NICE National Institute for Health and Care Excellence

Developing a novel model structure for evaluating treatment sequences in oncology pathways: a pilot in non-small cell lung cancer

Lucy Beggs[1], Kusal Lokuge[1], Yuanyuan Zhang[1], Lindsay Claxton[1]

[1] National Institute for Health and Care Excellence (NICE)

Background

NICE technology appraisals (TAs) evaluate cost-effectiveness of interventions at individual decision points in a treatment pathway. However, the treatment sequence generated from a series of TAs may differ from the optimal sequence estimated for the whole pathway. Generating individual models for separate TAs in the same indication may also be less resource efficient and consistent than using a single multi-use model. In 2023, NICE commissioned a pilot "pathways approach" for evaluating sequences in advanced non-small cell lung cancer (NSCLC). Existing model structures to evaluate pathways require patient-level simulation[1] or use numerous tunnel states, with potential operational barriers around data access, ease of use and interpretation. We aimed to develop a novel model structure in Microsoft Excel for evaluation of oncology treatment sequences to overcome these challenges.

Model requirements

Standard cohort model structures used in oncology are typically based on progression-free (PFS) and overall survival (OS) data. Clinical trial estimates of OS are affected by any subsequent treatments received after an intervention. This is not usually problematic as the distributions of subsequent treatments are assumed to be broadly balanced across randomised arms and are unlikely to affect the relative treatment effect between comparators.

However, in this pilot we wanted to explicitly model each treatment in the pathway with enough detail to inform a TA. This meant we needed to isolate line-specific treatment effects; using OS data was not appropriate for any interventions that had subsequent treatments because it would be unclear where the survival benefit was being derived from.

Patients in the model progress from first-line to subsequent treatments at different points in time. To estimate time spent in each health state, the model needed to account for this.

The context of the pilot meant that model needed to be accessible to stakeholders who were familiar with Microsoft Excel and needed a cohort structure.

The model scope covered 7 separate decision points ("nodes") across 3 lines of treatment, varying by histology and PD-L1 status.

Table 1: Model structure

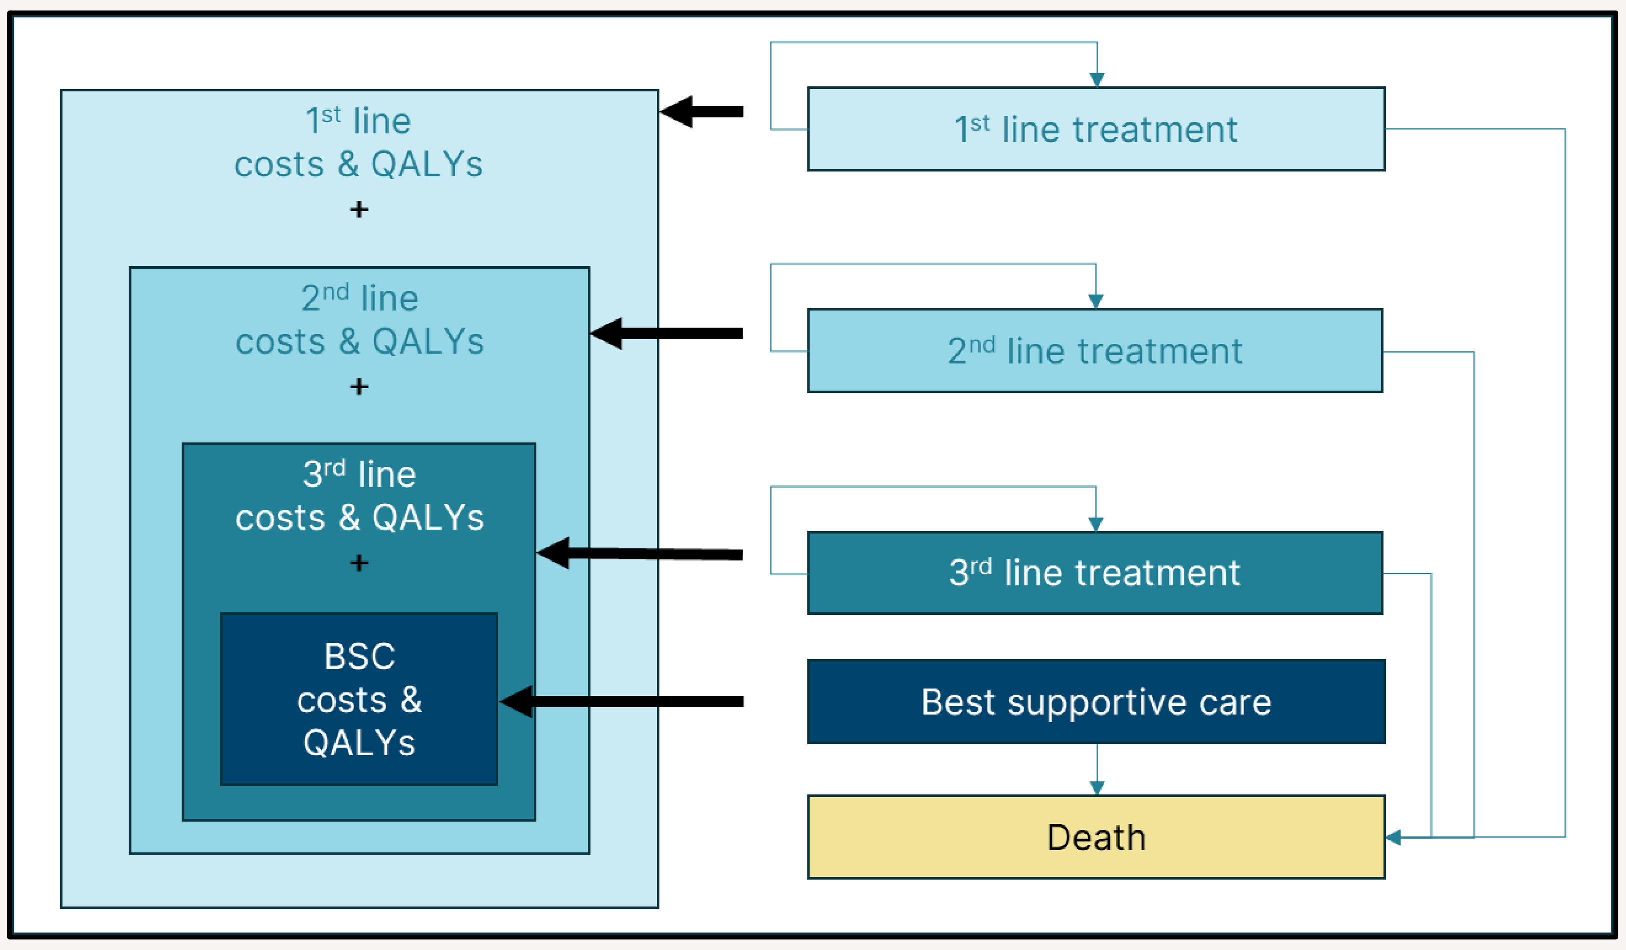

Model structure

We developed a novel '**nested partitioned survival**' model structure, where the last-line is built first and earlier lines are added in reverse order.

OS data was suitable for use in modelling of last-line best supportive care (BSC) as this had no subsequent treatments. For this node, a simplified two-state survival model was developed to capture discounted costs and QALYs for patients on BSC.

Third-line was modelled using PFS and pre-progression death (PPD) analyses to generate 3 health states: progression-free, progressed and dead. Patients generated costs and QALYs for any time in the progression-free health state. PPD was used to estimate transitions to either death or the progressed state. Patients moved onto BSC at progression, so instead of generating costs and QALYs for the progressed state, the aggregate values from the BSC component of the model were applied.

The same approach was taken for second- and first-line modelling. Patients start in a progression-free state, and move to progressed or death based on PPD data. Costs and QALYs are explicitly generated in the progression-free state, and an aggregate of the costs and QALYs generated from later line models is applied to the progressed state to capture all downstream consequences. A weighted average was applied to this aggregate estimate to reflect subsequent treatment distributions.

What we learnt

The key output of this work was the conceptualization of a novel model structure for modelling sequences in oncology and development of a fully-executable model demonstrating the concept. This structure can be implemented in Excel without access to patient-level data, and can answer questions about optimal treatment choice at a specific node, optimal positioning of a treatment in a pathway and optimal treatment sequences.

Results from the model were based on list prices without commercial arrangements applied and so do not reflect true cost-effectiveness. For non-squamous patients of any PD-L1 status, the most cost-effective sequence was pemetrexed + platinum chemotherapy → docetaxel → nintedanib + docetaxel. For squamous patients of any PD-L1 status the most cost-effective sequence was platinum chemotherapy → docetaxel → no treatment.

The model was complex and computationally intensive, with 119 different sequences across multiple subgroups and additional placeholders for future model extensions.

PFS data came from an evidence synthesis of aggregate trial results, but PPD was not widely reported. As this is a critical input for the model structure, SACT data was used to estimate the proportion of people who died before progression for each treatment, allowing PPD estimates to be derived from PFS. The project illustrates how real-world evidence can be used alongside trial data to model treatment pathways.

Using the model at specific decision nodes

The model allows users to select the specific decision node of interest and relevant comparators for economic evaluation. The model has separate engines for each line of treatment, and based on the selection, the relevant engines of the model are run to calculate costs and QALYs for the selected node and a weighted average of downstream nodes.

Conceptually, costs and QALYs are generated in reverse order (i.e. BSC results forming the progressed-state of the third-line engine, the aggregate of BSC and third-line forming the progressed-state of second-line etc). However, the proportion of patients at each line of treatment varies depending on the user-defined decision node of interest.

If a user only wanted to generate results for BSC, the BSC engine would start with a full cohort. However, if a user wanted to generate results for a third-line node it would be this node that starts with a full cohort and the membership of the BSC state would be determined by transitions based on pre-progression death data.

References

1. Tappenden P, Chilcott J, Brennan A, Squires H, Stevenson M. Whole Disease Modeling to Inform Resource Allocation Decisions in Cancer: A Methodological Framework. Value in Health. 2012;15(8):1127-1136.

Lucy Beggs
lucy.beggs@nice.org.uk

National Institute for Health and Care Excellence
Health Economics and Decision Modelling team

nice.org.uk