Delta C_0 \tag{16}$$

or

$$\frac{\Delta C_0}{\mu_p\delta + p_1\phi\mu_B\epsilon} \leq K\omega_H R \tag{17}$$

The left-hand side of (17) is the ratio of incremental medical spending to units of certainty-equivalent QoL improvement. The right-hand side of (17) is the Risk-Aversion and Severity Adjusted WTP (RASA-WTP).

This generalizes the standard incremental cost-effectiveness ratio (ICER), by accounting for stochastic changes in QoL, diminishing returns to QoL improvement, and the severity-of-illness adjustment R. It also permits application of traditional ICER-based decision analyses as carried out in current CEA studies. Once the new parameters δ, ϵ, ω_H, and R are estimated (see Section 3), these cost-effectiveness ratios are straightforward to obtain for a given technology.

Note that, if important, differences in variability of treatment cost (ΔC_0) could readily be introduced at this point with similar Taylor Series estimates using financial risk measures such as relative risk aversion and relative prudence. Since cost-effectiveness analysis typically assumes risk-averse consumers are insured against healthcare spending, and since third-party payers are typically thought of as risk-neutral, risk-aversion over costs is likely to be a more specialized matter than over QoL. Therefore, we leave this analysis to future research and focus on the particular issues of risk aversion in QoL.

Our framework nests the traditional GP decision rule as a special case. By assuming linearity in the utility of QoL, traditional cost-effectiveness analysis imposes functional form restrictions that imply $\delta = H_1$ and $\omega_H = R = 1$. By abstracting from the effects of uncertainty,

traditional analyses also presume that $\epsilon = 1$. For these parameter values, Equation (17) collapses down to the GP decision rule, where $\frac{\Delta C}{\Delta QALY} \leq K$.

Our analysis also shares with GP the applicability either to welfare-maximizing social planners or budget-constrained payers seeking second-best allocations conditional on fixed budgets. The risk-aversion and severity adjusted threshold, $K\omega_H R$, depends on the chosen levels of consumption. Budget-constrained payers will have smaller $K\omega_H R$ thresholds than unconstrained welfare-maximizing decision makers (Phelps, 2019a). However, within such budget-constrained environments, better resource allocation occurs using our model by incorporating both severity of illness and uncertain treatment outcomes rather than by ignoring these fundamental issues brought about by risk aversion in QoL.

3. CALIBRATION AND ESTIMATION

Incorporating risk-aversion and uncertain treatment benefits requires additional parameters. We now turn to calibrating several of these parameters and providing guidance for the estimation of others. We begin by discussing the calibration of the cost-effectiveness threshold, $K\omega_H R$. We then discuss calibration of δ, the marginal rate of substitution between survival and QoL, and finally, the certainty-equivalence ratio, ϵ.

3.1. Calibrating $K\omega_H R$

Some existing estimates for the value of QoL improvements may remain viable as estimates of $K\omega_H$. Suppose we have empirically valid[17] reduced-form estimates for the value of a quality-adjusted life-year.[18] By nature, reduced-form estimates do not specify underlying utility functions, but instead yield atheoretical estimates of willingness to pay. Assuming such estimates come from a population of people with diminishing marginal utility in QoL, then $K\omega_H R$ more naturally describes these estimated parameters, presumably at average values of R.[19] In contrast, studies that take structural approaches estimating only the component where

[17] As Phelps (2019a) has emphasized, care must be taken in assessing and interpreting estimated values of a QALY from VSL studies.

[18] See Hirth et al. (2000) for a review of the literature.

[19] We are aware of no analyses that directly measure severity of illness effects (R) when estimating WTP in such studies.

$K = \frac{C}{\omega_C}$ or equivalently $\frac{K}{C} = \frac{1}{\omega_C}$ (Phelps, 2019a) must be adjusted by the factor $\omega_H R$. Regardless, more can be said about the calibration of $K\omega_H R$ from a theoretical perspective.

3.1.1. Calibrating R

We begin by calibrating $R \equiv \frac{W'(\mu_H)}{W'(H_0)}$, the disease severity ratio. The full Taylor Series expansion for $W'(\mu_H)$ has a convenient structure. First, define the average health loss $\ell^* \equiv H_0 - \mu_H$. We can expand $W'(\mu_H)$ around H_0:[20]

$$W'(\mu_H) = \{W'(H_0) + W''(H_0)(\mu_H - H_0) + \frac{1}{2}W'''(H_0)(\mu_H - H_0)^2 + \frac{1}{6}W''''((H_0)(\mu_H - H_0)^3 + \cdots.\} = W'(H_0)\left\{1 + r_H^*\ell^* + \frac{1}{2}r_H^*\pi_H^*\ell^{*2} + \frac{1}{6}r_H^*\pi_H^*\tau_H^*\ell^{*3} + \cdots.\right\} \quad (18)$$

Thus,

$$R = \frac{W'(\mu_H)}{W'(H_0)} = \left\{1 + r_H^*\ell^* + \frac{1}{2}r_H^*\pi_H^*\ell^{*2} + \frac{1}{6}r_H^*\pi_H^*\tau_H^*\ell^{*3} + \cdots.\right\} \quad (19)$$

This expansion simplifies further for any CRRA utility function, wherein $\pi_H^* = r_H^* + 1, \tau_H^* = r_H^* + 2, \ldots$. (see Appendix 7.3). Thus, with CRRA utility, the disease severity ratio, *R*, becomes simply a function of the CRRA parameter, r_H^*. Since it is common in applied economics to calibrate utility functions using the CRRA assumption, this approach provides a well-established strategy for calibrating R.

When $r_H^* = 1$, Equation (19) further collapses to the familiar value of a perpetuity, since the products of the relative risk parameters in the numerator of each term involving r_H^* just cancel the factorial terms in the denominator, leaving a value of $R = 1 + \left(\frac{\ell^*}{1-\ell^*}\right) = \left(\frac{1}{1-\ell^*}\right)$. Thus, for a small health loss, $\ell^* = .1$, we obtain the severity ratio $R = 1.11$. For the moderate health loss, $\ell^* = \frac{1}{2}$, we see $R = 2$. For a very large health loss, $\ell^* = 0.9$, we obtain $R = 10$.

Evaluating R becomes slightly more complex when $r_H^* \neq 1$. Using Equation (19), which applies to all utility functions (not just CRRA), Table 2 shows values of R across a range of values for ℓ^* and r_H^*.[21] QoL loss, ℓ^*, runs from 0 to 0.9. Relative risk-aversion varies from zero

[20] The signs in each term are all positive because of alternating powers of $\bar{\ell} = -(\mu_H - H_0)$.
[21] For non-CRRA utility functions, the values of r_H^* are assumed to obtain locally.

(to incorporate the GP case) to 1.5, which lies beyond consensus estimates of consumption risk-aversion, r_C^* (Chetty, 2006; Phelps, 2019a).

Two general results emerge. First, the multiplier R rises with r_H^*. In the extreme case of $r_H^* = 0$, the total multiplier is 1, and $W'(\mu_H) = W'(H_0)$, the (nested) result from GP. Second, R rises with ℓ^*, increasingly so as ℓ^* rises.

[INSERT TABLE 2 HERE]

To provide readers with context, Table 3 reports example diseases across various values of ℓ^*. Examples with ℓ^* at or below 0.1 include peptic ulcer disease and benign prostatic hyperplasia. At the other end of the spectrum, with ℓ^* greater than 0.7, examples include Alzheimer's disease and metastatic colorectal cancer.[22]

[INSERT TABLE 3 HERE]

For average values of r_H^* near H_0, it can also be shown that when utility is DRRA (IRRA), R will be larger (smaller) than shown in Table 2.[23] For r_H^* in the vicinity of 1, R is roughly unity for very mild diseases with $H_{1s} \approx 1$, roughly equal to 2.0 for diseases with $H_{1s} \approx 0.5$, just over 3.0 for $H_{1s} \approx 0.3$, and over 4.5 for $H_{1s} \approx 0.2$.

Finally, we note that this discussion does not preclude the possibility of estimating R in the usual stated preference fashion. The term $R \equiv \frac{W'(\mu_H)}{W'(H_0)}$ represents WTP for QoL improvements in the sick state, in terms of foregone QoL in the baseline healthy state. The question is how many units of sick state QoL would be required in order to compensate a consumer for one unit of lost QoL in the healthy baseline.

3.1.2. Calibrating $K\omega_H$

We now turn to the other part of the cost-effectiveness threshold expression, $K\omega_H \equiv [\frac{U(C_1)}{U'(C_0)}][\frac{W'(H_0)}{W(H_0)}]$. Define $G \equiv \frac{U(C_1)}{U(C_0)}$ as the relative growth in period consumption utility and $\omega_C \equiv U'(C_0)\frac{C_0}{U(C_0)}$ as the elasticity of period consumption utility with respect to consumption. The

[22] The example diseases and associated untreated baseline QALYs are taken from cost-effectiveness studies reported in the Tufts Cost-Effectiveness Analysis Registry (as updated in February 2014).

[23] Suppose utility is DRRA. Appendix 7.3 shows that $\pi^* > r^* + 1$ so the product $\pi^* r^*$ is larger than when utility is CRRA and similarly for all successive higher levels of risk parameters. Thus, the Taylor series sum is larger with DRRA than with CRRA. The reverse holds when utility is IRRA. The spreadsheet that created Table 2 demonstrates this phenomenon when the gap between $r^* and\ \pi^*$ (and successive higher order terms) is allowed to differ from 1.0.

term ω_C is analogous to ω_H, the elasticity of *W(H)* with respect to QoL. Some algebraic manipulation reveals that:

$$K\omega_H = C_0\left(\frac{\omega_H}{\omega_C}\right)G \tag{20}$$

In most applications, analysts presume little or no change over time in consumption, implying $G \approx 1$ and $K\omega_H \approx C_0\left(\frac{\omega_H}{\omega_C}\right)$. Thus, $K\omega_H$ is proportional to C_0 where the factor of proportionality is approximately $\left(\frac{\omega_H}{\omega_C}\right)$.

While utility elasticities like ω_H and ω_C are not commonly estimated in the literature,[24] we can characterize them for the salient case in which $U(C)$ and $W(H)$ are Hyperbolic Absolute Risk-Aversion (HARA) utility functions. HARA functions include constant, increasing, and decreasing relative risk-aversion cases, along with many commonly used functions – exponential utility, power utility, linear utility, quadratic utility, and logarithmic utility, and are commonly used in the finance literature studying risk aversion (Merton, 1971).

To pursue this point further, note that for $z \equiv \frac{aC}{1-\gamma_C} + b > 0$, HARA utility is:

$$U(C) \equiv \left[\frac{1-\gamma_C}{\gamma_C}\right] z^{\gamma_C}, \tag{21}$$

Relative risk-aversion over consumption is defined as $r_C^* = \frac{aC}{z}$, and the elasticity with respect to consumption is $\omega_C = \left[\frac{\gamma_C}{1-\gamma_C}\right]\left[\frac{aC}{z}\right] = \left[\frac{\gamma_C}{1-\gamma_C}\right] r_C^*$. Similar relationships hold for $W(H)$, so that $\omega_H = \left[\frac{\gamma_H}{1-\gamma_H}\right] r_H^*$. Thus, $\frac{\omega_H}{\omega_C}$ moves with changes in r_C^* (inversely) or r_H^* (directly), as modified by the ratios involving the respective γ values. Further, if $U(C)$ and $W(H)$ exhibit identical (similar) relative risk preference parameters, r^* and π^*, they will have identical (similar) values of γ.[25] This similarity in both r^* and π^* suffices to make the relationships between r^* and ω similar for $U(C)$ and $W(H)$.[26]

---

[24] A distinct approach arises from the literature using reported "happiness" measures, as developed originally by Easterlin (2004).

[25] Phelps and Cinatl (2020) prove that for HARA utility, $\gamma = (\pi^* - 2r^*)/(\pi^* - r^*)$, so γ can be recovered from π^* and r^*.

[26] In the special case of CRRA, $\omega_C = \gamma_C = (1 - r_C^*)$ and $\omega_H = \gamma_H = (1 - r_H^*)$, so $\frac{\omega_H}{\omega_C} = \frac{(1-r_H^*)}{(1-r_C^*)} = \frac{\gamma_H}{\gamma_C}$.

Lacking direct measurement of risk attitudes in health in the current empirical literature, an interim approach is to assume similar risk attitudes in H and C, i.e., $\frac{\omega_H}{\omega_C} \approx 1$, and to conduct sensitivity analysis around this placeholder assumption. As new evidence emerges on its actual value, the $\frac{\omega_H}{\omega_C}$ ratio can easily be adjusted appropriately.

In sum, the WTP threshold can be approximated as $K\omega_H R \approx C_0 \left(\frac{\omega_H}{\omega_C}\right) GR$. This expression is a multiple of base period consumption, C_0, where the multipliers are $G \approx 1$, R (values of which can be read off Table 2), and $\left(\frac{\omega_H}{\omega_C}\right)$. R specifies how the optimal WTP increases with untreated illness severity. G reflects the growth in the value of QoL that results from consumption growth over time. The ratio $\left(\frac{\omega_H}{\omega_C}\right)$ illuminates how WTP changes with risk aversion in $W(H)$ and $U(C)$.

3.2.Estimating the GRA-QALY

The GRA-QALY associated with a technology that increases QoL by B and survival by an average of μ_p is given by $GRA\text{-}QALY(\mu_p, B) = \{\mu_p \delta + p_1 \phi \mu_B \epsilon\}$. This requires estimates of δ, the marginal rate of substitution between survival and QoL, and of ϵ, the certainty-equivalence ratio.

3.2.1. Marginal rate of substitution between survival and QoL

Recall the definition $\delta \equiv \frac{[\phi E(W(H_{1S}+B))+(1-\phi)W(H_{1W})]}{W'(\mu_H)}$, which represents the ratio between the marginal utility of survival gains and the marginal utility of QoL improvement. In equilibrium, this marginal rate of substitution will also satisfy $\delta = \frac{dH}{dp^S}$, which measures the units of QoL that an individual will give up in exchange for a unit gain in survival probability. Using this latter representation, δ can be recovered via standard time tradeoff (TTO) approaches to valuing different health states: specifically, one needs to estimate the rate of trade between survival and expected quality of life, holding consumption fixed. TTO methods rely on the concept of compensating variation (see Figure 1a). Consumers are given a choice between one fixed alternative (baseline survival probability and expected QoL) and another with less remaining survival probability but perfect health (QoL = 1). They are asked to state what

survival probability they would give up to move from their current health condition to perfect health, thus representing WTP in the same way that Figure 1a considers.[27]

To illustrate how a TTO study would estimate δ, define the baseline expected QoL in the study as $E(H^B)$. Define the survey respondent's "stated" change in the survival probability as $|\Delta p^S|$. Based on the results of the TTO study design above, δ could be estimated as, $\hat{\delta} = \frac{1-E(H^B)}{|\Delta p^S|}$. For example, suppose survey respondents were asked to consider end-stage osteoarthritis of the knee, which presents with $H_{1s} \approx 0.75$. We would then ask how much survival probability they would give up in exchange for restoring perfect QoL. Suppose they state a willingness to reduce survival by 5 percentage points in exchange for this increase in QoL. In this case, therefore, we would estimate $\hat{\delta} = \frac{1-E(H^B)}{|\Delta p^S|} = \frac{1-0.75}{0.05} = 5.0$.[28]

As Figure 1 illustrates, δ varies with disease severity, and thus is essentially a disease-specific parameter. Health states involving high longevity but low QoL will likely result in lower values of δ, because individuals in such states are more willing to trade away longevity in exchange for QoL. As an example, suppose we surveyed respondents about severe Alzheimer's disease rather than osteoarthritis, so that $H_{1s} \approx 0.25$. The QoL loss is three times what it was in the osteoarthritis example. Therefore, if diminishing marginal utility of QoL holds, we expect that respondents will give up more than three times the survival reduction they agreed to in the osteoarthritis case, or more than $3 * 5\% = 15\%$ points of survival.

3.2.2. Calibrating the certainty-equivalence ratio

The certainty-equivalence ratio is defined as $\epsilon = \left\{1 + [\frac{1}{\mu_B \mu_H}][-\frac{1}{2} r_H^* \Delta\sigma^2 + \frac{1}{6}\pi_H^* r_H^* \left(\frac{1}{\mu_H}\right)\Delta(\gamma\sigma^3) + \cdots]\right\}$. Therefore, ϵ depends on consumer risk-attitudes (r_H^*, π_H^*) and differences in statistical risks between patients treated by T and the relevant control group (C). While ϵ is technically a disease-specific parameter, it can in practice be recovered quite readily

---

[27] Earlier methods simply asked consumers for their estimated tradeoff value. More modern approaches use discrete choice experiments to show subjects continuously modified pairs of choices until they reach the point of indifference. This process is logically equivalent to standard methods for selecting corrective lenses for eyeglasses, e.g. "Which do you prefer, A or B" with the pairs adjusted according to prior responses.

[28] For simplicity, we consider a deterministic example in which $\phi = 1$ and $\sigma_H^2 = \sigma_{H+B}^2 = 0$. Stochastic cases would simply fold the relevant uncertainties into the problem. E.g., one might ask, "suppose you have an X% chance of osteoarthritis," or "suppose you have a Y% chance of mild osteoarthritis with QoL=0.85 and a 1-Y% chance of severe osteoarthritis with QOL=0.7," and so on.

for each decision context. First, risk attitudes over consumption are typically assumed to be constant and context-independent (Chetty, 2006; Noussair et al., 2013), and we would propose a similar approach to risk attitudes over health. Second, the parameters μ_B and μ_H are already routinely estimated on a disease- and treatment-specific basis. Finally, as we discuss below, changes in variance and skewness could be recovered from clinical trials or cost-effectiveness simulation models.

Estimating risk-attitudes in health.

To date, there are no estimates of relative risk-aversion and relative prudence over QoL or any measure of health. Given the evident importance of risk preferences over health, we hope this deficiency will be remedied in the near future. The current cost-effectiveness literature assumes $r_H^* = \pi_H^* = 0$. As a placeholder, we suggest assuming that risk attitudes over QoL and consumption are similar and conducting sensitivity analyses that allow the ratio between the two to vary. In other words, the analyst can explore what happens if relative risk-aversion over QoL is X% higher or lower than relative risk-aversion over consumption, and so forth.

Estimating variability in QoL benefits.

The statistical parameters to calculate $\Delta\sigma^2$ and $\Delta(\gamma\sigma^3)$ can be estimated from the data readily available in normal HTA studies, e.g., Randomized Controlled Trials (RCTs). Such studies routinely measure average differences in QoL outcomes (μ_B) and the variance of these outcomes in treatment and control populations, in order to estimate standard errors of mean differences. Thus, nearly all current HTA efforts have the data to incorporate the higher-order risk terms necessary for computing the certainty equivalence ratio, ϵ. To incorporate effects of changes in skewness will require (at a minimum) further data analysis to measure skewness of outcome distributions for alternative treatments and then their differences. We leave to relevant experts (e.g., biostatisticians and statisticians) the proper study design and calculations for these parameter estimates. We note, however, that if QALYs are not estimated in clinical trials, cost-effectiveness studies will typically produce them via simulation models; in this case, the cost-effectiveness researcher should estimate the variance and skewness in QoL benefits.

3.3. Consequences for CEA and HTA

3.3.1. Implications for Cost-Effectiveness Thresholds

GP assumed $r_H^* = 0$ and $\omega_H = 1$, with the resulting value of QoL improvement approximately $K \approx C_0 \frac{1}{\omega_C}$, as long as C_1 is sufficiently close to C_0. Once we allow for risk aversion over health, this approximation becomes $K\omega_H R \approx C_0 \frac{\omega_H}{\omega_C} R$.

Imposing the placeholder assumption that $\frac{\omega_H}{\omega_C} \approx 1$, the cost-effectiveness threshold becomes $K\omega_H R \approx CR$. For mild-severity diseases (where $R \approx 1$), this is approximately $C \frac{\omega_H}{\omega_C} = C$, well below common estimates in the current literature, which suggests $2C$ to $3C$ as appropriate values (Phelps, 2019a). Therefore, our analysis suggests that the correct WTP threshold might fall for mild conditions, relative to current practice.

However, for severe illnesses, the threshold might well be larger than current methods suggest, because the multiplier R rises with illness severity. For values of r_H^* meaningfully different from zero, R rises geometrically with illness severity (see Table 2). Instead of applying a single threshold K to every illness condition, it will be necessary to evaluate the level of health for untreated patients in various health conditions, and to adjust the threshold as illustrated in Table 2.

Using $r_H^* = 1, \frac{\omega_H}{\omega_C} \approx 1$, and using the diseases in Table 3 as examples, severe Alzheimer's disease ($H_{1S} \approx 0.2$, or $\ell^* \approx 0.8$) calls for a threshold of 5 times annual consumption, pressure ulcers in nursing home residents ($H_{1S} \approx 0.35$) calls for a threshold around 2.85 times annual consumption, acute lung injury ($H_{1S} \approx 0.6$) calls for 1.67 times annual consumption, and peptic ulcers ($H_{1S} \approx 0.97$) requires a threshold approximately the same as annual consumption.[29] The severity gradient steepens as r_H^* increases, as Table 2 shows. As more specific information becomes known about $\frac{\omega_H}{\omega_C}$, these thresholds can be adjusted more precisely.

Care must be taken to ensure objectivity in the evaluation of illness severity, and the associated computation of cost-effectiveness thresholds. Both third-party payers and life sciences manufacturers benefit from changes in these severity estimates, but in opposite directions. Similar tables of severity are in widespread use in hospital payment systems in the

[29] For $r_H^* = 1$, these values can all be calculated from the perpetuity value given previously.

United States, such as the Diagnosis-Related-Groups (DRG) system in Medicare. Creating such a table for general use would seem relatively straightforward. Building upon registries of cost-effectiveness analyses – e.g., the Tufts Cost-Effectiveness Analysis Registry, used for the examples in Table 3 – would seem a natural way to proceed.

3.3.2. The Consequences of Permanent Disability

Cost-effectiveness thresholds will also need to vary for those with permanent disability. To pursue this point, we now relax the assumption that $H_0 = 1$. Consider a permanent disability that reduces H_0 to some other level. Recall from Equation (1) that $K \equiv \frac{U(C_1)}{U'(C_0)}\left[\frac{1}{H_0}\right]$. When we allow for $H_0 < 1$, $K\omega_H R = \frac{U(C_1)}{U'(C_0)}\left[\frac{1}{H_0}\right]\omega_H R$. Relying on the arguments above, this can be rewritten as $K\omega_H R = G\left[\frac{\omega_H}{\omega_C}\right]R\left[\frac{C_0}{H_0}\right]$. Assuming that consumption growth is approximately zero, G can be ignored, and $K\omega_H R \approx \left[\frac{\omega_H}{\omega_C}\right]R\left[\frac{C_0}{H_0}\right]$. The value-per-QALY measure is a multiple of consumption units per QALY, $\frac{C_0}{H_0}$, where $\left[\frac{\omega_H}{\omega_C}\right]R$ is the multiplication factor.

To keep the key issue in focus, we assume for the moment that $\left[\frac{\omega_H}{\omega_C}\right]$ remains constant across health and consumption levels; for instance, this would be true if both $U(C)$ and $W(H)$ exhibited constant relative risk aversion. Under this assumption, permanent disability increases the cost-effectiveness threshold for two reasons. First, since $K\omega_H R \approx \left[\frac{\omega_H}{\omega_C}\right]R\left[\frac{C_0}{H_0}\right]$, reductions in H_0 proportionally inflate WTP by reducing the denominator. Second, permanent disability also increases R. To see why, return to equation (19), which defines R. Notice that the parameter ℓ^* representing illness severity is normalized by H_0. As H_0 shrinks from $H_0 = 1$, therefore, ℓ^* increases, and the Taylor series for R grows.

To give an example of the combined magnitudes of these effects, suppose that a permanent disability reduces H_0 from 1 to 0.8. Now introduce an acute illness that has $\ell^* = 0.4$ when H_0 = 1, but that results in $\ell^* = 0.5$ when $H_0 = 0.8$. We can readily calculate the effect of disability by assuming CRRA for $W(H)$. Then, the perpetuity value of R increases from $\frac{1}{1-0.4} = \frac{1}{0.6} = 1.667$ to $\frac{1}{1-0..5} = \frac{1}{0.5} = 2$. This disability increases $K\omega_H R$ by a factor of $\left[\frac{2}{1.667}\right]\left[\frac{1}{0.8}\right] = 1.5$. This effect expands geometrically as the disability severity increases (H_0 falls).

Through this analysis, introducing risk aversion in QoL resolves a long-standing dilemma in the practice of CEA, wherein people with permanent disabilities are disfavored. Indeed, the Affordable Care Act explicitly prohibits using any CEA model that discriminates against disabled persons. Several "workaround" efforts have emerged to attempt to repair this problem (Basu et al., 2020; Nord et al., 1999), altering QALY equivalents received by persons with permanent disability. Our approach shows that this may not be necessary. Our model of value does not discriminate against disabled persons when valuing QoL improvements. Rather, it adds value to curing their disabilities and makes it more (not less) valuable to cure them of any acute illness.

Some complexities continue to arise in the valuation of survival gains for the permanently disabled. Diminishing returns to QoL gains mean that rational consumers demand more QoL in exchange for a given gain in survival, when QoL falls. Our model then implies that a unit change in survival is worth fewer QoL units to a disabled person. Theoretically, this is a fairly direct and uncontroversial implication of diminishing returns to QoL.

To see this precisely, from equation (14), the component valuing survival gains is $K\omega_H R\mu_p\delta$. Looking at Equation (13) that defines δ, define $D = \frac{[\phi E(W(H_{1S}+B))+(1-\phi)W(H_{1W})]}{W(H_0)}$, i.e. the change over time in utility from health-related QoL. Further, to emphasize the role of disability, consider the case where $\phi = 0$, eliminating acute illnesses from the analysis. Now define $H_{1W} = H_0 - d$, where d is the QoL loss from disability.

A bit of algebraic manipulation shows that $\frac{\delta}{H_0} = \frac{D}{\omega_h R}$. Therefore, the expression for the value of survival improvement—$K\omega_H R\delta\mu_p$—reduces to $KD\mu_p$ In the general evaluation of GRA-QALYs, the concavity effects of $W(H)$ are removed when evaluating the deterministic components of life expectancy gains. This is intuitively plausible, since we have assumed risk neutrality in survival. With these simplifying assumptions that remove acute illness from the picture, the WTP for survival becomes:

$$KD = K\frac{W(H_0-d)}{W(H_0)} \tag{22}$$

WTP for survival declines as disability worsens. This is wholly expected, as portrayed in Figure 1a. People with lower QoL would be more willing to give up LY to gain QoL. Preferences change as disability alters health status.

As logical as this conclusion is, decision makers might wish to remove permanent disability from evaluation of technologies that extend life expectancy. For the sake of consistency, this would then imply that the WTP for QoL improvements is the same for the permanently disabled as the non-disabled, thus discarding our earlier conclusion at Equation (8) that value of QoL gains rises with disability severity. The ethical question becomes whether or not to acknowledge the higher relative value of QoL improvements compared to survival gains among the disabled. The answer to this question lies beyond the scope of economic analysis.

3.4. Summary of Parameters to be Estimated

Equation 17 provides the generalized ICER decision rule for our context:

$$\frac{\Delta C_0}{\mu_p \delta + p_1 \phi \mu_B \epsilon} \leq K \omega_H R$$

Leveraging the approximations in the prior section, we can rewrite this as:

$$\frac{\Delta C_0}{\mu_p \delta + p_1 \phi \mu_B \epsilon} \leq \left[\frac{\omega_H}{\omega_C}\right] R \left[\frac{C_0}{H_0}\right] \quad (23)$$

The parameters, ΔC_0, μ_p, p_1, and μ_B are routinely estimated in the current literature, and the probability of illness, ϕ, is also widely available in burden of illness estimates. Furthermore, C_0 and H_0 can be taken from a variety of existing sources on income, consumption, and QoL for a subgroup of interest. This leaves δ, ϵ, $\frac{\omega_H}{\omega_C}$, and R to be estimated.

In summarizing the estimation approaches, we suggest restricting attention to the case of HARA utility, which covers a wide variety of common utility functions. We also suggest borrowing from the empirical literature on consumption risk-aversion the assumption that relative risk-aversion over health, r_H^*, is roughly constant empirically. This leads to the following estimation approaches.

$\boldsymbol{r_H^*, \pi_H^*}$: Relative risk attitudes over health need to be estimated, if we are to measure accurately the costs and benefits of risky health outcomes. These can be estimated using methods from the happiness economics literature (Easterlin, 2004) or discrete choice experiments (Dohmen et al., 2011; Ebert and Wiesen, 2011, 2014; Eckel and Grossman, 2008; Eisenhauer and Ventura, 2003; Harrison et al., 2007; Meyer and Meyer, 2006; Noussair et al., 2013). As an interim step, we recommend using the corresponding risk attitudes over consumption as placeholders, with appropriate sensitivity analysis around this benchmark.

$\boldsymbol{R}$: Under CRRA utility, R depends entirely on r_H^*. Thus, once r_H^* is known, so is R. More general estimates are possible when relative prudence, π_H^*, is also known, according to Equation 19 and with assumptions about whether r_H^* is constant, increasing or decreasing in H.

$\frac{\boldsymbol{\omega_H}}{\boldsymbol{\omega_C}}$: Under HARA utility and constant relative risk-aversion, ω_H depends on r_H^* and the HARA utility parameter γ_H. Moreover, in this case, γ_H depends on r_H^* and π_H^* (see footnote 25). Analogous results hold for ω_C. Therefore, $\frac{\omega_H}{\omega_C}$ is identified as soon as estimates of r_H^* and π_H^* are in hand.

$\boldsymbol{\epsilon}$: The certainty-equivalence ratio depends on relative risk preferences, r_H^* and π_H^*, along with μ_B, μ_H, $\Delta\sigma^2$, and $\Delta(\gamma\sigma^3)$. As discussed earlier, the latter two parameters must be estimated in randomized trials or in cost-effectiveness simulation models that predict QoL improvements.

$\boldsymbol{\delta}$: The marginal rate of substitution between longevity and QoL likely varies with disease severity, by the logic of Figure 1. As discussed in Section 3.2.1, it can be recovered via time trade-off survey methods implemented in relevant patient populations.

4. CONCLUSIONS

Cost-effectiveness analysis has become one of the most successful economic methods in real-world applications to evaluate medical technologies, and its use continues to expand. However, standard CEA frameworks fail to adequately account for the role of risk aversion in QoL and attendant uncertainty in treatment effects. This can lead to misallocation of resources by health insurers and/or health care systems that rely on it. We develop a relatively straightforward and tractable way for analysts and real-world decision makers to account for these limitations. While several new parameters are needed, once these estimates are in place, our method is no more taxing to carry out than current approaches.

How often will our generalized methods matter? Our calibration analysis provides some insight. Momentarily ignoring changes in uncertainty of treatment outcomes, cost-effectiveness decision thresholds should be about 5 times higher for severe Alzheimer's disease than for peptic ulcer disease. Currently, these and every other disease are treated uniformly. These adjustments, of course, vary with the degree of risk-aversion over health, r_H^*. In addition, as we discussed, treatments with highly variable outcomes will have their values reduced from current levels.

Our analysis has implications for both public-sector and private-sector payers. Public-sector payers would benefit from more accurate assessments of value to avoid setting prices that are too high or low from the perspective of encouraging efficient innovation. Private-sector payers could use our approach to match their coverage and reimbursement decisions more closely to the preferences of their beneficiaries. Future research ought to consider how to aggregate our representative consumer model into a payer population of heterogeneous types.

These issues regularly appear in real-world coverage determinations and other policy choices affecting people's access to various healthcare innovations. These appear most clearly in decision processes of the British National Health Service and NICE. While they have an announced CEA threshold of £20,000 to £30,000 per QALY, they have formal exceptions for end of life care, rare diseases, and other circumstances. A separate "Cancer Fund" was established in 2011 to provide access to new cancer drugs that did not meet NICE thresholds (Chambers et al., 2020). These and related "adjustments" signal the inadequacy of the current model, wherein one threshold K is applied to all diseases. Our generalized model makes it abundantly clear why that approach is fundamentally flawed. Severity of illness adjustments are necessary to align decision thresholds with preferences of risk-averse people.

Finally, we note that these results could help to focus R&D efforts and guide design of new technologies in beneficial ways. Different diseases have health consequences varying from relatively small to very large. Interventions that affect the lower end of severity produce less total expected utility gain than those treating patients with the most severe disease, holding constant the magnitude of average improvement (μ_B). A properly designed reimbursement system would alter incentives towards assisting those most-afflicted by illness or injury, whether through chronic disability or acute illness. We believe that such an approach will improve population health efficiently and humanely.

BOX A: THREE WAYS TO INTERPRET PRUDENCE AND TEMPERANCE

Most economists are familiar with the standard measure of risk aversion (Arrow, 1965; Pratt, 1964), but less so (perhaps wholly unfamiliar) with the higher-order terms of prudence and temperance. We discuss these further here.

1. Rates at which absolute and relative risk aversion change. Consider the common question of whether absolute risk aversion is constant (CARA), decreasing (DARA) or increasing (IARA). The economics literature widely assumes that utility is DARA (Gollier, 2001). In parallel comes the question of whether relative risk aversion is constant (CRRA), increasing (IRRA) or decreasing (DRRA). Here the literature is less settled (Meyer and Meyer, 2006). Where relative prudence is $\pi_C^* = -C(\frac{U'''(C)}{U''(C)})$, Appendix 7.3 proves that:

If $(\pi^* - r^*) = 1$ then utility is CRRA,
$(\pi^* - r^*) > 1$ then utility is IRRA,
$(\pi^* - r^*) < 1$ then utility is DRRA.

Concurrently, since $\epsilon(r, M) = \epsilon(r^*, M) - 1$,

If $\pi^* = r^*$ then utility is CARA,
$\pi^* > r^*$ then utility is DARA,
$\pi^* < r^*$ then utility is IARA.

Similar relationships hold between π^* and relative temperance, $\tau^* \equiv -C(\frac{U''''(C)}{U'''(C)})$.

As a useful benchmark, for iso-elastic utility where $r^* = \gamma$, we have that $\pi^* = \gamma + 1, \tau^* = \gamma + 2$, ... For example, when $U(C) = \ln(C)$, $r^* = 1,\ \pi^* = 2, \tau^* = 3$, ...

2. Savings Behavior. In his pioneering work on the concept, Kimball (1990) defines prudence as "the sensitivity of a decision variable to risk." An individual with positive prudence will respond to increases in the variance of future income by saving more today, known as "precautionary savings." Similarly, for $\tau \equiv -\frac{u''''(Y_T)}{u'''(Y_T)}$ (Kimball, 1990, 1992), individuals with positive temperance seek to moderate their total exposure to risk (Kimball, 1992). In risk aversion in QoL, this involves investment in reductions of future risk (e.g., diet, smoking cessation).

3. Mean-Preserving Spreads in Risk. The degree of absolute risk-aversion measures a consumer's distaste for mean-preserving spreads in the distribution of consumption (Rothschild and Stiglitz, 1970). As Eeckhoudt et al. (1995) show, risk-averse but prudent people dislike mean-preserving spreads, but if they must accept one, they prefer that they occur in positive rather than negative outcomes. Prudence represents the strength of their preferences in this respect (Eeckhoudt et al., 1995). This links directly to skewness of outcomes' distributions.

Finally, temperate, prudent, and risk-averse people dislike mean-preserving spreads, prefer that any such spreads apply to more positive outcomes, but derive diminishing marginal utility from successive rightward shifts of mean-preserving spreads of positive outcomes (Eeckhoudt et al., 1995). This links to kurtosis, the "fatness" of tails of distributions compared with normal distributions.

BOX B: Summary of Parameter Definitions

$H_0 =$ Health in base period

$H_{1w} =$ Health in period 1, no illness

$H_{1s} =$ Health in period 1, sick

$\phi =$ probability of illness in period 1

p_1 = probability of survival to period 1

$\mu_H =$ mean health level in sick state (0 = death)

σ_H^2 = variance of health in sick state

$\gamma_{1H} =$ Pearson skewness of health in sick state

$\mu_p =$ mean survival benefit (non-stochastic)

$\mu_B =$ mean QoL treatment benefit

σ_B^2 = variance of treatment benefit

$\gamma_{1B} =$ Pearson skewness of treatment benefit

r_H^* = relative risk aversion in health

π_H^* = relative prudence in health

$\sigma_{H+B}^2 =$ variance of treated patients

γ_{S+B} = Pearson skewness for treated patients

5. REFERENCES

Alatorre C, Carter G, Chen C, Villarivera C, Zarotsky V, Cantrell R, Goetz I, Paczkowski R, Buesching D. A comprehensive review of predictive and prognostic composite factors implicated in the heterogeneity of treatment response and outcome across disease areas. International journal of clinical practice 2011;65; 831-847.

Arrow K. 1976. Welfare analysis of changes in health coinsurance rates. In: Rosett RN (ed.), The role of health insurance in the health services sector. NBER; 1976.

Arrow KJ. 1965. The Theory of Risk-Aversion, Aspects of the Theory of Risk-Bearing (Lecture 2). Essays in the Theory of Risk-Bearing. Markham: Chicago; 1965.

Asaria M, Griffin S, Cookson R, Whyte S, Tappenden P. Distributional cost-effectiveness analysis of health care programmes–a methodological case study of the UK bowel cancer screening programme. Health economics 2015;24; 742-754.

Basu A, Carlson J, Veenstra D. Health Years in Total: A New Health Objective Function for Cost-Effectiveness Analysis. Value in Health 2020;23; 96-103.

Bleichrodt H, Rohde KI, Wakker PP. Koopmans' constant discounting for intertemporal choice: A simplification and a generalization. Journal of Mathematical Psychology 2008;52; 341-347.

Cantrell RA, Alatorre CI, Davis EJ, Zarotsky V, Le Nestour E, Carter GC, Goetz I, Paczkowski R, Sierra-Johnson J. A review of treatment response in type 2 diabetes: assessing the role of patient heterogeneity. Diabetes, Obesity and Metabolism 2010;12; 845-857.

Carter GC, Cantrell RA, Zarotsky V, Haynes VS, Phillips G, Alatorre CI, Goetz I, Paczkowski R, Marangell LB. Comprehensive review of factors implicated in the heterogeneity of response in depression. Depression and anxiety 2012;29; 340-354.

Chambers JD, Silver MC, Berklein FC, Cohen JT, Neumann PJ. Orphan Drugs Offer Larger Health Gains but Less Favorable Cost-effectiveness than Non-orphan Drugs. Journal of general internal medicine 2020.

Chetty R. A new method of estimating risk aversion. American Economic Review 2006;96; 1821-1834.

Córdoba JC, Ripoll M. Risk aversion and the value of life. The Review of Economic Studies 2016;84; 1472-1509.

Cutler DM, McClellan MB, Newhouse JP, Remler D. Are Medical Prices Declining? Evidence from Heart Attack Treatments. Quarterly Journal of Economics 1998;113; 991-1024.

Dohmen T, Falk A, Huffman D, Sunde U, Schupp J, Wagner GG. Individual risk attitudes: Measurement, determinants, and behavioral consequences. Journal of the European Economic Association 2011;9; 522-550.

Easterlin RA. The economics of happiness. Daedalus 2004;133; 26-33.

Ebert S, Wiesen D. Testing for prudence and skewness seeking. Management Science 2011;57; 1334-1349.

Ebert S, Wiesen D. Joint measurement of risk aversion, prudence, and temperance. Journal of Risk and Uncertainty 2014;48; 231-252.

Eckel CC, Grossman PJ. Men, women and risk aversion: Experimental evidence. Handbook of experimental economics results 2008;1; 1061-1073.

Eeckhoudt L, Gollier C, Schneider T. Risk-aversion, prudence and temperance: A unified approach. Economics letters 1995;48; 331-336.

Ehrlich I, Chuma H. A Model of the Demand for Longevity and the Value of Life Extension. Journal of Political Economy 1990;98; 761-782.

Eisenhauer JG, Ventura L. Survey measures of risk aversion and prudence. Applied Economics 2003;35; 1477-1484.

Galama T, Kapteyn A. Grossman's missing health threshold. J Health Econ 2011;30; 1044-1056.

Garber AM, Phelps CE. Economic foundations of cost-effectiveness analysis. J Health Econ 1997;16; 1-31.

Garrison LP, Zamora B, Li M, Towse A. Augmenting Cost-Effectiveness Analysis for Uncertainty: The Implications for Value Assessment—Rationale and Empirical Support. Journal of Managed Care & Specialty Pharmacy 2020;26; 400-406.

Goetz I, Carter GC, Lucero M, Zarotsky V, Alatorre CI, Cantrell RA, Paczkowski R, Sterling KL. Review of treatment response in rheumatoid arthritis: assessment of heterogeneity. Current medical research and opinion 2011;27; 697-711.

Gollier C. The economics of risk and times. MIT Press: Cambridge, Mass.; 2001.

Green C, Gerard K. Exploring the social value of health-care interventions: a stated preference discrete choice experiment. Health Econ 2009;18; 951-976.

Grossman M. On the Concept of Health Capital and the Demand for Health. Journal of Political Economy 1972;80; 223-255.

Hall RE, Jones CI. The Value of Life and the Rise in Health Spending. Quarterly Journal of Economics 2007;122; 39-72.

Harrison GW, Lau MI, Rutström EE. Estimating risk attitudes in Denmark: A field experiment. scandinavian Journal of Economics 2007;109; 341-368.

Hirth RA, Chernew ME, Miller E, Fendrick AM, Weissert WG. Willingness to pay for a quality-adjusted life year: in search of a standard. Med Decis Making 2000;20; 332-342.

Jones BF. The burden of knowledge and the "death of the renaissance man": Is innovation getting harder? The Review of Economic Studies 2009;76; 283-317.

Kimball MS. Precautionary Saving in the Small and in the Large. Econometrica 1990;58; 53-73.

Kimball MS. 1992. Precautionary Motives for Holding Assets. In: Newman P, Milgate M, Eatwell J (eds.), The New Palgrave Dictionary of Money and Finance, vol. III. Macmillan: London; 1992.

Kimball MS. Standard Risk-Aversion. Econometrica 1993;61; 589-611.

Koopmans T. 1972. Representation of Preference Orderings Over Time, Chapter 4 in CB McGuire and R. Radner, Eds., Decision and Organization. (Ed)^(Eds). New York: North-Holland; 1972.

Lakdawalla DN, Doshi JA, Garrison LP, Jr., Phelps CE, Basu A, Danzon PM. Defining Elements of Value in Health Care-A Health Economics Approach: An ISPOR Special Task Force Report [3]. Value Health 2018;21; 131-139.

Lakdawalla DN, Malani A, Reif J. The insurance value of medical innovation. Journal of Public Economics 2017;145; 94-102.

Lakdawalla DN, Romley JA, Sanchez Y, Maclean JR, Penrod JR, Philipson T. How cancer patients value hope and the implications for cost-effectiveness assessments of high-cost cancer therapies. Health Affairs 2012;31; 676-682.

Linley WG, Hughes DA. Societal views on NICE, Cancer Drugs Fund, and Value Based Pricing Criteria for Prioritising Medicines: A Cross-Sectionalk Survey of 4118 Adults in Great Britain. Health Economics 2013;22; 948-964.

Lising A, Rosner A, Gladman J, Drummond M. 2016. Payers' Use of ICER Reports in Decision Making. (Ed)^(Eds), Association of Managed Care Pharmacy Nexus. Journal of Managed Care and Specialty Pharmacy: National Harbor, Maryland; 2016.

Markowitz H. Portfolio selection. The journal of finance 1952;7; 77-91.

Maxwell A, Özmen M, Iezzi A, Richardson J. Deriving population norms for the AQoL-6D and AQoL-8D multi-attribute utility instruments from web-based data. Quality of Life Research 2016;25; 3209-3219.

Meltzer D. Accounting for future costs in medical cost-effectiveness analysis. J Health Econ 1997;16; 33-64.

Meltzer DO, Smith PC. 2011. Theoretical issues relevant to the economic evaluation of health technologies. Handbook of health economics, vol. 2. Elsevier; 2011.

Merton RC. Optimum Consumption and Portfolio Rules in a Continuous-Time Model. Journal of Economic Theory 1971;3; 373-413.

Meyer DJ, Meyer J. Measuring risk aversions. Now publishers inc; 2006.

Murphy KM, Topel RH. The Value of Health and Longevity. Journal of Political Economy 2006;114; 871-904.

Muurinen JM. Demand for health: a generalised Grossman model. J Health Econ 1982;1; 5-28.

Nord E, Pinto JL, Richardson J, Menzel P, Ubel P. Incorporating societal concerns for fairness in numerical valuations of health programmes. Health economics 1999;8; 25-39.

Nord E, Richardson J, Street A, Kuhse H, Singer P. Maximizing health benefits vs egalitarianism: an Australian survey of health issues. Soc Sci Med 1995;41; 1429-1437.

Noussair CN, Trautmann ST, Van de Kuilen G. Higher order risk attitudes, demographics, and financial decisions. Review of Economic Studies 2013;81; 325-355.

Phelps CE. A new method to determine the optimal willingness to pay in cost-effectiveness analysis. Value in Health 2019a;22; 785-791.

Phelps CE. When Opportunity Knocks, What Does It Say? Value Health 2019b;22; 851-853.

Phelps CE, Cinatl C. 2020. A Generalized Method for Estimating Optimal Willingness to Pay Thresholds for Health Care Technology Assessment. (Ed)^(Eds). University of Rochester: Rochester, NY; 2020.

Pratt JW. Risk Aversion in the Small and in the Large. Econometrica 1964;32; 122-136.

Rosen S. Valuing Health Risk. American Economic Review 1981;71; 241-245.

Rosen S. The Value of Changes in Life Expectancy. Journal of Risk and Uncertainty 1988;1; 285-304.

Rothschild M, Stiglitz JE. Increasing risk: I. A definition. Journal of Economic theory 1970;2; 225-243.

Weinstein MC, Stason WB. Foundations of cost-effectiveness analysis for health and medical practices. N Engl J Med 1977;296; 716-721.

Yu Y, Cui J. Present and future of cancer immunotherapy: A tumor microenvironmental perspective. Oncology letters 2018;16; 4105-4113.

6. TABLES

Table 1: Summary of variance and skewness effects on economic welfare

	$\Delta\sigma^2 > 0$	$\Delta\sigma^2 < 0$
$\Delta\gamma_1 > 0$	*Uncertain* Consumer dislikes higher variance but values higher skewness	*Good* Consumer values lower variance and higher skewness
$\Delta\gamma_1 < 0$	*Bad* Consumer dislikes higher variance and lower skewness	*Uncertain* Consumer values lower variance but dislikes lower skewness

Table 2: R-Multipliers for alternative values of r_H^* and ℓ^*.

	r_H^*										
ℓ^*	**0**	**0.25**	**0.5**	**0.7**	**0.8**	**0.9**	**1**	**1.1**	**1.2**	**1.3**	**1.5**
0	1	1	1	1	1	1	1	1	1	1	1
0.1	1	1.03	1.05	1.08	1.09	1.1	1.11	1.12	1.13	1.15	1.17
0.3	1	1.09	1.2	1.28	1.33	1.39	1.43	1.48	1.53	1.59	1.71
0.5	1	1.19	1.41	1.62	1.74	1.87	2	2.14	2.3	2.46	2.83
0.7	1	1.35	1.83	2.32	2.62	3.01	3.33	3.76	4.24	4.78	6.09
0.9	1	1.78	3.15	5.01	6.29	7.91	9.95	12.51	15.73	19.79	31.2

Notes: These are estimates from a 50-term Taylor Series expansion. They converge rapidly for values of $\ell^* < 0.85$ and have small downward bias for the row for $\ell^* = 0.9$. For example, for $r_H^* = 1$, the infinite series value is 10. All relative risk aversion measures are constant, so $\pi^* = r^* + 1, \tau^* = r^* + 2, ...$

Table 3: Baseline QoL levels for a set of example diseases.

ℓ^*	Example Diseases
0.0-0.1	Peptic Ulcer Disease; Stress Urinary Incontinence; Benign Prostatic Hyperplasia
0.1-0.2	Grave's Disease; Hypertension with no complications; Sleep Apnea
0.2-0.3	Familial Hypercholesterolemia; Peripheral Arterial Disease; End-Stage Knee Osteoarthritis
0.3-0.5	Type 1 Diabetes; Acute Lung Injury; Moderate to Severe Rheumatoid Arthritis
0.5-0.7	TIA[a] and Carotid Stenosis; TBI[b]; Nursing Home Residents at Risk of Pressure Ulcers
0.7-1.0	Alzheimer's disease; Metastatic Colorectal Cancer; Acute Pulmonary Embolism

[a] Transient ischemic attack

[b] Traumatic brain injury

Notes: Example diseases are taken from the Tufts Cost-Effectiveness Analysis Registry.

7. APPENDIX

Here, we discuss the convergence properties of Taylor Series expansions in our application.

7.1. Convergence of Taylor Series for HARA Utility

Proposition 1: Suppose the health index has support on the interval [0,1], and the utility function takes the form $V(c,H) = U(c)W(H)$, where W belongs to the class of HARA (hyperbolic absolute risk-aversion) utility functions. If H is a random variable with support [0,1] and mean μ_0, then the Taylor expansion of $W(H)$ around μ_0 converges.

Proof: Any utility function in the HARA family can be written as $U(H) = \frac{1-\gamma}{\gamma}\left(\frac{\beta H}{1-\gamma}+\eta\right)^{\gamma}$ (Merton, 1971). The Taylor expansion around the mean of H, μ_0, takes the form:

$$EU(H) \approx U(\mu_0) + U'(\mu_0)E(H-\mu_0)) + \frac{1}{2!}U''(\mu_0)E(H-\mu_0)^2 + \frac{1}{3!}U'''(\mu_0)E(H-\mu_0)^3 + \cdots$$

Since H has support on the unit interval, it is evident that $\lim_{n\to\infty}\frac{E(H-\mu_0)^{n+1}}{E(H-\mu_0)^n} = 0$. Moreover, defining the nth derivative of U as $U^{(n)}$, we can write:

$$\frac{U^{(n+1)}}{U^{(n)}} = \frac{\beta^{n+1}(\gamma-2)\ldots\frac{\gamma-n}{(1-\gamma)^{n-1}}\left(\frac{\beta\mu_0}{1-\gamma}+\eta\right)^{\gamma-(n+1)}}{\beta^{n}(\gamma-2)\ldots\frac{(\gamma-(n-1))}{(1-\gamma)^{n-2}}\left(\frac{\beta\mu_0}{1-\gamma}+\eta\right)^{\gamma-n}} = \frac{\beta(\gamma-n)}{\left(\frac{\beta\mu_0}{1-\gamma}+\eta\right)(1-\gamma)}$$

Thus, L'Hôpital's rule implies that $\lim_{n\to\infty}\frac{U^{(n+1)}}{(n+1)U^{(n)}} = \lim_{n\to\infty}\frac{\beta(\gamma-n)}{(n+1)\left(\frac{\beta\mu_0}{1-\gamma}+\eta\right)(1-\gamma)} = \frac{-\beta}{\left(\frac{\beta\mu_0}{1-\gamma}+\eta\right)(1-\gamma)}$.

Since β and γ are both finite scalars, since $0 \le \mu_0 \le 1$, and since $\lim_{n\to\infty}\frac{E(H-\mu_0)^{n+1}}{E(H-\mu_0)^n} = 0$, it follows that:

$$\lim_{n\to\infty}\frac{U^{(n)}E(H-\mu_0)^{n+1}}{(n+1)U^{(n+1)}E(H-\mu_0)^n} = 0$$

This proves the claim.

7.2. Assessing Convergence of the Taylor Series Expansion

Next, we assess the convergence characteristics of our Taylor Series expansion using four-parameter Beta distributions with varying skewness and kurtosis. We calculated the absolute value of the second, third and fourth-order Taylor series terms to compare various technologies (T) with a common control (S), where the control intervention always has a four parameter Beta distribution with parameters $\alpha = 2, \beta = 2$ and support that is 0.5 units wide.

Similarly, the T distributions feature $\alpha = 2$ and vary β between 5 and 10 in increments of 1, with 0.5-unit width of support. Each increase in β increases the skewness and alters the kurtosis in complex ways (since β appears in both the numerator and denominator of the expression for kurtosis).

We characterize speed of convergence by calculating ratios between the 3rd-order and 2nd-order Taylor Series terms, and between 4th-order and 2nd-order Taylor Series terms. Our simulations illustrate that the speed of convergence varies with the average level of health in the untreated sick state (μ_H).

Figure A-2a shows the ratio of the 3rd to the 2nd order terms for distributions of T with successively increasing skewness (β ranging from 5 to 10 in increments of 1, with the lowest skewness in the bottom line, increasingly for higher lines).

As predicted, convergence is always faster when average health in the untreated sick state is higher. In Figure A-2a (the ratio of 3rd to 2nd order terms) the ratio is about 0.32 for $\mu_H = 0.1$, rapidly falling below 0.15 as μ_H exceeds 0.2.

Figure A-2b similarly graphs the ratio of the 4th order (kurtosis-related) term to the 2nd order term. There, the lines reverse in sequence, with the smallest β parameters in the higher lines. The ratios are all below 0.75 and rapidly fall below 0.17 for values of μ_H exceeding 0.2. These simulations demonstrate a reasonably rapid rate of convergence with these specific parameters, particularly when we limit the analysis to values of $\mu_H \geq 0.2$.[30] Since QALY values below 0.2 correspond to extremely poor health states, e.g., for a person with an untreatable and highly aggressive cancer, this is a helpful result.

Figure A-2 also demonstrates that statistical moments produce uncertain effects on speed of convergence. In our four-parameter beta distributions, higher positive skew (larger values of β) results in slower convergence for a third-order Taylor expansion, but faster convergence for a

[30] The width of the support interval affects
Figure A-2. As the support interval narrows for any 4-parameter Beta distribution, the variance falls with the square of that interval's width. Since
Figure A-2a contains σ^3 in the calculations of the 3rd moment (and similarly Figure 5b contains σ^4 to calculate the 4th moments), the ratios shown in these figures become smaller as the support interval narrows. The effect is linear in the support width in
Figure A-2a and quadratic in
Figure A-2b. The most pessimistic convergence would occur for support intervals near 1.0.

fourth-order expansion. Nonetheless, our calculations illustrate how practitioners can readily calculate speed of convergence using their specific statistical moments and parameter values.

Sensitivity analyses (not shown) identify situations where the higher-order terms have ratios of the 3rd to 2nd order terms exceed 1.0, indicating slower convergence. The most important of these cases occurs when the two distributions have nearly identical skewness, so that their difference is very small, and hence the ratio of the 2nd order to the 3rd order term relatively large. Even in these cases, it requires a relatively high degree of variance, which further increases the 3rd and 4th moment values. A similar phenomenon occurs when the two distributions have nearly identical values of kurtosis.

Figure A-2. Speed of convergence in Taylor Series approximation to value of technology.

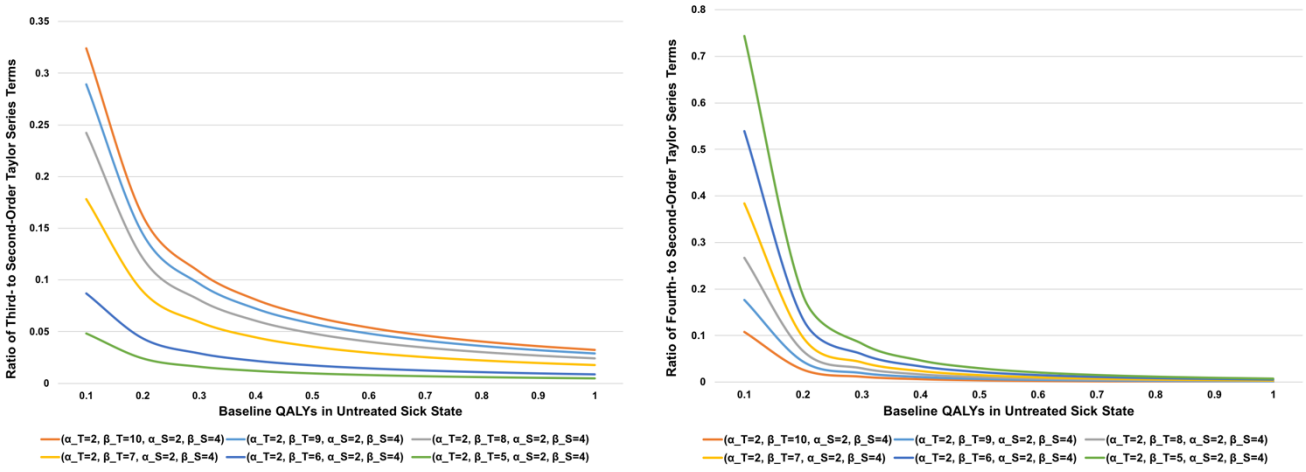

Figure A-2a Figure A-2b

Notes: These figures demonstrate the rate of convergence of the Taylor series by taking the ratio of the 3rd-order to the 2nd-order Taylor Series terms (Figure A-2a) and the 4th-order to the 2nd-order terms (Figure A-2b). Smaller ratios, and y-axis values, demonstrate faster convergence. In both panels, the control group (S) outcomes have Beta distribution parameters of $\alpha = 2; \beta = 4$. In Figure A-2a, the topmost curve has the greatest skewness ($\alpha = 2; \beta = 10$), with the value of β declining (serially) in the lower curves by one unit (10, 9, 8, 7, 6, 5) holding α constant at 2. Therefore, the skewness declines as one moves from the top to the bottom curves. In Figure A-2b, the sequence is reversed (following the formula for kurtosis in the 4-parameter Beta distribution). Apart from this difference, Figure A-2a and Figure A-2b share a similar structure.

7.3.Relative Risk Preferences and Income

Consider the basic question of how r^* changes with consumption—whether r^* is IRRA, CRRA or DRRA. And, consider the related question of how r_H^* changes with health. Since the analysis for both these cases is identical, we will present only the case of r_C^* and consumption. Define the elasticity of r^* with respect to consumption as $\epsilon(r_C^*, C)$. Following Kimball (1999):

$$r_C = -\frac{U''}{U'} \text{ and } r_C^* = rC \tag{24}$$

$$\pi_C = -\frac{U'''}{U''}, \text{ and } \pi_C^* = \pi C \tag{25}$$

Here, r_C^* is relative risk aversion and π_C^* is relative prudence, as defined by Kimball (1993).

The derivative of r_C^* with respect to C is:

$$\frac{dr_C^*}{dC} = -\frac{U''}{U'} + C\left[\frac{[U'U''' - U''U'']}{[U'U']}\right] = r_C + C\left(\frac{U'''}{U'} + r_C^2\right) \tag{26}$$

Hence the elasticity of r_C^* with respect to C is:

$$\epsilon(r_C^*, C) = 1 + C\left[\frac{U'''}{U''} + r_C\right] = 1 - (\pi_C^* - r_C^*) \tag{27}$$

or

$$\epsilon(r_C^*, C) = 1 - (\pi_C^* - r_C^*) \tag{28}$$

From this, we can infer that if:

$$(\pi_C^* - r_C^*) = 1, \text{ then utility is CRRA} \tag{29}$$

$$(\pi_C^* - r_C^*) > 1, \text{ then utility is DRRA} \tag{30}$$

$$(\pi_C^* - r_C^*) < 1, \text{ then utility is IRRA} \tag{31}$$

Similarly, since $\epsilon(r_C, C) = \epsilon(r_C^*, C) - 1$, if:

$$\pi_C^* = r_C^*, \text{ then utility is CARA} \tag{32}$$

$$\pi_C^* > r_C^*, \text{ then utility is DARA} \tag{33}$$

$$\pi_C^* < r_C^*, \text{ then utility is IARA} \tag{34}$$

Condition (34) is widely considered as empirically implausible, all known estimates of π_C^* and r_C^* indicate that utility is DARA (condition (33)).

An identical relationship exists between relative prudence and relative temperance when the utility function is differentiable four times and similarly for all higher-order ratios of utility-function derivatives. Again following Kimball (1999), define relative temperance as

$$\tau_C^* = C\,\frac{U''''}{U'''} \tag{35}$$

$$\epsilon(\pi_C^*, C) = 1 - (\tau_C^* - \pi_C^*) \tag{36}$$

Following the same method used to derive Equation (27), we can solve Equation (36) for τ_C^*:

$$\tau_C^* = 1 + \pi_C^* - \epsilon(\pi_C^*, C) \tag{37}$$

Thus if π_C^* is constant (i.e., $\epsilon(\pi_C^*, C) = 0$) then $\tau_C^* = 1 + \pi_C^*$. The same relationships hold for all higher-order risk parameter similar to those defined by Kimball (1999).