Aalto University
School of Business

Bachelor's Programme in Economics

Effects of risk preferences on the value of healthcare interventions

Assessing cost-effectiveness analysis tools in the medical field

Cristian Musio

Bachelor's Thesis
2025

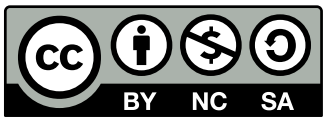

Author Cristian Musio

Title Effects of risk preferences on the value of healthcare interventions — Assessing cost-effectiveness analysis tools in the medical field

Degree programme Economics

Major Economics

Advisor Pauli Murto

Date 28 April 2025 **Number of pages** 32+9 **Language** English

Abstract

This bachelor thesis delves into the effects of risk preferences on the value of healthcare interventions. It consists of a literature review comparing the traditional Cost-Effectiveness Analysis (CEA) approach and the more novel approach proposed by Lakdawalla and Phelps (2021), the Generalized Risk-Adjusted Cost-Effectiveness model (GRACE). The two are compared both through a theoretical lens and through a review of empirical estimates. Furthermore, the thesis includes a theoretical case illustration meant to compare the two approaches, through the use of a Markov model coded in R, with the support of the heemod package. In this illustrative case, I consider the disease of multiple sclerosis and its treatments, using Australian data. I compare the cost-effectiveness results under CEA and under GRACE, to highlight that the latter is a more accurate methodology. I then perform a probabilistic sensitivity analysis with a Monte Carlo simulation, meant to be tied together with the insights from the theoretical discussion and those from the literature review.

From these insights, I provide implications and suggestions for policy decisions in the field of healthcare, specifically in the assessment of the value of medical interventions. The results of this work suggest that severely ill patients and disabled patients are discriminated against by conventional methods, and that some treatments that could add value to society, specifically to patients, have unwarranted restricted use. Moreover, risk preferences appear to play a key role in influencing the value-assessment of such treatments. Thus, I suggest that GRACE should replace CEA as the new standard, and that more research in this field should be conducted, as I believe it would increase societal welfare.

Keywords Risk preferences, healthcare evaluations, welfare, GRACE, CEA

Contents

Abbreviations

AUD	Australian Dollar
CEA	Cost Effectiveness Analysis
$CEAC$	Cost Effectiveness Acceptability Curve
DMT	Disease-Modifying Treatment
FDA	Food and Drug Administration (USA)
GBP	Great Britain Pound
$GRACE$	Generalized Risk-Adjusted Cost-Effectiveness
H	Health Level
$HRQoL$	Health-Related Quality of Life
$ICER$	Incremental Cost-Effectiveness Ratio
$ICER_{ORG}$	Institute for Clinical and Economic Review (USA)
K	Threshold
LE	Life Expectancy
MS	Multiple Sclerosis
MRS	Marginal Rate of Substitution
MUH	Marginal Utility of Health
MUC	Marginal Utility of Non-Health-Related-Consumption
$NICE$	National Institute for Health and Care Excellence (UK)
$PBAC$	Pharmaceutical Benefits Advisory Committee (Australia)
$PPMS$	Primary Progressive Multiple Sclerosis
PSA	Probability Sensitivity Analysis
$QALY$	Quality-Adjusted Life-Years
$RRMS$	Relapsing Remitting Multiple Sclerosis
$SPMS$	Secondary Progressive Multiple Sclerosis
USD	United States Dollar
WHO	World Health Organization
WTP	Willingness To Pay

1 Introduction

Every healthcare treatment, whether it is a new drug or a new medical technology, needs an economic assessment to determine its benefits to society and its related costs. This is of great importance, because resources are limited and thus their efficient allocation has direct effects on societal welfare.

As described by Garber and Phelps (1997), the standard way to perform such assessment is through a Cost-Effectiveness Analysis (CEA), which takes as input the costs of the medical intervention and its health benefits, in order to output a variable representing the incremental costs per incremental health benefits. This variable is commonly referred to as the Incremental Cost Effectiveness Ratio (ICER).

The traditional CEA approach comes with strong assumptions: an incremental health benefit of 1 unit is always worth the same regardless of patient's illness-severity. This is in contrast to conventional economics literature, that states that if something is scarce (in this case health) it should be valued more. Moreover, CEA assumes that all patients are risk neutral. Therefore, potential differences in risk preferences are not accounted for in the analysis.

This thesis challenges these assumptions and thus aims to answer the following research question: "how does accounting for risk preferences change the valuation of healthcare treatments, and what is the resulting impact on patient welfare?".

To understand how these concepts relate to each other, it is important to present the context that raised the research question, which is the goal of this introduction. As previously mentioned, CEA returns a variable ICER = $\frac{\Delta\text{Cost}}{\Delta\text{Health}}$. One first challenge arises, which is how these health benefits (ΔHealth) should be measured. In fact, treatments could improve the patient's life in a multitude of ways: by improving their survival (that is, adding years to life) or by improving their Health-Related Quality of Life (HRQoL).

One might argue that it would not be appropriate to value five extra years in perfect health the same as five extra years in mediocre health. This is why CEA measures health benefits the following way: each extra year of life is multiplied by a factor between 0 and 1, where 0 represents a QoL equal to death, and 1 represents perfect health. The result is a Quality Adjusted Life Year (QALY). Thus, 5 years in perfect health equal 5 QALY, and 5 years in mediocre health equal, for example, 5*0.5=2.5 QALY.

Costs are expressed the conventional way, so in terms of monetary currency (e.g., dollars). Therefore, the ICER is expressed in costs/QALY (e.g., dollars/QALY).

The ICER is then compared to a threshold K determined by decision-making health authorities, with K also expressed in costs/QALY. Treatments that satisfy ICER<K are said to be cost-effective. Therefore, K can be considered as the willingness-to-pay (WTP) for an incremental QALY.

The cost-effectiveness of a treatment is not the only factor that determines whether a treatment enters the market, but it influences the decision-making process regarding potential investments into the treatment, reimbursement rates, and prices. For instance, if a potential drug has high costs/QALY then investors may see it as "not worth it" to pursue. Similarly, public decision-makers might consider that such drug should not have a high reimbursement rate, since it only provides a low amount of QALY per money spent, implying that money would be better spent on other treatments. For both actors, the potential drug would be deemed as having low "value for money". This concept of "value for money" is central in the field of health economics.

One can imagine that if the CEA approach were flawed, that would have drastic consequences on the well-being of society. For instance, that would mean that a drug that is actually cost-effective and beneficial to society might be wrongly deemed as not cost-effective, thus limiting access to it for patients who would need it. From a societal perspective, this would represent a deadweight loss, and from an equality perspective it would also be problematic because it would mean that individuals with low QoL are affected more than those in perfect health.

Those concerns are well-founded, as indeed the CEA approach comes with its shortcomings, which can make those hypothetical scenarios a reality. As Lakdawalla and Phelps (2021) point out, researchers have suggested that CEA discriminates against the severely ill and the disabled. The same authors suggest that CEA can be made more accurate by accounting for the diminishing returns on health (i.e., improvements in health are worth more to sicker patients) and for the difference in risk preferences that patients might have. They suggest that the health states influence risk preferences, which in turn influence health valuations. Thus, the model they propose is called the Generalized Risk-Adjusted Cost-Effectiveness framework (GRACE).

In this thesis, GRACE is the bedrock upon which my arguments are based on. My intention is to provide a literature review highlighting how GRACE addresses the shortcomings of the traditional CEA, and specifically how risk preferences play a crucial role. However, the literature review is only part of the thesis. In fact, I also provide a basic tool that can be used to compare CEA and GRACE. This tool is based on the "heemod" (Health Economic Evaluation Modelling) R package developed by Filipović-Pierucci et al. (2017). My goal is to use heemod to create a model that

performs an analysis of a hypothetical medical treatment under CEA assumptions, and then under GRACE assumptions. Subsequentially, compare the results. Although simplistic, the model allows for more complexity and nuances to be introduced. This approach serves as both a way to demonstrate how the mechanism behind GRACE works, and also as a way to provide supporting evidence to the claims of Lakdawalla and Phelps (2021), who suggest that patients' risk preferences play a key role and should be taken into account when assessing the value of healthcare treatments to society. As a side note, it is important to point out that the software used to conduct real CEAs is proprietary and only accessible to authorities and select researchers (including those cited). On the other hand, R is freely accessible, which is why I'm using it.

Through the literature review and the accompanying model analysis, my findings suggest that disease severity impacts the risk profile of patients, making gains for sicker patients more valuable than for healthier patients. Moreover, risk-averse patients might still value risky treatment choices if there is a probability (however small) to gain considerable health benefits. This is the "value of hope" described by Lakdawalla and Phelps (2021). These findings suggest that GRACE is a more accurate assessment method than traditional CEA, and that risk preferences play a crucial role on the perceived value of healthcare interventions, thus affecting society's welfare.

The thesis is structured as follows: section 1 is the introduction; section 2 presents the CEA approach and its limitations; section 3 presents the GRACE approach and how it addresses the shortcomings of CEA; section 4 compares CEA and GRACE through a case illustration for a hypothetical drug treating multiple sclerosis in Australia; section 5 provides a discussion of policy implications; finally, section 6 serves as the summary and conclusion.

2 CEA approach and its limitations

This section contextualizes the foundations of the traditional CEA, presents its methodological framework, and its limitations. This is helpful in answering the research question — that is, examining the effects of risk preferences on the valuation of healthcare treatments — as highlighting the CEA limitations brings to light gaps that can be filled by a risk-accounting model (such as GRACE).

2.1 Background

Many countries around the world, including the US and the UK, employ the CEA approach as a way to conduct assessments on the value of healthcare interventions. The results of such analyses help decision-makers evaluate the desirability of a given medical treatment against the available alternatives.

This approach was initially formalized by Garber and Phelps (1997) to address a specific issue at the time: analysts would deliver the ICER value (i.e., marginal cost to marginal benefit ratio) of a treatment, and live it up to social and political institutions to subjectively conclude whether the ICER fell below an acceptable threshold. In other words, the analysts would avoid explicitly providing a threshold recommendation. This in turn meant that there was no consistency or systematic process in assessing which treatments were cost-effective, as the threshold effectively used were arbitrary and devoid of justification.

Thus, in their paper Garber and Phelps (1997) present an expression that can be used to systematically deduce the optimal threshold level, based on the microeconomic principles of utility maximization. In the next section, I proceed to discuss said expression.

2.2 The optimal threshold

As previously mentioned, a key component to consider in CEA is the threshold K used to determine whether a treatment is cost-effective or not. In other words, K can be considered as the willingness-to-pay (WTP) for a QALY. Thus, treatments whose ICER is below this level represent "good value for money", since they provide a QALY at a cheaper cost than society's WTP. Such treatments are expanded in use, meaning they are offered to more patients or are used more frequently. Conversely, treatments with an ICER above the threshold are restricted in use or have their price reduced.

Therefore, it is of main interest to find the optimal threshold that maximizes

utility. To do so, Garber and Phelps (1997) assume a two-period model in which an individual has health H_0 in period 0, H_1 in period 1, with $0 \leq H \leq 1$. This value of health (H) is assumed to be a function of two inputs a and b (e.g., health treatments), meaning that $H = H(a, b)$. Inputs a and b have respective prices of p_a and p_b. Furthermore, the model considers non-health-related consumption C (that is income minus health-related consumption), utility of consumption $U(C)$ with diminishing returns, and total utility $V(H, C)$ that depends on C and H, such that $V(C, H) = U(C)H$. Meaning that total utility has constant returns to health. The authors assume that generally $H_0 = 1$.

The threshold K (or the WTP) should be equal to $K = MUH/MUC$, where MUH is the marginal utility of health, and MUC the marginal utility of consumption. Thus:

$$\text{MUH} = \frac{\partial V}{\partial H} = U(C)$$
$$\text{MUC} = \frac{\partial V}{\partial C} = U'(C)H$$

Assuming baseline health $H_0 = 1$, the optimal threshold is:

$$K = \frac{MUH}{MUC} = \frac{U(C)}{U'(C)} \tag{1}$$

From there, the authors define the elasticity of utility with respect to C such that $w_C \equiv \frac{U'(C)C}{U(C)}$. This tells how much utility changes when consumption changes. By rearranging the definition it gives $U(C) = \frac{U'(C)C}{w_C}$. Then, plugging it into equation (1) implies:

$$K = \frac{C}{w_C} \tag{2}$$

Insurance coverage policies are chosen in period 0, prior to knowing illness states, and they influence access to care in period 1. Therefore, expected utility is maximized when $K \geq \frac{p_a}{H'_a} = \frac{p_b}{H'_b}$. The right-hand side are essentially the prices of each input per marginal health benefit. If, for example, during one period input a has better "value for money" than input b, then resources are reallocated to a. These readjustments are done until the two inputs have the same value for money.

Empirically, it is estimated that $0.3 \leq w_c \leq 0.5$ and that $2C \leq K \leq 3C$ in industrialized countries (Phelps, 2019). This translates to roughly $110000 \leq K \leq 165000$ USD/QALY in the US in 2019 (Lakdawalla and Phelps, 2021).

According to the same authors, this range is within the recommendations provided by the competent authority in charge in the US, that is the Institute for Clinical and

Economic Review (whose acronym is also $ICER$, so to avoid confusion with the previously discussed ratio it is noted $ICER_{ORG}$), which recommends $50000 \leq K \leq 200000$ USD/QALY. In the literature, common values considered are $100000 \leq K \leq 200000$, with $K = 150000$ being the most common value, referred to as the "central value" in the US (Lakdawalla et al., 2024).

For comparison purposes, it is important to highlight how much these thresholds can change from country to country, as different healthcare systems distribute the burden of costs differently. For instance, in the UK the responsible body for recommendations is the National Institute for Health and Care Excellence (NICE) uses an official range of $20000 \leq K \leq 30000$ GBP/QALY, going as low as 13000 in certain cases (Lakdawalla and Phelps, 2021).

In the case of certain rare or severe diseases such as cancer exceptions can be made. For instance, the $ICER_{ORG}$ can consider a treatment cost-effective even up-to $K = 300000$ in the case of cancer. Furthermore, other types of exceptions can be made, or even ban the use of CEA altogether. One such instance is the US Affordable Care Act, which bans the use of CEA if the result is deemed to be discriminatory towards disabled patients and limits its use in Medicare evaluations and in other federal-related health programmes.

The fact that such exceptions and bans exist hints at the possibility that the CEA approach might be flawed. Therefore, in the next section I discuss the limitations of the CEA that could explain why such scenarios arise.

2.3 Limitations of CEA

The main limitation of the CEA approach is that it assumes constant returns to scale for health. It does not consider untreated illness severity, preexisting level of disability, or prognostic uncertainty (i.e., the uncertainty of future treatment benefits). In other words, a QALY is always worth the same, no matter who receives it. This is the famous motto of the CEA: "a QALY is a QALY" (Williams, 1992). So, a treatment that improves the life of a cancer patient by 0.2 QALY is considered equivalent to a treatment that improves the life of a mildly-ill patient by the same amount.

I argue that the cancer patient would value those 0.2 QALY higher than the mildly-ill patient, since the cancer patient has a scarcer amount of baseline health, and when something is scarce, it is usually valued more. This reasoning is consistent with the decision of the $ICER_{ORG}$ to increase K to 300000 for cancer patients. In fact, it implies that the organization recognizes that even if a treatment is considered

"expensive" under traditional conditions, if a disease is severe enough then it is still worth it (from a societal perspective) to pursue investment in such treatment.

Moreover, the CEA assumption of constant returns to scale for health also implies that the marginal rate of substitution (MRS) between gains in health-related quality of life (HRQoL) and life expectancy (LE) are always the same in all situations (Lakdawalla and Phelps, 2023). However, I believe that is not necessarily the case, because I hypothesize that the preferences between HRQoL and LE depend on many factors, including the present health status of the individual. For instance, according to CEA, extending life by 1 year for an individual in perfect health is equal to 1 QALY, and extending life by 2 years for an individual in 0.5 health is worth 2*0.5=1 QALY as well. My argument is that in the latter case, the individual might value 1 extra year more than chronic pain relief. This might be the case for example for a terminal cancer patient. By the same reasoning, an individual that is risk-averse might prefer an increase in QoL over an uncertain gain in LE. Thus, by not accounting for these nuances, I believe that CEA undervalues treatments for severe diseases and high-risk populations.

To address these concerns, I will discuss in the next section how implementing risk preferences into the CEA model might provide a more accurate and realistic representation of the value that individuals (and society) place on healthcare interventions.

3 GRACE approach and the role of risk preferences

In this section, I first provide the theoretical foundation of GRACE, then I present how these affect the framework of CEA.

3.1 Theoretical foundation

The GRACE approach introduces the idea of positive diminishing returns to scale to health, as opposed to the constant returns in the case of CEA. In other words, the total utility that was previously expressed as $V(H, C) = U(C)H$ under CEA becomes $V(H, C) = U(C)W(H)$ under GRACE, with $W(H)$ being the utility of health (Lakdawalla and Phelps, 2023).

As the authors explain, this change yields a new set of implications:

1. Under GRACE, the more severe the untreated illness is, the greater the WTP for health-related quality of life (HRQoL) improvements
2. Under GRACE, the more severe the disability status, the greater the WTP for both HRQoL and LE improvements
3. Under traditional CEA, health improvements are systematically overvalued, with the degree of overvaluation depending on how quickly marginal utility declines as health increases
4. Under traditional CEA, treatments for mild conditions are overvalued and treatments for severe conditions and disabled populations are undervalued
5. Under GRACE, contrary to CEA, the relative value of LE and HRQoL depends on patients' initial health status
6. Risk averse patients value treatment gains less, the more outcomes for HRQoL are uncertain. Similarly, they value treatment gains more if the probability of unusually beneficial outcomes increases

Mathematically, the GRACE assumptions can be expressed as follows:

$$\text{MUH} = \frac{\partial V}{\partial H} = W'(H)U(C)$$
$$\text{MUC} = \frac{\partial V}{\partial C} = U'(C)W(H)$$
$$\text{MRS} = \frac{MUH}{MUC} = \frac{W'(H_0)U(C_0)}{U'(C_0)W(H_0)}$$

A new variable is considered, the elasticity of utility with respect to health, defined as $w_H \equiv \frac{W'(H_0)H_0}{W(H_0)}$. That is, the utilities of health are functions of the baseline health state (at period 0). Plugging in that gives:

$$\text{MRS} = \frac{C_0}{H_0}\frac{w_H}{w_C}$$

Essentially, CEA is a restricted case of GRACE that assumes $w_H = 1$ (linear utility), whereas GRACE allows for $0 < w_H \leq 1$. Now, let $K*$ be the threshold under simplified GRACE, then: $K^* = Kw_H$. This implies $K^* \leq K$ (since $0 < w_H \leq 1$). Moreover, it shows that CEA overvalues gains in HRQoL compared to simplified GRACE. The latter is referred to as the simplified version of GRACE, because GRACE should also consider illness severity, which I discuss in the next section.

3.2 Impact of illness severity

To model the impact of illness severity, Lakdawalla and Phelps (2023) consider a new variable R representing the ratio of marginal utilities in the average untreated sick state to the healthy state. That is, $R = \frac{W'(\mu_H)}{W'(H_0)}$ with μ_H being the average untreated sick state. This variable indicates how much more society would value health gains for sick individuals compared to healthy ones. Thus, the optimal threshold (or WTP) according to GRACE becomes:

$$K_{GRACE} = K^*R = Kw_HR$$

As illness severity increases R, increases exponentially. On the other hand, for mild conditions R is close to 1.

3.3 Impact of disability status

The same authors model the impact of disability in a similar fashion, with a variable named D, such that $D = \frac{W(H_0)}{W(H_D)}$ where $H_D = H_0(1-d^*)$ with d^* being the proportional loss created by the permanent disability.

Therefore, when accounting for disability status, the optimal WTP under GRACE becomes:

$$K^D_{GRACE} = Kw_HRD$$

This explains why disabled individuals are discriminated against under CEA. To

better illustrate the magnitude of this effect I include figures from Lakdawalla and Phelps (2023). In Figure 1, l^* is the proportional loss in untreated health in period 1.

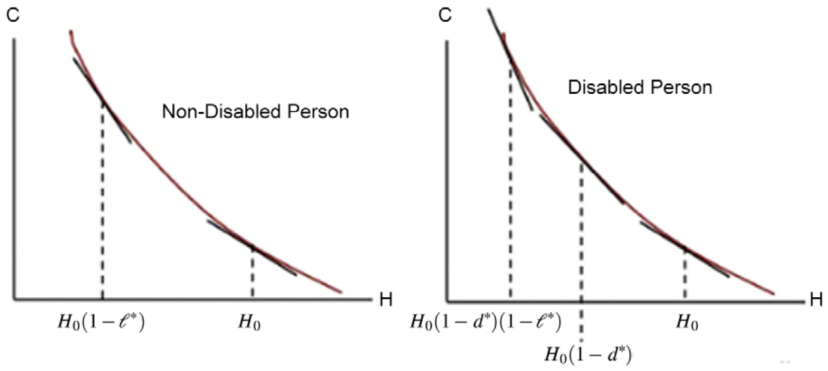

Figure 1: As illness severity and disability increase, the MRS between C and HRQoL increases (Lakdawalla and Phelps, 2023)

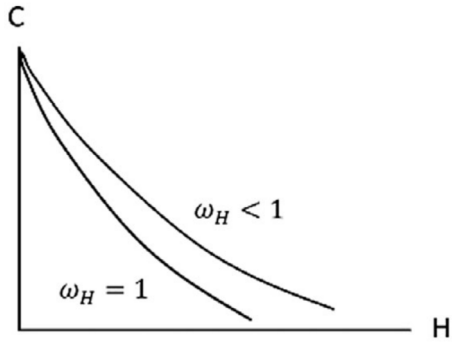

Figure 2: WTP for health improvements decreases as utility elasticity of health decreases, at any given illness severity level (Lakdawalla and Phelps, 2023)

In the next section, I introduce the effects of risk preferences.

3.4 Constant risk aversion

The GRACE framework, contrary to CEA, allows the risk aversion coefficient to be implemented. There are many forms of possible utilities, and Lakdawalla and Phelps (2023) provide an example with a constant relative risk aversion function (CRRA). In such function, the risk aversion coefficient is given by $r_H^* = (1-\gamma)$. So, the utility of health is given by $W(H) = \frac{1-\gamma}{\gamma} H^{\gamma}$, where $\gamma = w_H$ (the elasticity of utility with respect to health).

Considering this, the optimal WTP under GRACE becomes:

$$K^{D}_{\text{GRACE}} = K \cdot w_H \cdot D \cdot R = K \cdot \gamma \cdot \left(\frac{1}{1-d^*}\right)^{\gamma} \cdot \left(\frac{1}{1-l^*}\right)^{1-\gamma} \tag{3}$$

Therefore, the optimal threshold according to GRACE depends on risk aversion, disability status, and illness status. As illness and disability increase (d^* and l^*) the threshold increases.

However, I suggest there might be a limitation with this model. In fact, by definition the CRRA function holds risk aversion constant. So, the underlying assumption is that a given individual holds the same level of risk aversion throughout all possible health states he or she goes through. I suggest why this might not be an accurate representation in the next section.

3.5 Dynamic risk aversion

I define dynamic risk aversion as a risk aversion that changes depending on the health state. I hypothesize that the risk aversion level changes as one goes through different health states. The idea behind my hypothesis is that if an individual is in very poor health, I believe they would likely be more desperate for any improvement in health, and thus be willing to take more risk. In other words, their risk aversion would decrease as illness severity increases. For instance, I am considering the population that applies for clinical trials of treatments that are considered very risk. I hypothesize that this population had a higher risk aversion pre-disease or when their symptoms were milder.

This idea seems in line with the empirical estimates found by Mulligan et al. (2024), who found that individuals are risk-seeking at low levels of health, and become risk averse at a health level of $H = 0.485$ (the switching point), with the most risk averse individuals being at near perfect health. The same study found that health states below the switching point correspond to conditions like Alzheimer's disease, diabetes with end-stage renal disease, and Huntington's disease (a rare and severe genetic disease that causes nerve cells in the brain to degenerate over time), to name a few.

Figure 3 shows how risk aversion can change based on health status, from estimates gathered by Mulligan et al. (2024). This suggests my hypothesis might have some ground.

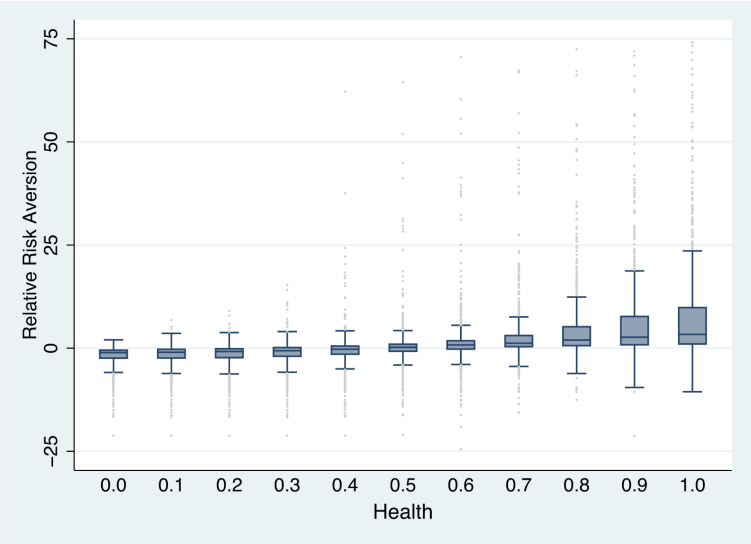

Figure 3: Individual risk aversion parameter estimates over the full range of health (Mulligan et al., 2024)

3.6 Limitations of GRACE

One limitation that is present in CEA, and still not adressed in GRACE, is the affect of dynamic pricing. In fact, both approaches do not account for the expected decrease in prices after patent expiry of drugs. These decreases can be significant, and can skew the results of analyses when calculating ICERs. For instance, Lakdawalla et al. (2024) estimate that in the US generics on average enter the market after 14 years, and prices can decrease by about 70 percent, with decreases over 90 percent not being rare.

4 Theoretical case illustration

In this section, I aim to compare CEA with GRACE using data, to show that in a real-world scenario it is possible that a drug is deemed cost-ineffective by CEA, while being cost-effective under GRACE. To do so, I consider the disease of multiple sclerosis, which is a severe neurological disease that has no cure and limited treatment options. In order to make the comparison, I run a Markov model using the "heemod" R package, where I implement two simulations — one with the CEA assumptions using linear utility, and one with the GRACE assumptions using a constant relative risk aversion utility function (CRRA). The CRRA utility function is defined as $W^{CRRA}(H) = \frac{H^{1-\rho}}{1-\rho}$ as per Mulligan et al. (2024). To measure the impact of different values of risk aversion on the WTP, I conduct a probabilistic sensitivy analysis with different values of the risk aversion parameter ρ, estimated by the aforementioned paper.

For this case study, I consider a hypothetical drug for MS called Aaltozumab, which has the same costs as the real MS drug Ocrelizumab. This is a category 3 DMT (disease-modifying-treatment), meaning it is part of the most effective options available. For some phenotypes of MS, such as primary progressive MS (PPMS), it is the only treatment available. Ocrelizumab has been approved in Australia in 2017, by the Pharmaceutical Benefits Advisory Committee (PBAC). This drug requires three injections the first year, and two each following year. One injection costs 16705.03 AUD, so annual costs come to be 50115.09 AUD for the first year, and 33410.06 AUD for the following years. However, most patients do not pay the full cost (PBS, 2025).

There is a lack of data on the impact Ocrelizumab or other DMTs have on other costs that MS patients face. This is why I am considering the hypothetical drug Aaltozumab, which I set to have the same profile as Ocrelizumab, except I make the assumption that it reduces the other costs borne by MS patients by 40 percent, thanks to its efficacy. That is, a reduction in nursing costs, hospital visits, as well as indirect costs, such as loss of productivity. I choose to focus on the Australian market to control for differences between healthcare systems and reimbursement rates between different countries, which would have introduced bias into the simulation. Although theoretical and illustrative in nature, the case study incorporates empirical data when it is available. This is discussed in the following sections, but first in the next section I will discuss the methodology used for my study: a Markov model.

4.1 Markov model

Markov models are extensively used in healthcare economics and in all sorts of cost-effectiveness analyses. Essentially a Markov process is a model that considers discrete states, such as "healthy", "sick", "dead". The simulated individual transitions through each state according to a matrix containing the transition probabilities between states. The future state depends only on the present state, and not on the past state (this is called the Markovian assumption, and it is a key principle, as it can come with its limitations). Each state is assigned a given cost and a given benefit (or utility). Then the simulation is run through software for a given number of cycles. This can be for example 30 cycles, representing 30 years. So, at each cycle, the model transitions through states according to the transition probability matrix, and records the values "cost" and "benefit", both discounted at a given rate to account for the value of time. Once all cycles are completed, heemod calculates the total cost and the total benefit. That gives the ICER for each treatment. From there, I change the model to include the GRACE assumptions (CRRA utility instead of linear utility) and run the simulation again with the same values. That gives the ICER under GRACE, which I then compare to the ICER under CEA to illustrate that the two can differ considerably.

For this case illustration, I choose to model multiple sclerosis (MS) because this disease has the particularity of having discrete disease progression categories, measured on the Expanded Disability Status Scale (EDSS) (MSTrustUK, 2025). The scale includes states from 0 (no disability) to 10 (death due to MS), increments of 0.5 (Figure 4). Thus, the discrete cases make it suitable to be modeled by a Markov process. However, including all 20 states would not be practical, because there is not enough data on the annual transition probabilities between all 20 states. To address this, J. A. Campbell et al. (2025) constructed a Markov model in which they grouped EDSS states into 4 categories: no disability (EDSS of 0.0), mild (EDSS of 1.0–3.5), moderate (EDSS of 4.0–6.0), and severe (EDSS of 6.5–9.5). Then they used a large cohort of over 6000 patients registered in the MS Database of Australia, spanning 49 years (from January 1973 to December 2021), equivalent to 38837 person-years, to estimate annual transition probabilities between the 4 categories. The predictive capacity of their model shows robustness, in that the predicted results are nearly identical to the observed results for each year through year 2 through year 10 (Figure 5). For this reason, I follow the methodology of J. A. Campbell et al. (2025) and construct my model to include 4 states: no disability, mild, moderate, and severe. The absorbing state of death is not included in the model, because J. A. Campbell et al. (2025) found that data on death rates was limited in the cohort (only $n = 102$ individuals with

all cause mortality), and including an absorbing state with such limited data caused spurious results (such as an expected life expectancy of 135 years). Therefore, the transition probabilities are conditional on no patient mortality, which is judged as a reasonable assumption by J. A. Campbell et al. (2025) because the mean age of the cohort (42.51 years) is below the average life expectancy of MS patients.

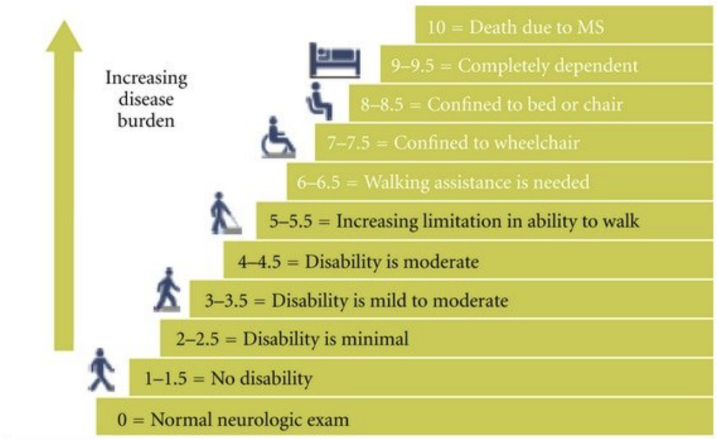

Figure 4: EDSS illustration by Singh et al. (2012)

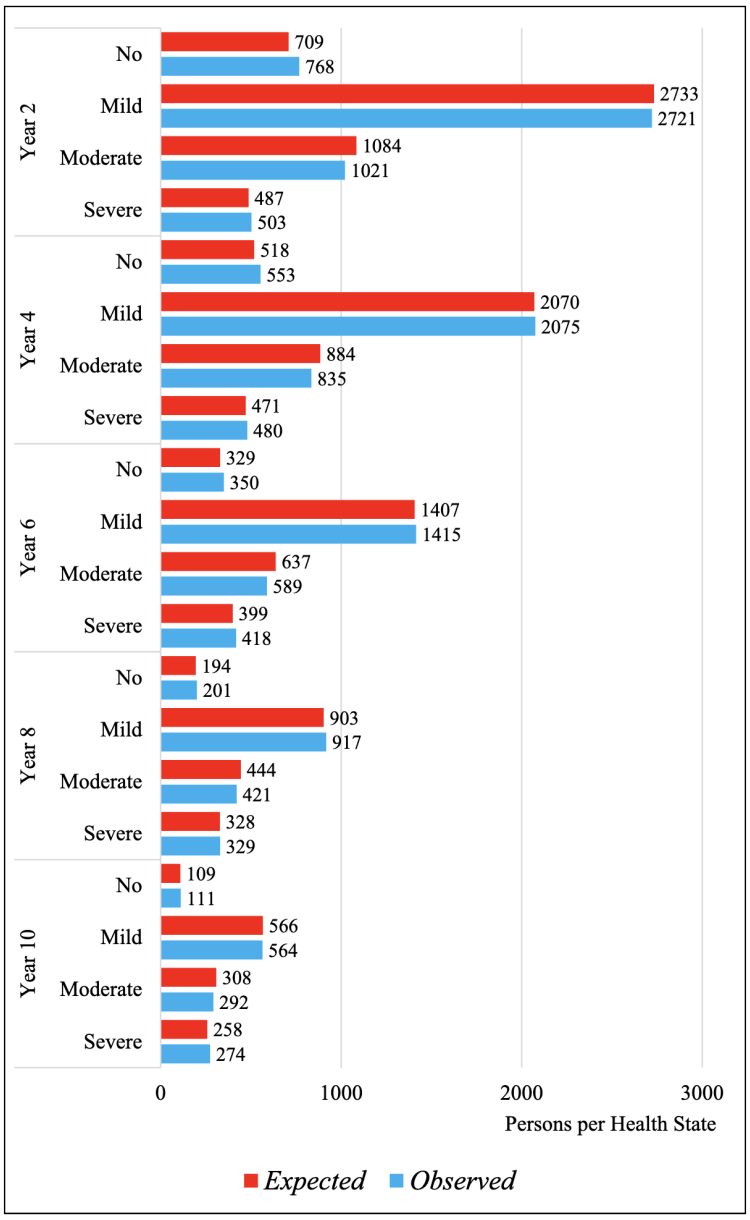

Figure 5: Expected and observed states based on the transition probabilities estimated by J. A. Campbell et al. (2025)

For my model, I consider two matrices of transition probabilities, one for Aaltozumab (which is a category 3 DMT) and one for no DMT usage. In fact, Aaltozumab (or other treatments) influence transition probabilities, since they are medically effective drugs that slow down the progression of the disease.

Using the matrices provided by J. A. Campbell et al. (2025), I construct the Markov model depicted in Figure 6 to represent the scenario when no DMT is used, and Figure 7 the scenario when the patient is using Aaltozumab. These two scenarios

are called strategies. So, I compare the strategy Aaltozumab (which uses the transition probabilities of a DMT of category 3, and costs the same as Ocrelizumab while reducing other costs by 40 percent) against the strategy of using no DMT (which uses the transition probabilities of no DMT usage by J. A. Campbell et al. (2025), and for which I assume costs are the average costs borne by MS patients in Australia in 2021 depending on their EDSS-category, estimated by J. Campbell et al. (2023)).

For utilities, I use the values estimated by J. A. Campbell et al. (2023). As the authors note, there are multiple ways to measure utilities, the most popular being the EQ-5D-5L framework, which is a questionnaire focused on the physical aspects of health utility. However, the paper suggests that such framework is not as accurate as the EQ-5D-5L-Psychosocial framework, which includes more questions aimed at capturing psychosocial aspects (such as sleep, loneliness, and relationships) into the health utility measure. The authors suggest that EQ-5D-5L-Psychosocial is more accurate for non-severe patients. Therefore, in my model I use the EQ-5D-5L utility for the severe state, and the EQ-5D-5L-Psychosocial utilities for all other states.

As such, all of the data that is incorporated in my model is from the same country — Australia. This should address a limitation that many models have: mixing international data, and thus having the effect of different healthcare prescription and reimbursement policies embedded in the results, which makes them less reliable. I believe that considering a homogeneous scenario yields more accurate results. Moreover, Australia has the largest database of MS patients (J. A. Campbell et al., 2025), which helps support the validity of the model.

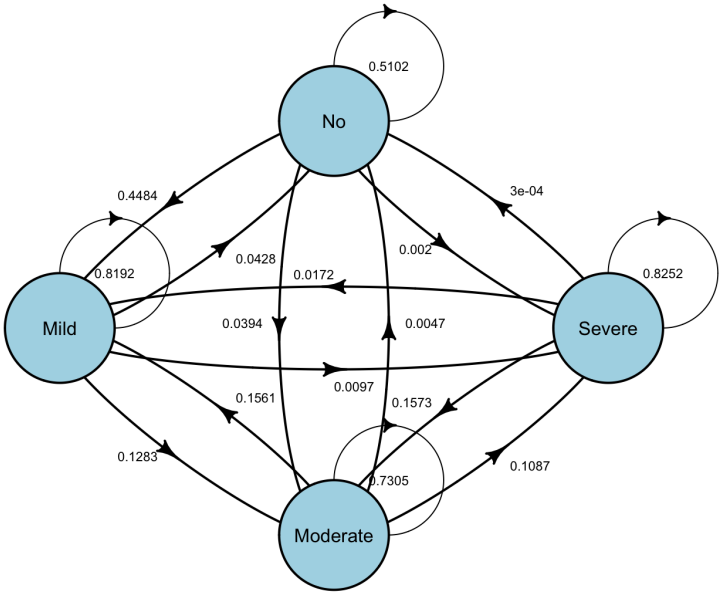

Figure 6: States of the Markov model, and transition probabilities when using no DMT. Code in the appendices.

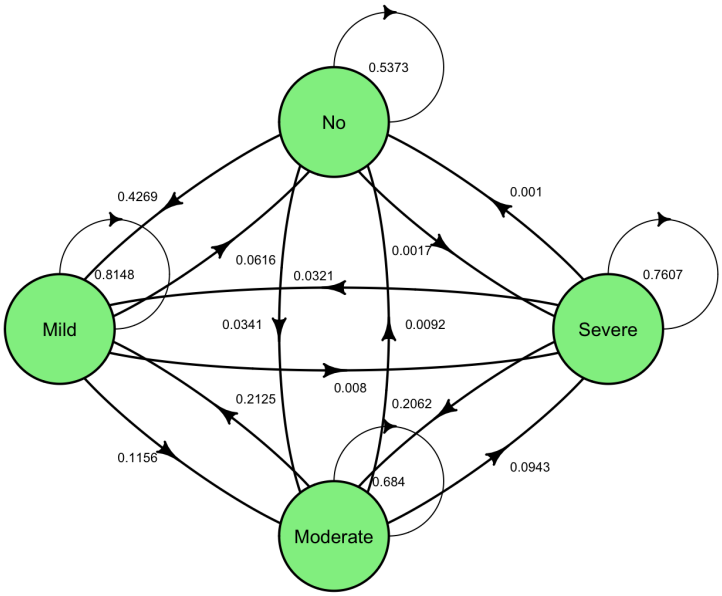

Figure 7: States of the Markov model, and transition probabilities when using Aaltozumab. Code in the appendices

4.2 Case under CEA

I ran the model for 30 cycles, which produced an ICER of 89498.85 AUD/QALY (Figure 8). The progression through various states over time is depicted in Figure 9. This would make Aaltozumab cost ineffective in Australia. In fact, while the PBAC does not openly disclose the threshold used, it is estimated to be between 45000-60000 AUD/QALY (Sampson et al., 2022).

```
Values:

                cost  utility
noDMT      1142481246 9347.164
Aaltozumab 1192141663 9902.036

Efficiency frontier:

noDMT -> Aaltozumab

Differences:

           Cost Diff. Effect Diff.     ICER  Ref.
Aaltozumab   49660.42    0.5548721 89498.85 noDMT
```

Figure 8: Results of CEA analysis: Aaltozumab vs no DMT

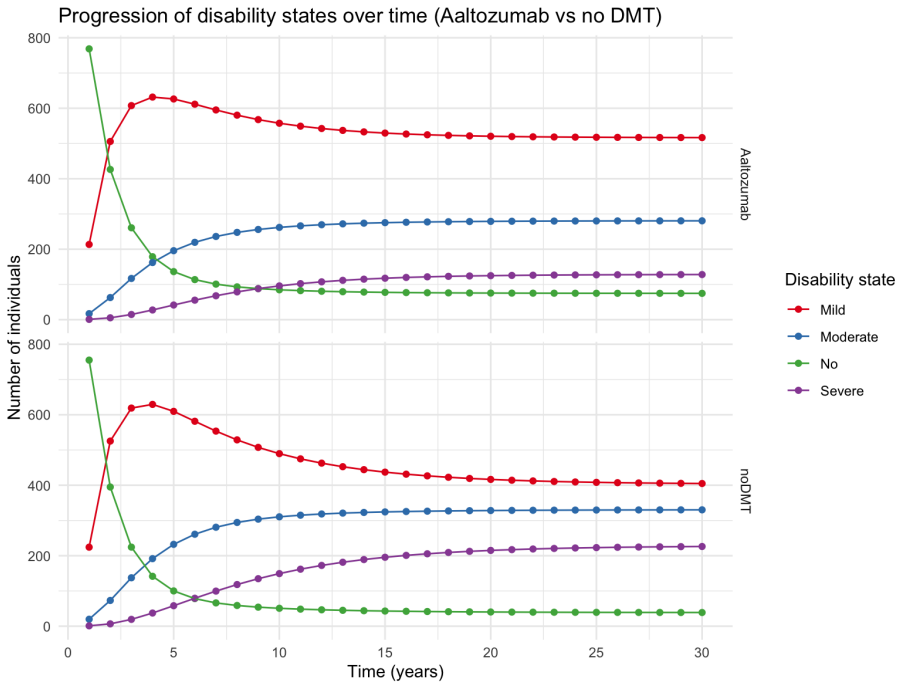

Figure 9: Progression of disability states over time (Aaltozumab vs no DMT) in the simulation

4.3 Case under GRACE

I implemented the same simulation again, except that I applied the CRRA utility function to the utility values. For the risk aversion parameter ρ, I used the mean value of 0.2822, which is the estimate from the pooled sample of Mulligan et al. (2024). With this value of ρ, the ICER is 71187.55 AUD/QALY. Considerably lower than CEA (Figure 10). However, it is still above the PBAC upper bound threshold of 60000 AUD/QALY.

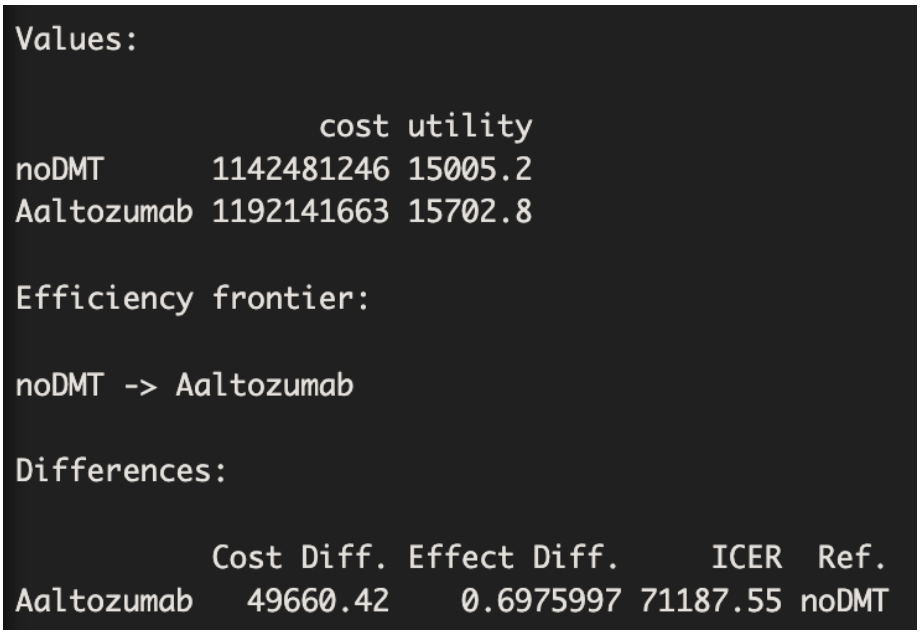

Figure 10: Results of GRACE analysis: Aaltozumab vs no DMT

4.4 Deterministic sensitivity analysis (DSA)

To delve deeper into the changes caused by the risk-aversion value, I conducted a deterministic sensitivity analysis to test the lower and upper bounds of the risk parameter estimated by Mulligan et al. (2024) and how they would influence the result. I find that, by using the upper-bound estimate of 0.349, Aaltozumab would have an ICER of 67311.37 (Figure 11), which is close to the threshold of cost-effectiveness of the PBAC.

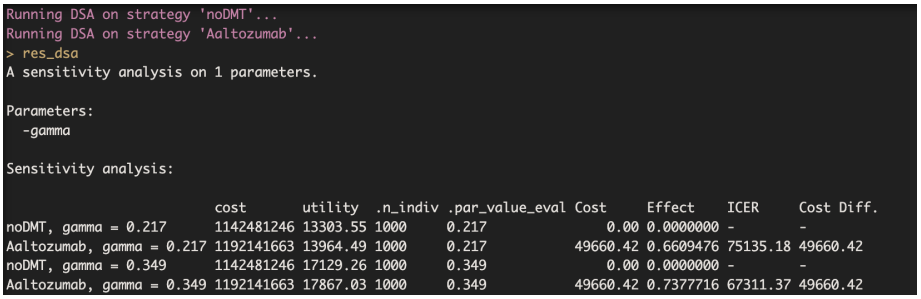

Figure 11: Deterministic sensitivity analysis on risk-aversion parameter

This shows that the higher the risk aversion, the lower the ICER. In other words, patients with a higher risk aversion (and thus healthier) should face a lower threshold (i.e., a less generous threshold), and those with lower risk aversion (sicker patients)

should face a higher (i.e., more generous) threshold, if the goal is to match their WTP with the appropriate ICER.

4.5 Probabilistic sensitivity analysis (PSA)

In this section, I conduct a probabilistic sensitivity analysis through a Monte Carlo simulation with 1000 simulations, assuming a normal distribution of the risk-aversion parameter, and using the mean value of 0.2822 discussed earlier.

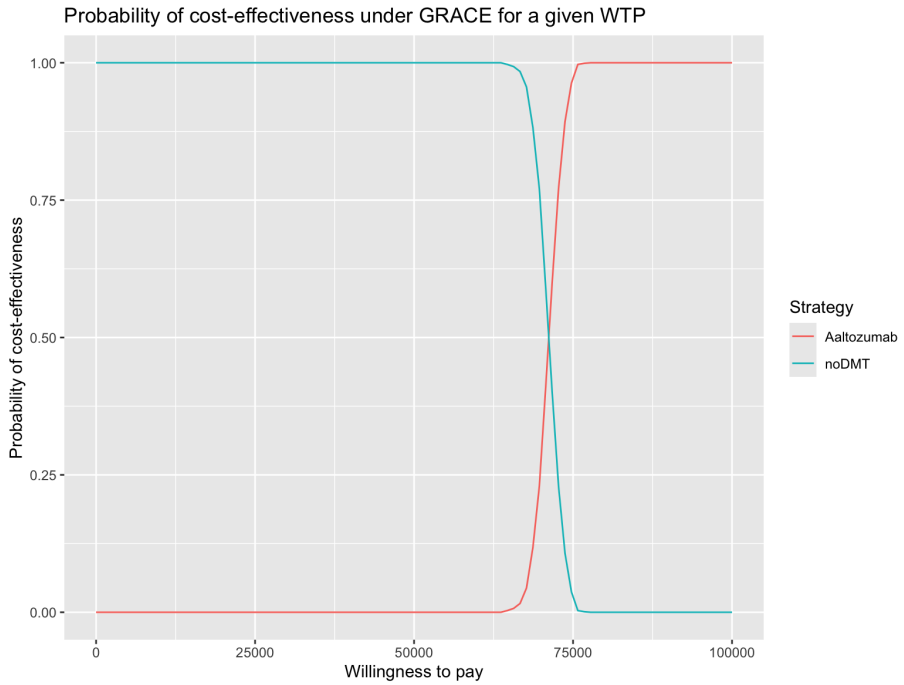

Figure 12: PSA on risk-aversion parameter

The results suggest that the distribution and standard deviation of the risk-aversion parameter greatly affect the shape of the cost acceptability curve (CEAC) depicted above. Therefore, it is essential for more research to be conducted in order to properly estimate the shape of distribution and the mean values for different subgroups of individuals, especially those suffering from rare and severe diseases. This would allow to produce more accurate results.

This case illustration showed that it is entirely possible for a treatment to fall below a given threshold under GRACE, and not under CEA. For a rare and severe disease such as MS, it is important to take this into account, as phenotypes such as PPMS only have one treatment option as of 2025: Ocrelizumab.

5 Discussion: policy implications

Competent authorities have already recognized the shortcomings and limitations of the CEA approach, by means such as the Affordable Care Act banning the CEA measures in federal programs. However, the GRACE model, which addresses these limitations, is still not widely adopted in place of the CEA.

I suggest that gradually phasing out CEA models and making GRACE the new standard can bring benefits to society through multiple pathways. First, assessing the cost-effectiveness of healthcare treatments more accurately can reduce the deadweight-loss caused by the mismatch between companies that could produce a new drug that would benefit patients, particularly those who are severely ill and for which treatment options are limited. An example would be the disease RRMS and the medication dimethyl fumarate, which was deemed cost ineffective. I hypothesize this assessment might have contributed to the significant price increase of the drug, which went from about 4500 USD/month in March 2013 to about 8200 USD/month in January 2020. Fortunately, generic versions entered the market and offered an alternative at 1 percent of the cost. However, this is not always the case. Another type of multiple sclerosis, primary progressive multiple sclerosis (PPMS), is a rarer version compared to RRMS, and the available treatments are much more limited. The FDA (US Food and Drug Administration) only approves one PPMS treatment: ocrelizumab (brand name Ocrevus). It is an infusion that is given every 6 months and costs 83000 USD/year in the US. It is deemed cost ineffective by traditional CEA, but it is the only FDA-approved so patients do not have much of a choice.

The difference between PPMS and RRMS is that PPMS only progresses, it never gets better. So, patients become progressively more ill over time, as opposed to RRMS where it is possible to reverse to prior stages of the disease. This is a relevant distinction to highlight, because if my hypothesis that risk-aversion declines as the disease severity increases holds, then that would imply that the threshold for ocrelizumab to be considered cost-effective should increase not only thanks to the GRACE assumptions, but also through the expected decrease in risk aversion. If policies address this issue, I argue that it might be possible to decrease the price of ocrelizumab and thus make it available to more patients, who desperately need it since there are no alternatives. The reason why I believe the price would decrease is because, as explained by Phelps (2019), medical interventions that are cost-effective "are expanded in use (either offered to more people or are used more frequently, or both)", and those that are not "have their use restricted in various ways or the price

falls, or both".

Another pathway through which wide adoption of GRACE can bring value to society is by better understanding the risk preferences of various individuals, based on their illness status or disability status. This is valuable because a care provider might be more prone to disclose the uncertainty around a prognostic, if said provider has a tool to assess whether the patient would benefit from knowing about the uncertainty level of the medical treatment. This tool can be the GRACE framework and its accompanying research. In fact, although GRACE can be considered to be in its early stages and presents some limitations, it provides a direction for researchers to explore, namely the field of risk preferences in the valuation of medical interventions, which can directly impact the wellbeing of a growing proportion of the population (as natality rates fall and the relative population of the elderly increases).

6 Conclusion

The traditional CEA approach presents a wide range of shortcomings, notably by not accounting for the diminishing returns to health, for the expected decrease in drug price after patent expiry, and for risk preferences, it skews the willingness-to-pay estimates, and thus the thresholds put in place by the authorities. This has an affect on society's welfare, as patients who would have benefitted from a medical intervention might not have access to it because of the systematic flaws embedded in the CEA model.

The GRACE framework provides an alternative approach that I believe is a more accurate representation of the value that patients place on their health and on the expected benefits of their prognostic. The literature reviewed in this thesis points in this direction, as well as the simulation case study, which highlighted the effects of wrongly assessing the cost effectiveness of a treatment. Similarly, the fact that many exceptions and restrictions are set in place for the use of CEA further suggests that the model is not the best alternative available.

To develop more accurate and realistic models, it is essential to conduct more research to estimate the parameter of risk aversion and the dynamics through which it changes, specifically when disease progresses. This would shed more light into ways on how the GRACE framework can be improved, with the goal of closing the gap between patients and their true valuations and eventually minimizing the deadweight loss experienced by society as a whole, all while decreasing discrimination towards the severely ill and the disabled population.

References

Campbell, J., van der Mei, I., Taylor, B., & Palmer, A. (2023). Health Economic Impact of Multiple Sclerosis in Australia in 2021: An Interim Update of Prevalence, Costs and Cost of Illness from 2017 to 2021 [Publisher: University of Tasmania]. *Multiple Sclerosis Australia*. Retrieved April 27, 2025, from https://ms-research.neonjungle.com.au/documents/2/health-economic-impact-of-multiple-sclerosis-in-australia-in-2021_final.pdf

Campbell, J. A., Ahmad, H., Chen, G., Van Der Mei, I., Taylor, B. V., Claflin, S., Henson, G. J., Simpson-Yap, S., Laslett, L. L., Hawkes, K., Hurst, C., Waugh, H., & Palmer, A. J. (2023). Validation of the EQ-5D-5L and psychosocial bolt-ons in a large cohort of people living with multiple sclerosis in Australia. *Quality of Life Research*, *32*(2), 553–568. https://doi.org/10.1007/s11136-022-03214-y

Campbell, J. A., Henson, G. J., Ngwa, V. F., Ahmad, H., Taylor, B. V., Van Der Mei, I., MSBase Australian Researchers, Butzkueven, H., Kermode, A. G., Fabis-Pedrini, M., Carroll, W. M., Kalincik, T., Roos, I., Lechner-Scott, J., Walt, A. V. D., Hodgkinson, S., John, N., McCombe, P., Macdonell, R., ... Palmer, A. J. (2025). Estimation of Transition Probabilities from a Large Cohort (> 6000) of Australians Living with Multiple Sclerosis (MS) for Changing Disability Severity Classifications, MS Phenotype, and Disease-Modifying Therapy Classifications. *PharmacoEconomics*, *43*(2), 223–239. https://doi.org/10.1007/s40273-024-01417-4

Filipović-Pierucci, A., Zarca, K., & Durand-Zaleski, I. (2017, April). Markov Models for Health Economic Evaluations: The R Package heemod [arXiv:1702.03252 [stat]]. https://doi.org/10.48550/arXiv.1702.03252

Garber, A. M., & Phelps, C. E. (1997). Economic foundations of cost-effectiveness analysis. *Journal of Health Economics*, *16*(1), 1–31. https://doi.org/10.1016/S0167-6296(96)00506-1

Lakdawalla, D. N., Phelps, C. E., Arndorfer, S., Incerti, D., & Masia, N. (2024). Getting the Math Right When Measuring the Value of New Medicines. Retrieved March 3, 2025, from https://nopatientleftbehind.docsend.com/view/889u6zs74tra9x4a

Lakdawalla, D. N., & Phelps, C. E. (2021). Health Technology Assessment With Diminishing Returns to Health: The Generalized Risk-Adjusted Cost-

Effectiveness (GRACE) Approach. *Value in Health*, *24*(2), 244–249. https://doi.org/10.1016/j.jval.2020.10.003

Lakdawalla, D. N., & Phelps, C. E. (2023). The Generalized Risk-Adjusted Cost-Effectiveness (GRACE) Model for Measuring the Value of Gains in Health: An Exact Formulation. *Journal of Benefit-Cost Analysis*, *14*(1), 44–67. https://doi.org/10.1017/bca.2023.6

MSTrustUK. (2025). Expanded Disability Status Scale (EDSS) for Multiple Sclerosis. Retrieved April 28, 2025, from https://mstrust.org.uk/a-z/expanded-disability-status-scale-edss

Mulligan, K., Baid, D., Doctor, J. N., Phelps, C. E., & Lakdawalla, D. N. (2024). Risk preferences over health: Empirical estimates and implications for medical decision-making. *Journal of Health Economics*, *94*, 102857. https://doi.org/10.1016/j.jhealeco.2024.102857

PBS. (2025). Price of Ocrelizumab. Pharmaceutical Benefits Scheme (PBS). Retrieved April 27, 2025, from https://www.pbs.gov.au/medicine/item/11237k-11242q

Phelps, C. E. (2019). A New Method to Determine the Optimal Willingness to Pay in Cost-Effectiveness Analysis. *Value in Health*, *22*(7), 785–791. https://doi.org/10.1016/j.jval.2019.03.003

Sampson, C., Zamora, B., Watson, S., Cairns, J., Chalkidou, K., Cubi-Molla, P., Devlin, N., García-Lorenzo, B., Hughes, D. A., Leech, A. A., & Towse, A. (2022). Supply-Side Cost-Effectiveness Thresholds: Questions for Evidence-Based Policy. *Applied Health Economics and Health Policy*, *20*(5), 651–667. https://doi.org/10.1007/s40258-022-00730-3

Singh, A. V., Khare, M., Gade, W. N., & Zamboni, P. (2012). Theranostic Implications of Nanotechnology in Multiple Sclerosis: A Future Perspective. *Autoimmune Diseases*, *2012*, 1–12. https://doi.org/10.1155/2012/160830

Williams, A. (1992). Cost-effectiveness analysis: Is it ethical? *Journal of medical ethics*, *18*(1), 7–11. https://doi.org/10.1136/jme.18.1.7

A Appendices

A.1 Code for CEA simulation

```r
#install if not already installed
install.packages("heemod")
install.packages("dplyr")
install.packages("ggplot2")
install.packages("diagram")
library(ggplot2)
library(diagram)
library(dplyr)
library(heemod)

#here I define the parameters for the simulation
param <- define_parameters(
  u_no = 0.79, #utility of no disability state under CEA
  u_mild = 0.66, #utility of mild disability state under CEA
  u_moderate = 0.53, #utility of moderate disability state under CEA
  u_severe = 0.18, #utility of severe disability state under CEA
  cost_no = 32829, #average annual cost for MS patients with no disability in 2021 in
    Australia
  cost_mild = 59957,
  cost_moderate = 82624,
  cost_severe = 123333
)
#Utilities are EQ-5D-5L-Psychosocial for no, mild, moderate disability cases,
#and EQ-5D-5L for severe cases, as per suggestions by Campbell, 2023.
#The utilities are mean values

#matrix containing the transition probabilities between the 4 states when using no
    treatment
mat_trans_noDMT <- define_transition(
  state_names = c(
    "No",
    "Mild",
    "Moderate",
    "Severe"
  ),
  0.5102, 0.4484, 0.0394, 0.002,
  0.0428, 0.8192, 0.1283, 0.0097,
  0.0047, 0.1561, 0.7305, 0.1087,
  0.0003, 0.0172, 0.1573, 0.8252
```

```
)
mat_trans_noDMT
plot(mat_trans_noDMT, cex.txt=0.7, box.col="lightblue")
#plots diagram with states and transition probabilities. Text size 0.7

#matrix containing the transition probabilities between the 4 states when using a
    category-3 DMT
#ocrelizumab is a category 3 DMT (as is the theoretical aaltozumab)
mat_trans_DMTcat3 <- define_transition(
  state_names = c(
    "No",
    "Mild",
    "Moderate",
    "Severe"
  ),
  0.5373, 0.4269, 0.0341, 0.0017,
  0.0616, 0.8148, 0.1156, 0.008,
  0.0092, 0.2125, 0.6840, 0.0943,
  0.0010, 0.0321, 0.2062, 0.7607
)
mat_trans_DMTcat3
plot(mat_trans_DMTcat3, cex.txt=0.7, box.col="lightgreen")

#next, I attach values to each state
#model_time is a built-in variable in heemod that keeps track of the cycle count during
    the simulation
#The first year of treatment requires 3 injections of ocrelizumab, hence costs 16705.03*
    3=50115.09 AUD
#Subsequent years require only 2 annual injections, hence 16705.03*2=33410.06 AUD
#Costs and utilities are discounted using the heemod discount function at an annual rate
    of 5%
#in line with current Australian discounting practices for healthcare
state_no <- define_state(
  cost = dispatch_strategy(
    noDMT = discount(cost_no, 0.05),
    Aaltozumab = discount(ifelse(model_time == 1, 0.6*cost_no+50115.09, 0.6*cost_no
    +33410.06), 0.05)
  ),
  utility = discount(u_no,
                     0.05) #5% discount rate
)
state_no

```

```r
state_mild <- define_state(
  cost = dispatch_strategy(
    noDMT = discount(cost_mild, 0.05),
    Aaltozumab = discount(ifelse(model_time == 1, 0.6*cost_mild+50115.09, 0.6*cost_mild
    +33410.06), 0.05)
  ),
  utility = discount(u_mild,
                     0.05)
)
state_mild

state_moderate <- define_state(
  cost = dispatch_strategy(
    noDMT = discount(cost_moderate, 0.05),
    Aaltozumab = discount(ifelse(model_time == 1, 0.6*cost_moderate+50115.09, 0.6*cost_
    moderate+33410.06), 0.05)
  ),
  utility = discount(u_moderate,
                     0.05)
)
state_moderate

state_severe <- define_state(
  cost = dispatch_strategy(
    noDMT = discount(cost_severe, 0.05),
    Aaltozumab = discount(ifelse(model_time == 1, 0.6*cost_severe+50115.09, 0.6*cost_
    severe+33410.06), 0.05)
  ),
  utility = discount(u_severe,
                     0.05)
)
state_severe

#next, I combine the transition matrix with the state values by using the define_strategy
    function
#strategy using no DMT
strat_noDMT <- define_strategy(
  transition = mat_trans_noDMT,
  No = state_no,
  Mild = state_mild,
  Moderate = state_moderate,
  Severe = state_severe
)
```

```r
strat_noDMT

#strategy using a DMT of category 3 (like ocrelizumab and aaltozumab)
strat_DMTcat3 <- define_strategy(
  transition = mat_trans_DMTcat3,
  No = state_no,
  Mild = state_mild,
  Moderate = state_moderate,
  Severe = state_severe
)
strat_DMTcat3

#next, I run the model for 30 cycles (ie, 30 years)
#by default, the model is run for 1000 persons starting in the first state (here the no
    disability state)
res_mod <- run_model(
  noDMT = strat_noDMT,
  Aaltozumab = strat_DMTcat3,
  parameters = param,
  cycles = 30,
  cost = cost,
  effect = utility
)
res_mod

plot(res_mod) +
  xlab("Time (years)") +
  ylab("Number of individuals") +
  ggtitle("Progression of disability states over time (Aaltozumab vs no DMT) - CEA") +
  theme_minimal() +
  scale_color_brewer(
    name = "Disability state",
    palette = "Set1"
  )
```

Listing 1: Code for CEA simulation

A.2 Code for GRACE simulation

```r
#install if not already installed
install.packages("heemod")
install.packages("dplyr")
install.packages("ggplot2")
install.packages("diagram")
library(ggplot2)
library(diagram)
library(dplyr)
library(heemod)

#here I define the parameters for the simulation
param <- define_parameters(
  gamma = 0.2822, #estimate from mulligan2024 pooled sample
  u_no = 0.79, #utility of no disability state under CEA
  u_mild = 0.66, #utility of mild disability state under CEA
  u_moderate = 0.53, #utility of moderate disability state under CEA
  u_severe = 0.18, #utility of severe disability state under CEA
  cost_no = 32829, #average annual cost for MS patients with no disability in 2021 in
    Australia
  cost_mild = 59957,
  cost_moderate = 82624,
  cost_severe = 123333
  )
#Utilities are EQ-5D-5L-Psychosocial for no, mild, moderate disability cases,
#and EQ-5D-5L for severe cases, as per suggestions by Campbell, 2023.
#The utilities are mean values

#matrix containing the transition probabilities between the 4 states when using no
    treatment
mat_trans_noDMT <- define_transition(
  state_names = c(
    "No",
    "Mild",
    "Moderate",
    "Severe"
  ),
  0.5102, 0.4484, 0.0394, 0.002,
  0.0428, 0.8192, 0.1283, 0.0097,
  0.0047, 0.1561, 0.7305, 0.1087,
  0.0003, 0.0172, 0.1573, 0.8252
)
mat_trans_noDMT
```

```r
plot(mat_trans_noDMT, cex.txt=0.7, box.col="lightblue")
#plots diagram with states and transition probabilities. Text size 0.7

#matrix containing the transition probabilities between the 4 states when using a category-3 DMT
#ocrelizumab is a category 3 DMT (as is the theoretical aaltozumab)
mat_trans_DMTcat3 <- define_transition(
  state_names = c(
    "No",
    "Mild",
    "Moderate",
    "Severe"
  ),
  0.5373, 0.4269, 0.0341, 0.0017,
  0.0616, 0.8148, 0.1156, 0.008,
  0.0092, 0.2125, 0.6840, 0.0943,
  0.0010, 0.0321, 0.2062, 0.7607
)
mat_trans_DMTcat3
plot(mat_trans_DMTcat3, cex.txt=0.7, box.col="lightgreen")

#next, I attach values to each state
#model_time is a built-in variable in heemod that keeps track of the cycle count during the simulation
#The first year of treatment requires 3 injections of the drug, hence costs 16705.03*3=50115.09 AUD
#Subsequent years require only 2 annual injections, hence 16705.03*2=33410.06 AUD
#Costs and utilities are discounted using the heemod discount function at an annual rate of 5%
#in line with current Australian discounting practices for healthcare
state_no <- define_state(
  cost = dispatch_strategy(
    noDMT = discount(cost_no, 0.05),
    Aaltozumab = discount(ifelse(model_time == 1, 0.6*cost_no+50115.09, 0.6*cost_no+33410.06), 0.05)
  ),
  utility = discount(ifelse(gamma == 1, log(u_no), (u_no^(1 - gamma))/(1 - gamma)), #applying CRRA utility for GRACE
                     0.05) #5% discount rate
)
state_no

state_mild <- define_state(
```

```
  cost = dispatch_strategy(
    noDMT = discount(cost_mild, 0.05),
    Aaltozumab = discount(ifelse(model_time == 1, 0.6*cost_mild+50115.09, 0.6*cost_mild
    +33410.06), 0.05)
  ),
  utility = discount(ifelse(gamma == 1, log(u_mild), (u_mild^(1 - gamma))/(1 - gamma)),
                     0.05)
)
state_mild

state_moderate <- define_state(
  cost = dispatch_strategy(
    noDMT = discount(cost_moderate, 0.05),
    Aaltozumab = discount(ifelse(model_time == 1, 0.6*cost_moderate+50115.09, 0.6*cost_
    moderate+33410.06), 0.05)
  ),
  utility = discount(ifelse(gamma == 1, log(u_moderate), (u_moderate^(1 - gamma))/(1 -
    gamma)),
                     0.05)
)
state_moderate

state_severe <- define_state(
  cost = dispatch_strategy(
    noDMT = discount(cost_severe, 0.05),
    Aaltozumab = discount(ifelse(model_time == 1, 0.6*cost_severe+50115.09, 0.6*cost_
    severe+33410.06), 0.05)
  ),
  utility = discount(ifelse(gamma == 1, log(u_severe), (u_severe^(1 - gamma))/(1 - gamma)
    ),
                     0.05)
)
state_severe

#next, I combine the transition matrix with the state values by using the define_strategy
    function
#strategy using no DMT
strat_noDMT <- define_strategy(
  transition = mat_trans_noDMT,
  No = state_no,
  Mild = state_mild,
  Moderate = state_moderate,
  Severe = state_severe
```

```r
)
strat_noDMT

#strategy using a DMT of category 3 (like ocrelizumab and aaltozumab)
strat_DMTcat3 <- define_strategy(
  transition = mat_trans_DMTcat3,
  No = state_no,
  Mild = state_mild,
  Moderate = state_moderate,
  Severe = state_severe
)
strat_DMTcat3

#next, I run the model for 30 cycles (ie, 30 years)
#by default, the model is run for 1000 persons starting in the first state (here the no
    disability state)
res_mod <- run_model(
  noDMT = strat_noDMT,
  Aaltozumab = strat_DMTcat3,
  parameters = param,
  cycles = 30,
  cost = cost,
  effect = utility
)
res_mod

plot(res_mod) +
  xlab("Time (years)") +
  ylab("Number of individuals") +
  ggtitle("Progression of disability states over time (Aaltozumab vs no DMT) - GRACE") +
  theme_minimal() +
  scale_color_brewer(
    name = "Disability state",
    palette = "Set1"
  )

#DSA for gamma (relative risk aversion parameter)
#lower and upper bounds suggested by table 4 of mulligan2024 (pooled sample)
se <- define_dsa(
  gamma, 0.217, 0.349
)

res_dsa <- run_dsa(
```

```
  model = res_mod,
  dsa = se
)
res_dsa

plot(res_dsa,
     result = "icer",
     type = "difference")

#PSA setting gamma as the mean rho CRRA parameter from table 4 of mulligan2024
rsp <- define_psa(
  gamma ~ lognormal(mean = 0.2822, sd = 0.033) #sd roughly approximated as (0.349-0.217)/
    4
)

pm <- run_psa(
  model = res_mod,
  psa = rsp,
  N = 1000
)
summary(pm)
plot(pm, type = "ac", max_wtp = 100000, log_scale = FALSE) +
  ggtitle("Probability of cost-effectiveness under GRACE for a given WTP")
```

Listing 2: Code for GRACE simulation