Effective Healthcare

TL;DR

- We can do healthcare better
- Proposal: the probability we assign a patient to a treatment, should match the probability that that treatment is 'best'
- This would create RCTs of similar patients, and we can use this to improve our understanding of which treatment works best. We can use this to update the assignment ratios
- This will create an automated process where
    - mistakes are quickly discovered and corrected
    - innovations are cautiously explored then, if warranted, quickly incorporated into the standard of care

Introduction

The modern healthcare ecosystem has substantial problems. Examples include the lack of research into off-label treatments, the high cost of running clinical trials, the difficulty that researchers and doctors face in staying up-to-date amidst the millions of papers published each year, and the fact that so little is learnt from the majority of treatment that happens. These problems are different and diverse, but often come back to two root causes: i) it is very expensive to do research, and ii) even when research does get done, it doesn't get used. I believe that both these problems are solvable. In this document, I will propose a unified approach for solving these problems: Probability Matching.

I suspect that running healthcare in this way would have two main effects:
- It would lead to great savings in clinical research spending
- It would lead to far better medical treatment in the medium to long-term.

At the moment, this is written in bullet point form. If anyone asks me to write this up properly, I will do so.

Optimal Healthcare

- The problem
    - We learn nothing from the majority of treatment
        - There is no automatic feedback loop that allows course correction
    - Research is very expensive
        - Few 'off-label' treatments get researched
        - Pharma companies set their prices super high to recoup costs
    - Research is unorganised

        - There are way more scientific papers published each year than anyone can possibly read. Automated systems don't seem to have made much progress (yet). Organisations like the Cochrane collaboration and guideline publishers have stepped up to fill this void, but many blindspots remain and they face an uphill battle to keep their work up to date.

- The problem (but more abstract)
    - This problem has the following properties
        - We have several options (which treatment to prescribe)
        - We don't know which option is best
        - We can gain information about an option by trying it (giving it to a patient)
    - If we suppose we are trying to maximise welfare, this is a 'multi-armed bandit problem'

- The solution (still abstract)
    - Probability matching
        - Multi-armed bandit problems are extremely well studied. Many algorithms for operating successfully have been proposed. I will be discussing Probability Matching (also known as Thompson Sampling)
        - Probability matching is simply the idea that we should balance exploring and exploiting in the following way: the probability that we pick each option should match the probability that that option is best. Each time we try an option we should use that additional information to update the probabilities. Repeat.
        - There are theoretical reasons to think that this is a good solution
        - There are empirical reasons (i.e., based on computer simulations) to think that this is a good solution

- The solution (more concrete)
    - If the problem that healthcare faces is analogous to the multi-armed bandit problem, perhaps a solution to the multi-armed bandit problem would be worth considering in the context of healthcare. Following this analogy, this suggests the following approach to healthcare: **the probability that we assign a patient to a treatment should match the probability that that treatment is 'best' for them**.
        - This would create RCTs of similar patients, and we can use this to improve our understanding of which treatment works best. We can use this to update the assignment ratios
        - This will create an automated process where
            - mistakes are quickly discovered and corrected
            - innovations are cautiously explored then, if warranted, quickly incorporated into the standard of care

Examples

What would this look like? Consider the following examples.

- Example 1 (Correcting a mistake)
    - Suppose the NHS is investigating antivirals for Covid.
        - Suppose they think that 'no treatment' is likely to be better than protocol A
        - Suppose that the evidence objects to this, and suggests that A is better than no treatment
        - Then the NHS would automatically allocate more people protocol A

- Example 2 (Confirming the prior)
    - Suppose a GP frequently treats young men with depression.
        - Suppose he initially thinks there is a 70% chance that treatment A is best
        - He will assign 70% of his patients to treatment A, 30% to other, and will observe how they do.
        - Suppose that the data provides a resounding answer: it is very likely that treatment A is best
        - He will then rejig the allocation ratios to assign more people to treatment A

- Example 3 (Discovering an innovation)
    - Suppose the NHS treats many women in their 60s with condition X for whom the off-label and relatively unknown drug A is best.
    - Suppose that initially the forecasters think there is a 2% chance that drug A is best
    - Since 2% of the women will be assigned drug A we will quickly get data about how the treatments perform. The evidence will (probably) suggest that drug A is best
    - The probability that drug A is best will increase as a result of this evidence to, say, 20%. As this process is repeated, the evidence will continue to mount until the point that it seems quite unlikely that drug A isn't best. Eventually, practically everyone will be treated by drug A

In all examples, understanding of which treatment is best has increased. In examples 1 and 3, the NHS were able to use this increased understanding to improve the treatment of their patients. In all examples, this happens relatively cheaply, and, once the system is in place, without the need for much thought.

More details

- Where will the probabilities come from?
    - We can elicit probabilities using prediction markets

- However, clearly we cannot simply use a prediction market to predict which treatment is best. We won't be able to tell which is best! The best we can do is tell which one seems best. This suggests we get forecasters to predict which treatment will appear best, according to the next patients.
- For example, we can subsidise prediction markets where we ask questions in the form "What is the probability that, for the next 200 men in their 20s being treated for muscle pain, treatment A will reduce muscle pain the most?" Since this will approximate the original question, and has an *objective* resolution criteria, we can use this question as a proxy for the original.
- For discussion, see the open questions at the bottom

- How should we group patients?
  - In the probability matching system, we are assigning people a treatment based on the probability that it is best for them and similar patients. What should that group be?
  - There is a tradeoff between having small fine grained groups where we distinguish between a 21 year old woman and 25 year old woman (which might have some clinical relevance), and having large sample sizes where we can quickly gather data and resolve the questions.
  - Because of this trade-off, it makes sense to focus on factors that will almost certainly be relevant: age, gender and medical history. Anything else, like ethnicity, occupation, etc., will reduce the sample size without (I suspect) adding too much information of clinical relevance.
  - This is just my first pass at answering this question. There should be a more principled answer that balances these trade-offs in a better way. See the open questions at the end.

- What is the scope of this system? What would it apply to?
  - This system is particularly suited to use in calm environments where the healthcare professional has a computer available, time to enter the patients data, and where the situation is similar to ones seen before.
  - It is badly suited for replacing the snap decisions made by paramedics and other first responders day-to-day.
  - These responders should continue doing their thing, but, in the medium to long-term, medical insights learned in the _non emergency_ hospital environment would be included in responders training.

- What would the role of the doctor be in this system?
  - We would still need people with great expertise and experience to act quickly, to respond in emergencies, and to ensure that surgeries and other complicated procedures are performed correctly. We would still need doctors.
  - The burden on the doctors, however, would be substantially reduced. No longer would they need to make mundane repetitive treatment decisions day after day -

the probability matching system would take care of that. Instead, they could focus their efforts on novel situations where speed and judgement matter most.

Open Questions

While I believe this proposal shows promise, I have only been thinking about it for a few weeks, and several open questions remain.

- How ethical is it? It seems easy to justify from a consequentialist perspective, but many deontological schools of thought disagree. This is on the grounds that it is unethical to assign someone a treatment that is believed to be suboptimal. (I pointed out that they are okay with normal clinical trails. They suggested that an opt-in clinical trial is different.)
- Should it be opt-in, opt-out, paid, forced, or other? As I've been thinking this through, I've been reading *Simple*~~*r*~~ by Cass Sunstein. Consequently, I suspect an opt-out approach would work best (although I am open to changing my mind). Is there an approach from mechanism design/social choice theory that would slot in here?
- Do we need diagnoses? Would this system work better if the predictions were based on diagnoses, or if they were based directly on the symptoms that caused the diagnoses?
- What exactly do we mean by 'best'? What should we mean? I am implicitly assuming that, if we knew the distribution of outcomes, it would be clear which treatment is best. That might not be the case. Even if it is the case, different patients might value different outcomes differently. Is this resolvable?
- The forecasts get resolved using 'seems best' rather than 'best'. We don't have an oracle who can tell which treatment is best, so we have to resolve the question according to which hypothesis is most supported by the data. Does this matter? I suspect it would lead forecasters to give less extreme probabilities than they would otherwise. [If they think there is a 100% chance a treatment will be best, they might think there is a 90% chance that that treatment will seem best; if they think there is a 10% chance, they might report a 15% chance; etc. They will be moderated towards 50:50.] Could we correct this? Is it a problem?
- What technology would we need to build for this to be easy? Could we try to investigate treatments for SAD/migraines/depression/etc by making a website that assigns you an off-the-shelf treatment for you to buy and try? Could we build tools to be used in hospitals and surgeries for doctors and other healthcare professionals to use (similar to that used in the RECOVERY trial)?
- How should this handle the complexity of guidelines? How can we tell which 26 step flow diagram for treatment is best? Obviously, we could do this the lazy way: we could assign people to be treated according to the different guidelines and monitor the results. This is possible, but probably not best. The guidelines probably have quite a lot in common. Can we exploit this (or anything else) to make better comparisons between guidelines.
- Is there a principled approach for choosing the patient groups we consider? For each condition, could we measure how much the different factors (age, sex, medical history, occupation, ethnicity, etc.) matter? Can we use this to make the factors that are

accounted for relevant to the condition? What would happen if we did cluster analysis on the whole population?