Social Investment Policy Flight Simulator

A user's guide to the trade-off in who to help

Dr. James Mansell
Noos Limited

June 2026

A User's Guide to the Targeting Simulator

*No maths background needed. This guide explains, in plain English, why deciding **who** to help is harder than it looks – and how to use the simulator to feel the trade-off for yourself.*

1. What this tool is about

Imagine you run a programme that can genuinely help people – but only some people in the population actually need it, you can't be sure exactly who they are, and every treatment costs money. You have to decide **how many people to treat, and which ones.**

Treat too few and you miss people who needed help. Treat too many and you pour money into people who were never going to benefit. There is no way to avoid this tension – it is built into the nature of the problem. This tool lets you *see* and *play* with it.

Five things shape the answer. The simulator gives you a control for each:

Control	In plain English
Population size	How many people you're choosing from.
Base rate	How common the need actually is (e.g. 8 in every 100).
Discrimination (AUC)	How good your "risk list" is at putting the people who need help near the top. 0.5 = a coin toss; 1.0 = a perfect list.
Who gets treated (the band)	How far down the ranked list you go – your policy choice.
Cost, Effectiveness, Forward liability	What each treatment costs, how often it works, and how much future harm/cost a success prevents.

2. A worked example — the "sweet spot" (scenario S2)

Click **S2 · The sweet spot** at the bottom of the screen. It sets up a realistic, middle-of-the-road situation. Let's walk through the numbers it shows.

The world. A population of **200** people. The condition affects **8%** of them – so **16 people are genuinely in need**, and **184 are fine**. The catch: you can't see who's who. You only have a **risk score** that ranks people from most to least likely to need help. It's a decent score but not perfect (AUC 0.80).

Your choice. You decide to treat the **top 7%** of the list – that's **14 people**, the ones the score rates highest.

What actually happens. Of those 14 treated:

- About **5 genuinely needed it** (true positives) – and you found them.
- About **9 did not** (false positives) – you treated them anyway, because the score flagged them.
- Meanwhile, about **11 people who needed help are still out there** (false negatives) – they were ranked too low to make your top 14.

So even with a decent list, treating the top 7% you **catch only about a third of those in need** (that's the *Sensitivity*: 32%), and **only about a third of the people you treat actually needed it** (that's the *Precision*: 36%). This is normal. It is the base-rate problem: when need is rare, even a good list mixes in a lot of false alarms.

Does it pay? Each treatment costs **$12,000**, so treating 14 people costs **$168,000**. The treatment works about **half the time**, so of your 5 true positives, about **2.5 succeed**. Each success prevents **$100,000** of future cost (the *forward liability*) – so you save about **$255,000**. Net of the cost, you're **$87,000 ahead**, a **return of about $1.52 for every $1 spent.**

That's the whole machine: **find → treat → some succeed → future harm avoided → compare to cost.**

---

3. The central trade-off: how far down the list do you go?

Now the key move. Keep everything else fixed and **drag the band's lower handle** up and down – treating more or fewer people – and watch the two right-hand charts.

- **Return on investment (ROI)** is *highest when you treat very few* – the people at the very top of the list are the surest bets, so every dollar works hardest there. As you treat more, you reach into less-certain cases, precision falls, and ROI steadily drops. Eventually it falls below **$1** – the break-even line – and you're losing money on the margin.
- **Total benefit** (the total dollars saved, all-in) tells a different story. Treat too few and the total is small (you helped only a handful). Treat more and the total climbs… up to a point… then falls as the cost of all those false positives overwhelms the savings. It has a **peak**.

That peak is the **optimal investment point**: the spot where treating one more person would return less than the dollar it costs. **The most *efficient* programme (best ROI) and the programme that does the most *total good* are not the same point** – and good policy usually aims near the total- benefit peak, not the ROI peak.

> **The one-sentence lesson:** better targeting and cheaper, more effective treatment all push the peak further down the list (you can afford to help more people) – but there is *always* a point beyond which casting the net wider destroys value.

A word on "cream-skimming"

There's a second, sneakier choice hidden in the band. Instead of treating the **top** of the list, a provider can slide the *whole band* down – treating, say, people ranked 30–45% instead of the top 15%. They can still claim "we treated 15% of the population," but they've quietly **dodged the hardest cases at the very top**.

Try it: drag the band down the list. Watch the dot on the ROC chart **fall away from the curve**. The curve is the best you can do by being honest (top-down); skimming drops you below it. Your sensitivity and precision collapse, ROI craters – because you skipped exactly the people who would have benefited most. Same apparent effort, far less real good. This is one of the central integrity problems in targeted services, and the chart makes it visible.

---

4. Exercises — work through the five scenarios

Each button at the bottom loads a different world. Load it, read the on-screen lesson, then try the suggested experiment and watch the charts.

S1 · The base-rate trap. *"Our tool is 90% accurate — let's roll it out widely."* It opens treating 20% of the population, **losing money** (ROI below $1) and drowning in false positives. **Try:** narrow the band toward the top. *Notice:* it only pays near the very top — a "90% accurate" tool still fails when the condition is rare. **Accuracy is not enough; the base rate decides precision.**

S2 · The sweet spot. The balanced case from Section 2. **Try:** drag the band wider and narrower and find where **Total benefit** peaks. *Notice:* ROI keeps falling the whole time, but total benefit has a clear summit. **There is an optimum — find it.**

S3 · Just go universal. Treatment is cheap ($300) and the payoff large. **Try:** drag the lower handle all the way to 100%. *Notice:* ROI stays well above $1 and total benefit keeps rising — you should simply **treat everyone**. When a false positive costs almost nothing, **targeting is wasted effort.** (Think immunisation or well-child checks.)

S4 · Surgical targeting only. A rare (1-in-500) but catastrophic outcome you *can* detect well. **Try:** widen the band past about 3%. *Notice:* the profitable window is tiny — it works only if you treat the top sliver. **Catastrophic-but-rare problems are worth tackling, but only with ruthless precision.**

S5 · Don't invest (yet). Same catastrophe, but now nearly **impossible to detect** (AUC 0.60). *Notice:* total benefit is negative at *every* setting — nothing you do with the treatment helps. **Try:** drag the **Discrimination (AUC)** slider upward. *Notice:* as detection improves, the programme becomes viable. **The fix isn't a new treatment threshold — it's better detection.**

---

5. Glossary

- **Base rate / prevalence** — how common the genuine need is in the population.
- **Risk score / risk list** — a ranking of people from most to least likely to need help.
- **AUC (discrimination)** — how good that ranking is. 0.5 = useless (coin toss), 1.0 = perfect.
- **Sensitivity** — of the people who genuinely need help, the share you actually treat.
- **Specificity** — of the people who are fine, the share you correctly leave alone.
- **Precision (PPV)** — of the people you treat, the share who genuinely needed it.
- **False positive** — someone treated who didn't need it. **False negative** — someone missed who did.
- **Effectiveness** — how often the treatment actually works on someone who needed it.
- **Forward liability** — the future cost/harm that a successful treatment prevents.
- **ROI** — dollars saved per dollar spent. Above $1 = paying off; below $1 = losing money.
- **Total benefit** — total dollars saved minus total cost. Its peak is the optimal investment point.
- **Cream-skimming** — treating an easier band lower down the list instead of the genuinely highest-risk people, while still claiming the same coverage.

> **A note on the numbers.** This tool measures benefit in **dollars of avoided future cost** only. Real programmes also create wider wellbeing — better employment, education and health, less crime and suicide — which is real and often far larger, and needs a broader, multi-dimensional approach. The money here is a teaching lens for the *targeting* concepts, not the whole value of helping people.