Late stage combination drug development for improved portfolio-level decision-making

Emily Graham[1], Thomas Jaki[1], Chris Harbron[2]

[1] Lancaster University, [2] Roche Products Ltd

June 5, 2018

Outline

- We are interested in the problem of decision-making for a portfolio of combination therapies.
- In this talk I will discuss:
  1. Motivation and background
  2. Portfolio-level decision-making
  3. Extension to combination therapies
  4. Using Gaussian Markov Random Fields
  5. Further work

Motivation and background

- **Combination therapies** are becoming increasingly used, especially in areas such as oncology.
- This has brought many new questions:
  - How do we optimise the outcome of a **portfolio** of combinations?
  - Can we **share information** between related combinations?
- However, there is little available methodology to answer these questions specific to combination therapies.

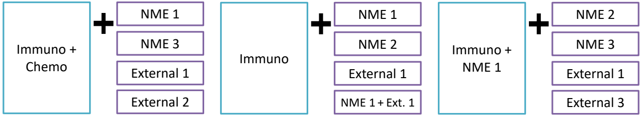

Portfolio-level decision-making - Overview

- Pharmaceutical **portfolio management** addresses decisions such as:
  - which drugs to include in the portfolio;
  - scheduling studies/tasks within the portfolio.
- Several methods exist for portfolio planning which use **stochastic programming**.
- Stochastic programming finds the **optimal decisions** under the **uncertain outcomes** of a process.

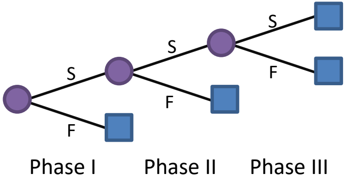

Portfolio-level decision-making - Colvin et al. (2008)

- Colvin et al. (2008) presented a scenario-based multi-stage stochastic programme for the pharmaceutical portfolio management problem.
- This approach:
  - models the uncertainty in **trial outcomes**;
  - aims to maximise the **expected net present value** of the portfolio;
  - returns an **optimal schedule** for each set of trial outcomes upon solving.

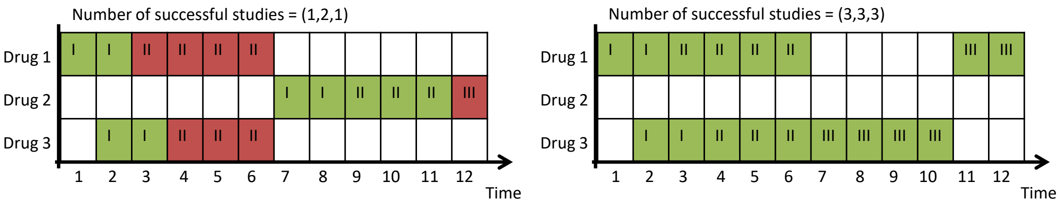

Portfolio-level decision-making - Colvin et al. (2008)

- The **scenarios**, $s \in S$, correspond to the sets of trial outcomes.
- The **stages** correspond to the times at which trials are completed.
- The **decision variables**, X_{ijts}, are binary and X_{ijts} is equal to 1 when trial (i, j) starts at time t in scenario s.
- The **objective function** maximises the expected net present value (ENPV).

$$\text{Maximise} \quad ENPV = \sum_{s} P(s) \left\{ Rv_s(X_{ijts}) + FRv_s(X_{ijts}) - Cst_s(X_{ijts}) \right\}$$

Portfolio-level decision-making - Colvin et al. (2008)

- The **constraints** ensure that:
  - each trial can be performed at most once;
  - resource requirements do not exceed limits;
  - trials are completed in the correct order;
  - programmes do not continue after an unsuccessful study;
  - future outcomes are not anticipated.
- The **optimal solution** is the feasible set of schedules which maximise the ENPV.

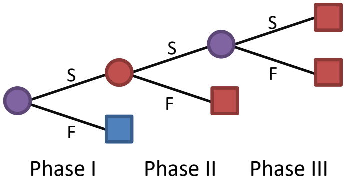

Portfolio-level decision-making - Colvin et al. (2008)

- This framework is very **flexible** and allows for many possible **extensions**.
- For a portfolio of 5 drugs and a planning horizon of 8 time periods this approach:
  - took 7 minutes to generate and 48 minutes to solve;
  - required 50668 variables and 362208 constraints.
- The full formulation can only be solved for portfolios containing up to 6 drugs.

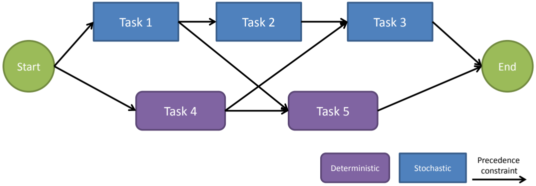

Portfolio-level decision-making - Christian et al. (2015)

- Christian et al. (2015) presented a **knapsack decomposition algorithm** (KDA) as a heuristic for this portfolio management approach.
- The stochastic programme is decomposed into a series of smaller knapsack sub-problems which are solved for each relevant time period.
- In each sub-problem:
  - each **eligible** study is assigned a **value** and a **weight**;
  - **constraints** on resources, overscheduling and eligibility are included;
  - the **objective** is to maximise the value of the selected studies.

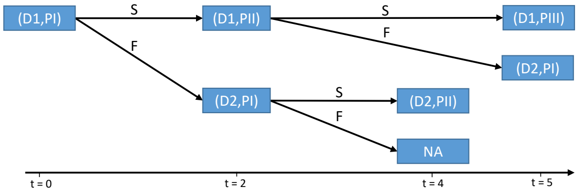

Combination therapies - Outline

- The existing methods do not consider the differences between **single agent** and **combination** drug development.
- Key points to consider include:
  - large possible number of combinations;
  - potential for **sharing information**.

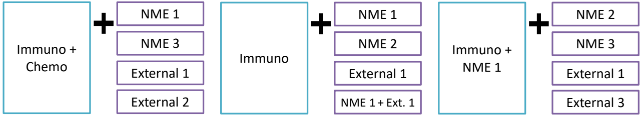

Combination therapies - Outline

- The existing methods do not consider the differences between **single agent** and **combination** drug development.
- Key points to consider include:
  - large possible number of combinations;
  - potential for **sharing information**.

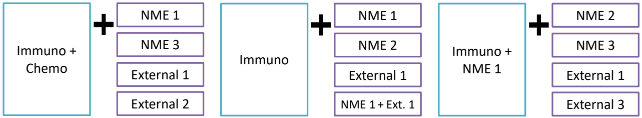

- We consider a Bayesian framework where the **success probability** of a particular combination study is:
  - found using **direct** and **indirect** data;
  - **updated dynamically** throughout the decision making process.

Combination therapies - Multivariate framework

- Let θ_1 and θ_2 represent the treatment effects of two **related combination therapies** which are measured on the same scale.
- Let our **prior beliefs** for $\boldsymbol{\theta} = (\theta_1, \theta_2)^T$ be represented by

$$\begin{pmatrix} \theta_1 \\ \theta_2 \end{pmatrix} \sim N_2 \left(\begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix}, \begin{pmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{pmatrix} \right).$$

- We then want to be able to **update** the prior distribution given observations.

Combination therapies - Multivariate framework

- We can use the **score statistic**, Z_i, and the **Fisher information**, V_i, to summarise the outcome of a study on combination i.
- We can represent the approximate distribution of Z_i by

$$Z_i|\theta_i \sim N\left(V_i\theta_i, V_i\right).$$

- However, we cannot update our prior in the standard Bayesian framework.
- Instead we will consider a **Gaussian Markov Random Field** (GMRF) framework.

Combination therapies - GMRF framework

- Let $\mathbf{x}$ be a finite dimensional random vector with a **conditional independence** structure represented by the graph G.

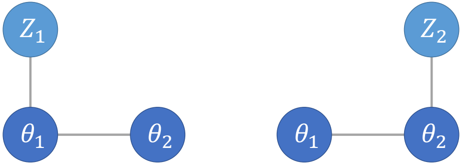

- Then $\mathbf{x}$ is a GMRF with respect to G if and only if

$$\mathbf{x} \sim MVN(\boldsymbol{\mu}, \boldsymbol{\Sigma} = \mathbf{Q}^{-1})$$

and

$$Q_{ij} \neq 0 \quad \Leftrightarrow \quad (i,j) \in E \quad \forall\ i \neq j.$$

Combination therapies - GMRF framework

- Let $\mathbf{x} = (\mathbf{x}_A, \mathbf{x}_B)^T$ be a GMRF with mean $\boldsymbol{\mu} = (\boldsymbol{\mu}_A, \boldsymbol{\mu}_B)^T$ and precision matrix

$$\mathbf{Q} = \begin{pmatrix} \mathbf{Q}_{AA} & \mathbf{Q}_{AB} \\ \mathbf{Q}_{BA} & \mathbf{Q}_{BB} \end{pmatrix}.$$

- Then the **conditional distribution** of $\mathbf{x}_A \mid \mathbf{x}_B$ will be MVN with

$$\boldsymbol{\mu}_{A|B} = \boldsymbol{\mu}_A - \mathbf{Q}_{AA}^{-1}\mathbf{Q}_{AB}(\mathbf{x}_B - \boldsymbol{\mu}_B)$$

$$\mathbf{Q}_{A|B} = \mathbf{Q}_{AA}.$$

Combination therapies - GMRF framework

- Suppose we observe the outcome of a study on combination 2.
- We can write the **score statistic** for this study as $Z_2 = V_2\theta_2 + N(0, V_2)$.
- Then we can write

$$\begin{pmatrix} \theta_1 \\ \theta_2 \\ Z_2 \end{pmatrix} \sim N_3 \left(\begin{pmatrix} \mu_1 \\ \mu_2 \\ V_2\mu_2 \end{pmatrix}, \begin{pmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 & V_2\rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 & V_2\sigma_2^2 \\ V_2\rho\sigma_1\sigma_2 & V_2\sigma_2^2 & V_2^2\sigma_2^2 + V_2 \end{pmatrix} \right).$$

- We can then find the **conditional distribution** of $\boldsymbol{\theta}|Z_2$ using the conditional properties of GMRFs.

$$\begin{pmatrix} \theta_1 \\ \theta_2 \end{pmatrix} \mid Z_2 \sim N_2 \left(\boldsymbol{\mu}_{\text{post}}, \boldsymbol{\Sigma}_{\text{post}} \right)$$

Combination therapies - GMRF simulation study

- We want to compare the decisions made when:
    - we only use **direct data**;
    - we use direct data and **indirect data** via the GMRF framework.
- Let the **probability of success** be given by

$$p_i = P(\theta_i > \theta_i^*)$$

  where θ_i^* is some prespecified threshold.
- **Decision rule:** If $p_i > p_i^*$, run a study on combination i; otherwise do not.

Combination therapies - GMRF simulation study

1. **Generate data** for a small study on each combination.
2. Using this information, specify **informative priors** for θ_1 and θ_2.

$$\begin{aligned}
&(\mathrm{M1}) : \theta_1 \sim N\left(\mu_1, \sigma_1^2\right) \text{ and } \theta_2 \sim N\left(\mu_2, \sigma_2^2\right) \\
&(\mathrm{M2}) : \boldsymbol{\theta} \sim N_2\left(\boldsymbol{\mu}, \boldsymbol{\Sigma}\right) \text{ s.t. } \rho > 0
\end{aligned}$$

3. **Generate data** for a large study on combination 2.
4. Update to find the **posterior** distributions of (M1) and (M2).
5. Find the **success probability**, p_1, under (M1) and (M2).
6. Repeat steps (1)-(5), record **decisions** under (M1) and (M2).

Combination therapies - GMRF simulation study

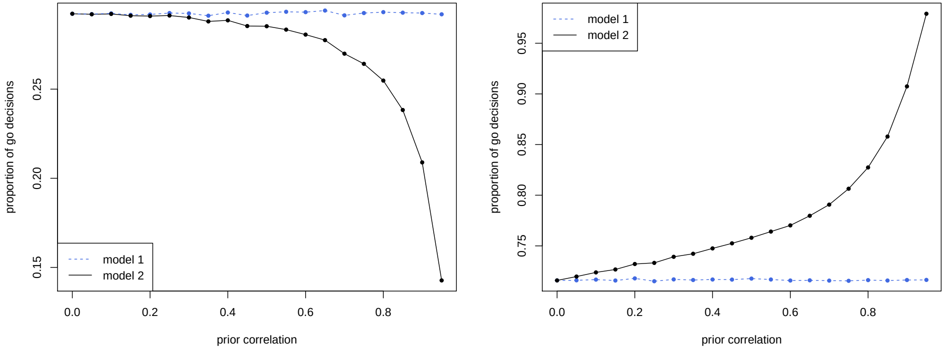

Figure: Proportion of “go” decisions for a study on θ_1 made when $\theta_1 = \theta_2 = 0$ (left) and $\theta_1 = \theta_2 = 0.5$ (right) for varied values of ρ.

Further work - Robustifying the prior

- (M2) leads to **better decision making** than (M1) when θ_1 and θ_2 are similar.
- However, we may want to use a **mixture model** in case θ_1 and θ_2 are different.
- We would consider both (M1) and (M2) and assign each a prior **weight**, ω_1 and ω_2.
- Given a new observation, we would **update** the models and the weights.
- This would allow discrepancies between the data and the correlation structure of (M2) to guide decision making by assigning a higher weight to (M1).

Further work - Portfolio-level decision-making

- Our aim is to improve the **portfolio-level decision-making** process for a portfolio of combination therapies.
- Therefore, our next steps include adding the GMRF method to the KDA.
- Currently, **study success probabilities** are specified at the beginning of the procedure and are fixed.
- We will extend the KDA to **update** the probabilities after each relevant observation.

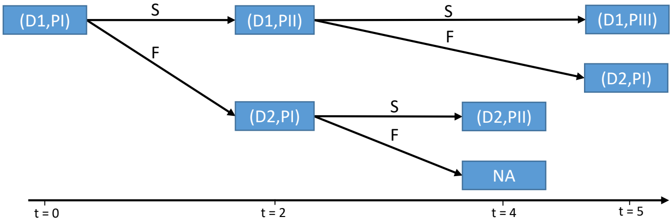

References:

- Christian, B., & Cremaschi, S. (2015). Heuristic solution approaches to the pharmaceutical R&D pipeline management problem. *Computers and Chemical Engineering, 74*, 34-47.
- Colvin, M., & Maravelias, C.T. (2008). A stochastic programming approach for clinical trial planning in new drug development. *Computers and Chemical Engineering, 32*(11), 2626-2642.
- Colvin, M., & Maravelias, C.T. (2011). R&D pipeline management: Task interdependencies and risk management. *European Journal of Operational Research, 215*(3), 616-628.
- Rue, H., & Held, L. (2005). *Gaussian Markov random fields : Theory and applications*. Boca Raton, Fla.: Chapman & Hall/CRC.

Thank you for listening.

Appendix - Christian et al. (2015)

Knapsack Decomposition Algorithm:

Time: $t := 0$
while $t < t^{\max}$ **do**
    Subproblem: $k := 0$
    while $k < |K_t|$ **do**
        Find the set of eligible studies
        Assign values, V_{ijt}, and weights, W_{ij}, to the studies
        Solve sub-problem to find the set of studies to run at time t
        Find the time, t', until the next observation
        Generate set, S, of sub-problems given observations at time $t + t'$
        $K_{t+t'} := K_{t+t'} \cup S$
        $k := k + 1$
    end while
    $t := t + 1$
end while

Appendix - GMRF framework

- Applying the properties of GMRFs we find that

$$\begin{pmatrix}\theta_1\\ \theta_2\end{pmatrix} \mid Z_2 \sim N_2\left(\boldsymbol{\mu}_{\text{post}}, \boldsymbol{\Sigma}_{\text{post}}\right)$$

where

$$\boldsymbol{\mu}_{\text{post}} = \begin{pmatrix}\mu_1 - \frac{\rho\sigma_1\sigma_2 V_2}{1+V_2\sigma_2^2}\mu_2 + \frac{\rho\sigma_1\sigma_2}{1+V_2\sigma_2^2}Z_2\\ \frac{1}{1+V_2\sigma_2^2}\mu_2 + \frac{\sigma_2^2}{1+V_2\sigma_2^2}Z_2\end{pmatrix}$$

$$\boldsymbol{\Sigma}_{\text{post}} = \begin{pmatrix}\sigma_1^2 - \frac{V_2\rho^2\sigma_1^2\sigma_2^2}{1+V_2\sigma_2^2} & \frac{\rho\sigma_1\sigma_2}{1+V_2\sigma_2^2}\\ \frac{\rho\sigma_1\sigma_2}{1+V_2\sigma_2^2} & \frac{\sigma_2^2}{1+V_2\sigma_2^2}\end{pmatrix}$$