Optimizing over Interventions within a Budget

Brian Hunt
University of Maryland
AMSC/MATH 420, Spring 2013

Budget Constraint

- Recall our 2-group SI model with type a and b interventions:

$$
\begin{aligned}
dS_1/dt &= -p_{11} S_1 \mathcal{I}_1 - p_{12} S_1 \mathcal{I}_2 - a_1 S_1 \\
d\mathcal{I}_1/dt &= p_{11} S_1 \mathcal{I}_1 + p_{12} S_1 \mathcal{I}_2 - (a_1 + b_1)\mathcal{I}_1 \\
dS_2/dt &= -p_{21} S_2 \mathcal{I}_1 - p_{22} S_2 \mathcal{I}_2 - a_2 S_2 \\
d\mathcal{I}_2/dt &= p_{21} S_2 \mathcal{I}_1 + p_{22} S_1 \mathcal{I}_2 - (a_2 + b_2)\mathcal{I}_2 .
\end{aligned}
$$

- The impact $M(a_1, a_2, b_1, b_2)$ (fraction of the inital at-risk population saved from infection by the intervention) is an increasing function of each parameter a_1, a_2, b_1, b_2.
- We now consider the problem of maximizing M subject to a “budget” constraint $K(a_1, a_2, b_1, b_2) \leq K_{\max}$ where $K(a_1, a_2, b_1, b_2)$ is the cost of achieving parameters a_1, a_2, b_1, b_2.

Linear Budget Functions

- The simplest class of cost functions are linear functions
  $K(a_1, a_2, b_1, b_2) = c_{a1}a_1 + c_{a2}a_2 + c_{b1}b_1 + c_{b2}b_2 + K_0$
  where the c's are positive constants representing "marginal" costs.
- To simplify further, let's assume $c_{a1} = c_{a2}$ and $c_{b1} = c_{b2}$ and let $c = c_{a1}/c_{b1} = c_{a2}/c_{b2}$.
- Let's normalize (choose units of cost) so that $c_{b1} = c_{b2} = 1$ and choose $K_0 = 0$ (a nonzero K_0 can be subtracted from the constraint, with $K_{\max}$ adjusted accordingly).
- The main flaw in a linear cost function is that it doesn't have the "diminishing returns" observed in real life.

Constrained Optimization

- We are considering a constrained optimization problem; in addition to the constraint $K(a_1, a_2, b_1, b_2) \leq K_{\text{max}}$ we have $a_1, a_2, b_1, b_2 \geq 0$. These inequalities describe a 4-dimensional simplex over which we want to maximize the impact M.
- Since M is an increasing function of the parameters a_1, a_2, b_1, and b_2, the maximum impact will occur when the entire budget is used: $K(a_1, a_2, b_1, b_2) = K_{\text{max}}$. This equality allows one parameter to be determined from the other three, reducing the domain to a 3-dimensional simplex – a tetrahedron.
- The maximum of M often occurs when one or more of the parameters is zero, meaning that the maximum occurs on one of the faces, edges, or vertices of the tetrahedron.

Optimization Strategy

- Iterative, “guess-and-perturb” optimization algorithms are problematic for constrained optimization because the allowed perturbations depend on whether one is inside the constraint domain or on its boundary, and where on the boundary.
- A simpler approach, feasible with a few parameters, is to search the entire domain to a certain resolution – choose a closely-spaced grid and search over the grid points in the domain.
- For best results, you should try to sample all parts of the boundary of the domain. For the tetrahedron described on the previous slide, a normal (rectangular) grid will sample the faces and edges that are aligned with the coordinate axes, but you may need to choose additional points to sample the diagonal face and edges.