Social Science & Medicine 276 (2021) 113830

Contents lists available at ScienceDirect

Social Science & Medicine

journal homepage: http://www.elsevier.com/locate/socscimed

Portfolio frontier analysis: Applying mean-variance analysis to health technology assessment for health systems under pressure

Darrin Baines [a], Marta Disegna [a], Christopher A. Hartwell [b,*]

[a] Department of Accounting, Finance and Economics, The Bournemouth University Business School, 89 Holdenhurst Road, Bournemouth BH8 8EB, UK
[b] Zurich University of Applied Sciences (International Management Institute) and Kozminski University, Poland

ARTICLE INFO

JEL classification:
I10
I11
G11

Keywords:
HTA
Portfolio
Mean variance
Health economics
Decision-making

ABSTRACT

The COVID-19 pandemic is challenging how healthcare technologies are evaluated, as new, more dynamic methods are required to test the cost effectiveness of alternative interventions during use rather than before initial adoption. Currently, health technology assessment (HTA) tends to be static and a priori: alternatives are compared before launch, and little evaluation occurs after implementation. We suggest a method that builds upon the current pre-launch HTA procedures by conceptualizing a mean-variance approach to the continuous evaluation of attainable portfolios of interventions in health systems. Our framework uses frontier analysis to identify the desirability of available health interventions so decision makers can choose diverse portfolios based upon information about expected returns and risks. This approach facilitates the extension of existing methods and assessments beyond the traditional concern with pre-adoption data, a much-needed innovation given the challenges posed by COVID-19.

1. Introduction

The coronavirus 2019 (COVID-19) pandemic will have unprecedented health, economic, and social consequences globally (Fernandes, 2020). The current crisis is severely challenging resource-limited health systems worldwide with demands for novel interventions and extra capacity, and the economic slowdown caused by country-wide lockdowns and other measures is diminishing the ability to fund future service provision (Emanuel et al., 2020; Legido-Quigley et al., 2020). Globally, COVID-19 is testing the ability of countries to quickly develop, test, and deploy new medications, with serious concerns being raised about the processes for evaluating and approving drugs (Rome and Avorn, 2020). Given the role of health economics in assessing the cost effectiveness of available interventions, COVID-19 will also create challenges for the health technology assessment (HTA) of innovative treatments (Mukherjee, 2021).

As an applied discipline, health economics has evolved rapidly over recent decades, with notable advances in the development of evaluation techniques (Paulden, 2020). Accompanying this growth has been the increasing use of economic methods in the evaluation of health technologies by bodies such as the UK's National Institute for Health and Care Excellence (NICE) and the Institute for Clinical and Economic Review in America (Thokala et al., 2020). The growing use of economic methods in applied evaluations has led to the continual modification of techniques originally recommended by authorities such as the U.S. Panel on Cost-Effectiveness (Gold, 1996). For instance, best practice originally recommended the use of the cost-effectiveness ratio, which was soon replaced by the Incremental Cost Effectiveness Ratio (ICER) accompanied by the Cost Effectiveness Acceptability Curve (CEAC). Recently, the modelling of Net Health Benefits has been recommended. Approaches that summarize uncertainty in cost-effectiveness analysis using these methods focus only on uncertainty associated with costs and effects of programs under consideration. In the real world, most decision-makers have to fund a portfolio of health care programs. Therefore, a more comprehensive approach would analyze the uncertainty of costs and effects of all programs supported by fixed healthcare budgets (Sendi et al., 2003). In response, we suggest a novel method for supporting health systems decision-makers, which uses real world data to analyze the uncertainty and cost-effectiveness of portfolios of available interventions.

Portfolio theory has an established pedigree as an approach for establishing the efficiency of new health care investments; in particular, Birch and Donaldson proposed integer programming as a method for selecting between mutually exclusive projects (Birch and Donaldson,

* Corresponding author.
E-mail address: christopher.hartwell@zhaw.ch (C.A. Hartwell).

https://doi.org/10.1016/j.socscimed.2021.113830
Received in revised form 24 February 2021; Accepted 5 March 2021
Available online 8 March 2021

1987). Similarly, the mean-variance approach central to portfolio theory has been included in several papers, with Sendi, Gafni and Birch suggesting analysing uncertainty in cost-effectiveness analysis by introducing a new graphical framework (the "decision making plane") for communicating with policy makers about the opportunity costs related to their decisions (Sendi et al., 2002). The authors later extended this approach to the context of a portfolio of programs when costs and effects are uncertain and resources are constrained (Sendi et al., 2003). Gafni, Walter and Birch develop the literature further by illustrating how a decision-making plane may be used to explicitly incorporate opportunity costs into healthcare decision-making, as well as relaxing the assumptions of perfect divisibility and constant returns to scale of the cost-effectiveness plane (Gafni et al., 2013). Our approach builds upon the work of these authors but attempts to add to the literature by suggesting adjustments to portfolio analysis that may make the method more popular to analysts and health systems decision-makers. In particular, we suggest that the concept of "technical change", which is a commonly used framework in mainstream economics, should be adopted by health economists during their HTA work (Elster, 1983).

Akehurst et al. (2017) suggest that all HTA evaluations generally follow five distinct phases: horizon scanning, topic determination, evidence collection and assessment, appraisal, and funding and policy determination. However, the economic evidence currently collected and assessed during these stages is "static" in that only current technologies (and the costs of their associated working practices and means of production) are compared. Even if future costs and benefits are modelled and factor prices are altered, the submitted economic evidence rarely captures the evolving dynamics of technical change. Indeed, a major limitation of current mainstream HTA work is the inability of its static procedures to analyze continuous economic costs and benefits through time.

The purpose of this paper is to present a framework for conceptualizing a mean-variance approach to the continuous evaluation of attainable portfolios of health interventions. The proposed framework is developed using frontier analysis to represent the evolving efficiency of health technologies while in use. The validity of using frontier analysis has been widely established, with robust examples from both HTA (Bradford et al., 2001; Caro et al., 2010) and the measurement of health systems efficiency (Hollingsworth, 2008; Alonso et al., 2015; Sun et al., 2017; Baines, 2018a).

We base our work on the method of representing expected values and uncertainty developed within modern portfolio theory (MPT) as shown in Markowitz (1952). We believe that a modified version of the "mean-variance analysis" performed in MPT can successfully form the basis of what we call "portfolio frontier analysis" (PFA), which we present here as a novel application of MPT for use in economic decision making in continually disturbed health systems.

PFA could make a significant contribution to both research and practice, with a clear potential to generate measurable improvements in health systems performance worldwide. The framework is designed to appeal to a broad interdisciplinary and international readership by avoiding the mathematical treatment of risk and pay-offs found in the finance literature. The approach has relevance for research in health systems and health economics, as well as management science and public health. As health systems in all countries (regardless of their structure and public/private mix) can influence the cost effectiveness of interventions supplied to patients, PFA is applicable internationally and the concept is transferrable to all health settings.

To demonstrate the potential of the PFA approach, this paper proceeds as follows: Section 1.1 discusses static health technology assessment; Sections 1.2 and 1.3 outline how data and modelling may be combined with expert opinion to construct beliefs about the future performance of selected health techniques. Sections 1.4 to 1.6 lay out the PFA framework and its relation to MPT as a way of representing the expected returns and uncertainties of individual health technologies, while Sections 1.7 and 1.8 discusses the usefulness of the PFA approach as a form of post-adoption HTA process and how it could support health system decision-making worldwide. Section 2 concludes.

2. HTA and the static approach

An issue underlying all current approaches to HTA assessment is the use of static projections of dynamic processes. While the health economics literature has attempted to overcome this issue by better modelling of uncertainty, especially in a Bayesian framework (Briggs et al., 1994), there is a manifest need for insights from mainstream economics to better integrate dynamism into evaluation models. A salient example comes from macroeconomics and the approach of Salter (1960), who considers the limitations of using static equilibrium concepts to analyze continuous economic processes and suggests a dynamic alternative. He reports that techniques of production change over time for two reasons: improving technical knowledge and changing factor prices. These are ongoing processes that together create continuous streams of new working methods. Because of this process, Salter argues that "once-over" analysis using comparative statics is only appropriate to changes in techniques that are sufficiently great to totally displace all pre-existing methods. For instance, the replacement of the typewriter by the personal computer would have been worth evaluating within a static framework because typewriter technology came to an end as a viable form of production with this innovation.

By applying Salter's reasoning, new health technologies should not be evaluated using once-over analysis unless the new technology is sufficiently disruptive (this is important because data from clinical trials and results from pre-adoption modelling rarely capture the cumulative benefits slowly delivered by improving technical knowledge and changing factor prices). On the other hand, Salter (1960) argues that the cumulative benefits of small unnoticed modifications and improvements in production methods or adaptation of existing technologies may be equally as great as the significant changes created by discrete large-scale innovations. This dynamic may improve health system efficiency by increasing benefits and controlling rising costs through reallocation of resources, recognizing the actual changes occurring in factor prices. Indeed, the slow but continuous change of factor prices is often sufficient to produce a constant stream of new techniques of production with significant economic benefits. For example, with no intervention for COVID-19 available, health systems may find that operational improvements may initially stem from thousands of small changes in working practices rather than from large gains from a "blockbuster" vaccine.

Any technology may be evaluated post-launch. However, Fig. 1 suggests that technologies in the first and third quadrants of the cost-effectiveness plane may be particularly good candidates for post-adoption analysis because their economic profiles may change over time. Indeed, HTA bodies cannot know, a priori, how the process of (what Salter calls) "continuous disturbance" will affect the discounted net value of technologies over time. Long-term patient flows may be represented in Markov modelling or budget impact analysis, but the dynamics of constant technical change are much harder to predict and capture. Consequently, post-adoption economic analysis may be required. Moreover, guidelines that recommend the universal adoption of new technologies soon after a positive HTA assessment may "crowd-out" the slow process of change that may have occurred otherwise. For instance, the requirement that the UK National Health Service must fully implement NICE guidance within three months may affect the ways in which knowledge generation and price modifications occur, particularly for novel interventions. As these options suggest, PFA may be used on all technologies post-launch, and HTA bodies must decide policies that determine their approach to this type of continuous evaluation.

The PFA offers an important methodological development for HTA, particularly during the COVID-19 pandemic. The approach is important because it shifts the emphasis from static, pre-launch evaluation to the dynamic analysis of interventions in use. Doing so, the approach moves

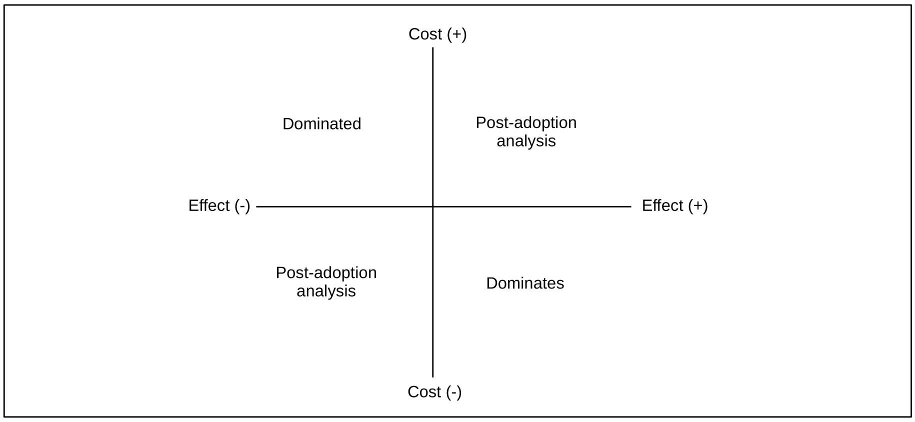

Fig. 1. Cost effectiveness plane for post adoption analysis.

attention from efficacy to effectiveness, and takes advantage of the growing availability of real-world data (RWD). At present, there is no "gold standard" for using post-launch RWD for evaluating the relative cost-effectiveness of healthcare interventions (Makady et al., 2017). PFA could become the method of choice for using real world evaluation in HTA work.

To date, portfolio approaches to HTA may have failed to gain popularity because of a lack of robust theoretical support. Economic evaluations in health care are usually supported by welfarist and extra-welfarist theoretical frameworks (Brouwer et al., 2008). In contrast, Salter's approach to conceptual approach to technical change offers a new theoretical foundation for HTA, which employs methods to capture dynamic gains accrued during use. As an example, many small changes to dietary advice, the use of statins and vitamins, better guidelines for obesity care and diabetes, combined, may be more cost-effective in "at risk" patient groups at preventing COVID-19 related deaths than vaccination. Therefore, we believe that applying mean-variance analysis to HTA for health systems under pressure may benefit both decision-makers and patients.

3. Shifting the curve: incorporating a portfolio approach

Salter's analytical framework is based upon the concept of "best-practice techniques" (BPTs), which he defines as the methods that yield "minimum costs in terms of the production function and relative factor prices of each date." (Salter, 1960) (p.23) In economic terms, we could say that BPTs create outputs located on the boundary of an efficiency frontier, with less productive methods creating outputs located inside, at a distance from what could be achieved by being like the best.

3.1. Mean-variance analysis

Given the importance of BPTs in evaluating efficiency in continually disturbed economic systems, we adopt this approach in our framework but operationalize Salter's concept specifically in the context of healthcare. We focus on employing the data generated, tested, and/or verified by health economists on the post-adoption evaluation of health technologies so that the measurement of their economic value and understanding of uncertainty becomes an ongoing process.

To explain how these existing methods of data generation and evidence verification may be applied differently, we propose using the seminal work on modern portfolio theory by Markowitz (1952). His key insight is that expected returns should not be evaluated in isolation but should be simultaneously assessed alongside their associated risks. To do this, Markowitz suggests "mean-variance analysis" as a way of representing investment portfolios in terms of expected risks and returns. In this theory, the rate of returns (r) on assets are random variables of which the first two moments (i.e., expected value and variance) are used to obtain the optimal weighted portfolio. This theory can be analytically formulated as the following constrained optimization problem:

$$\begin{cases} min : \dfrac{1}{2}\sum_{i\epsilon P}\sum_{j\epsilon P} w_i w_j \sigma_{ij} \\ s.t. \sum_{i\epsilon P} w_i \mu_i \geq \mu_b, \sum_{i\epsilon P} w_i = 1 \end{cases} \tag{1}$$

where P represents the number of assets constructing a portfolio, w_i is the weight representing the investment proportion for the ith asset ($i = 1, \ldots, P$), σ_{ij} indicates the covariance between the ith and jth assets, $\mu_i = E(r_i)$ is the expected value of the rate of return from asset i, and μ_b is the target baseline expected rate of return.

Markowitz's framework has significant similarities to the processes usually used for HTA. For instance, both measure "discounted expected returns" to establish the efficiency of investments or spending. Similarly, both use a form of "variance of return" to represent the reliability of deterministic estimates of discounted investment yields or cost-effectiveness calculations. Markowitz (1999) reports that the standard deviation is the most intuitively meaningful measure of dispersion in investment portfolios.

Our framework modifies Markowitz's work on finance for use in HTA and health systems. However, Markowitz focuses on formal financial markets that usually have lower and upper limits on the rate of returns and risks acceptable to investors. In health systems, minimum levels of expected cost effectiveness and maximum degrees of uncertainty are not institutionalized in the same way. The absence of these limits makes Markowitz's model harder to employ in a health systems context. In response, we address this problem by introducing the notion of an "economic floor" and an "uncertainty ceiling" into our framework.

A key feature of Markowitz's work is the way he links risk-free investments with an efficient frontier to create his "capital market line". We modify his approach to introduce a similar decision-making frontier into PFA and innovate the practical application of MPT by representing the dynamic movements of health technologies within the payoff-uncertainty space.

A brief example of the problem of estimating technological effectiveness under uncertainty and, more importantly, volatility, is shown in Fig. 2 where rate of return (r) is a variable that, in repeated sampling, will provide different numerical estimates that may be represented diagrammatically as a probability distribution. As r is not constant, healthcare purchasers will face uncertainty in the rate of returns they can accrue from the interventions they purchase. This is shown in the diagram where interventions one, two and three all have the same dispersion in r (shown by similar normal distribution curves), even though they have differing expected values. In contrast, intervention four has a much larger return but is accompanied by greater uncertainty because r is more variable around its expected value. As this example

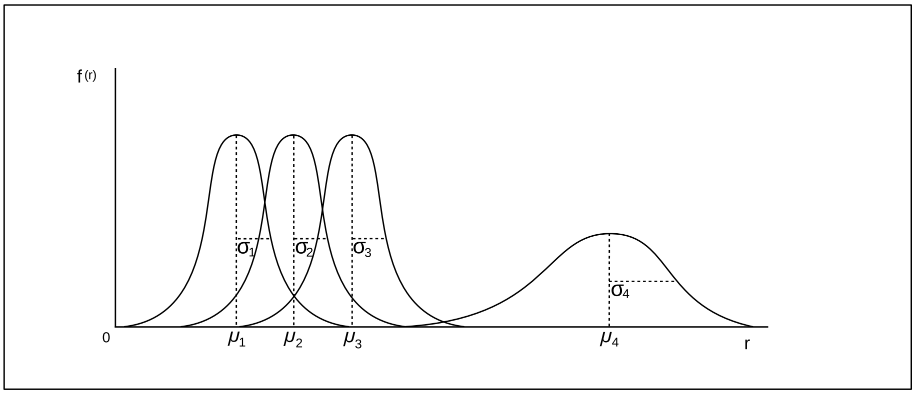

Fig. 2. Probability distribution of rate of return (r).

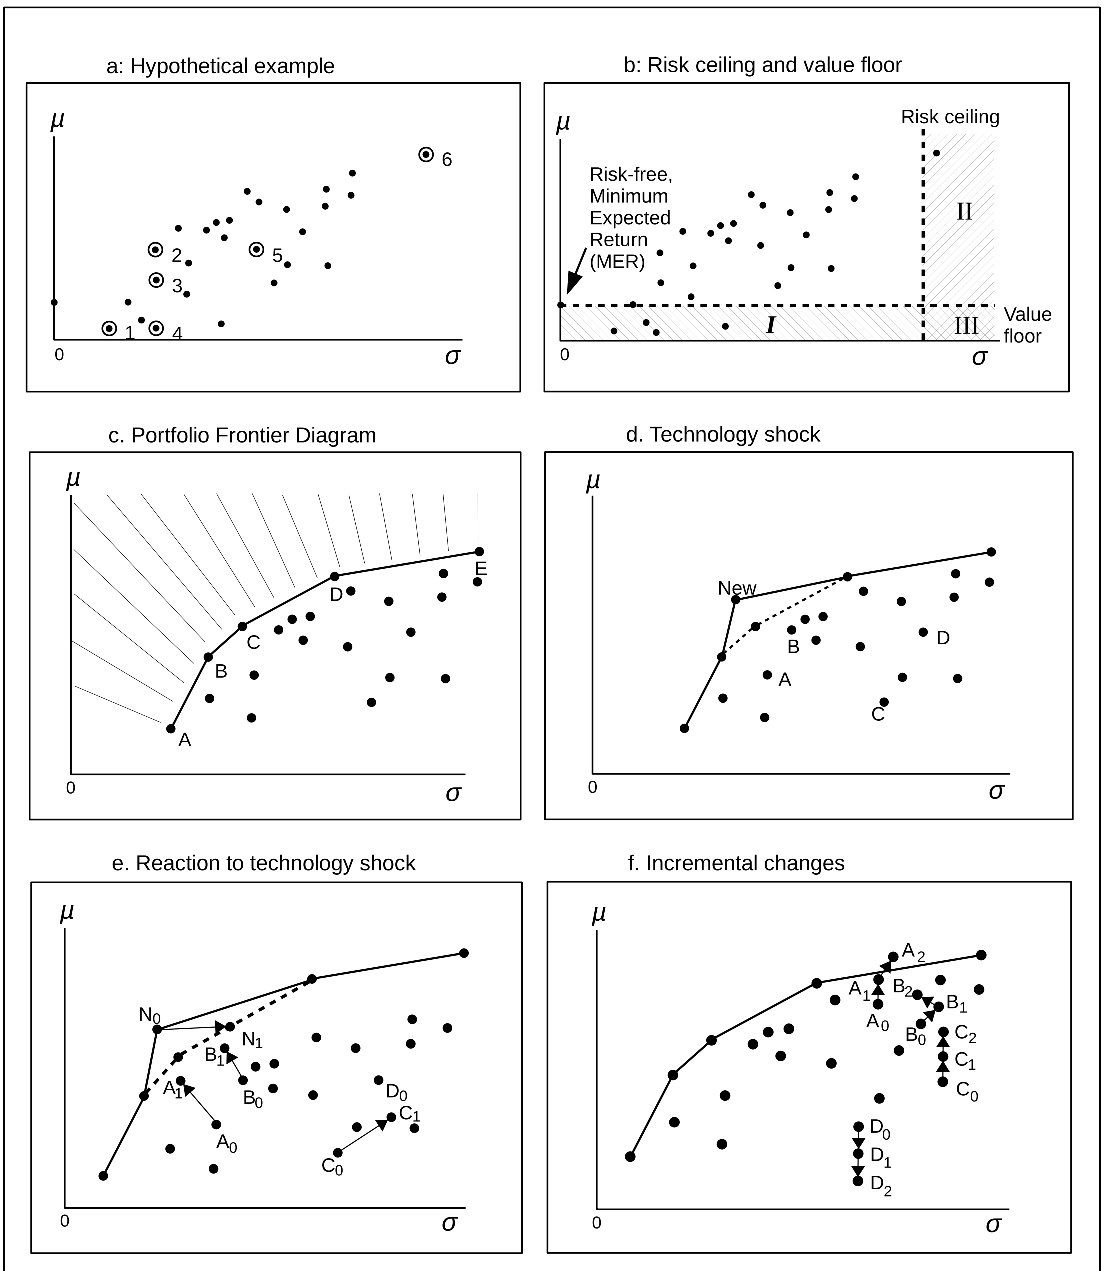

Fig. 3. Graphical representation of expected value (μ) and risk (σ).

suggests, r is not the only variable that healthcare purchasers should consider. Just like financial investments, differing rates of return possess different risks. For instance, the options that offer the largest payoffs often have the greatest dispersion of observed rewards.

The challenge is thus (following Markowitz) to create an optimal portfolio as in Equation (1), which achieves an acceptable baseline expected rate of return on an asset with minimal variance (for the purpose of this research, an "asset" in the healthcare sense of the optimal portfolio is any health intervention, such as a new technology, drug, or policy). In the health system, a portfolio of chosen health interventions is "mean-uncertainty efficient" if it maximizes the expected value of the rate of return of the portfolio for a given uncertainty (i.e., portfolio variance) or, equivalently, minimizes the uncertainty for a given portfolio expected rate of return.

A difficulty in this model is that the expected value of each asset and the variance–covariance matrix of the portfolio are assumed to be known. However, following Salter's (Salter, 1960) description of dynamic technical change, it is unlikely that these parameters will be known until an intervention is used. This is the case for three main reasons. First, pre-adoption evidence is usually based upon clinical trial or modelling results that are isolated from the continual economic disturbances experienced in the real world. Next, manufacturers may suppress unfavorable evidence during HTA submissions. Finally, health systems are dynamic and reactive, so technical processes of production and factor prices will continuously alter in unpredictable ways.

4. From portfolios to frontiers

Fig. 3a shows the first step in our application of Markowitz's approach to portfolio decision making in health systems, which involves plotting all available interventions in an expected value (μ) - risk (σ) space. If we compare points one and two, the former intervention has the same expected return as the latter, with both having an almost zero yield. However, the former intervention has a significantly lower dispersion of expected returns. In contrast to point two, points three and four have identical levels of uncertainty but higher expected returns. Therefore, points three and four are preferred to point two because they perform better. If we continue applying this logic, the diagram shows that point four is preferred to point five because they have identical returns, but the former has lower uncertainty. Next, point six is interesting as it has the largest return and uncertainty by far. Given its outlying position, a comparator is not obvious, and we need to employ an appropriate decision rule to determine whether its expected returns are worth the large variability in its results. As a means of dealing with such situations, the next stage in the presentation of our PFA is showing how appropriate decision rules may be applied to portfolio choice.

4.1. Floor and ceiling constraints

To generate more realistic models, one can consider additional constraints, among which the floor and ceiling constraints are relatively common in the literature. The floor and ceiling constraints define lower and upper levels per intervention investment representing health system decision makers' preferences. In our PFA, the floor constraint is an "economic" constraint that we can call "economic floor", defined by the lowest acceptable level of expected value. Most HTA processes reject interventions with low expected economic returns, so not all interventions assessed are subsequently adopted. In our framework, we follow this approach and suggest that health system decision makers should consider excluding the least cost-effective interventions from their portfolios. In making this recommendation, we are assuming perfect divisibility and constant returns to scale. If these assumptions are relaxed, it may be efficient to include the least cost-effective interventions in a portfolio depending upon their size and their (average) rate of return (Gafni et al., 2013). However, we do not explore this option here.

As Fig. 3b shows, we can apply a minimum threshold to our hypothetical data, which sets a floor for economic acceptability. In our example, the chosen floor leads to the rejection of the points in the shaded areas I and III. As well as determining which interventions to exclude, the floor also helps set (what we call) the minimum expected return (MER), which occurs where the floor intersects the horizontal axis. This position is chosen because it is located where the floor coincides with the point representing zero uncertainty. Consequently, MER is an important reference point because all other acceptable points have either greater returns, higher levels of uncertainty, or a combination of both.

The ceiling constraint is an "uncertainty" constraint that we call the "uncertainty ceiling", defined as the maximum acceptable level of risk. As well as avoiding treatments with low expected returns, health systems may consider some interventions too unstable in their expected results to be used routinely (outside of experimental or emergency cases). In response, our framework contains an uncertainty ceiling above which decision makers are unwilling to fund treatments. As Fig. 3b shows, any intervention in areas II or III are considered too uncertain for routine use because they lie above the ceiling. Therefore, interventions in these areas should be excluded from purchasing plans. Finally, if we combine the "uncertainty ceiling" with the "economic floor", Fig. 3b shows which interventions should and should not be considered for purchasing. In practice, the uncertainty ceiling will not be a fixed value but will depend on the opportunity cost of choosing one intervention over another (Palmer and Raftery, 1999). Consequently, PFA raises many new questions and calls for a new, pragmatic approach to HTA. For instance, for smaller investments, would a decision-maker be willing to accept a higher level of uncertainty than she would accept for larger investments? Moreover, new means of quantifying benefits forgone are needed for PFA and other forms of health economics analysis (Sculpher et al., 2017).

4.2. PFA diagram

Fig. 3c shows the interventions available for evaluation after ineffective and risky treatments have been removed (i.e., any interventions in areas I, II, and III of Fig. 3b are not considered). We can now see points in the expected value (μ) - risk (σ) space that have different combinations of efficiency and uncertainty. Points with higher returns or lower levels of uncertainty are always preferred. Consequently, we draw an efficiency frontier that links the most preferred points, which is shown by the series of straight lines connecting A, B, C, D, and E. The area to the north-west of the frontier is unattainable because interventions do not exist in this space. This is identified by the shading outside of the frontier. In contrast, the points shown within the frontier are attainable. However, they have lower expected returns and/or higher levels of uncertainty than those at the frontier.

4.3. Examples of application of the basic model

Fig. 3d shows our hypothetical data, with points A, B, C, and D representing alternative interventions for treating the same condition, with the four interventions lying within the portfolio frontier. Following a scientific breakthrough, a new intervention is launched that pushes the frontier outwards. This is a technological shock because its combination of expected return and uncertainty far outperforms those offered by the existing interventions. The points represented in Fig. 3d do not portray the dynamic nature of continually disturbed health systems. In contrast, Fig. 3e shows how the five interventions evolve from their baseline positions at A_0, B_0, C_0, D_0, and N_0 to their locations at A_1, B_1, C_1, D_0, and N_1. First, the launch of the new intervention stimulates the manufacturers of A to make their intervention more cost effective and less variable in its expected return. As a result, A moves to just inside the dotted line and rests at A_1, which coincides with the frontier attainable before the new treatment was launched. Similarly, changes in technical processes make B_1 more cost effective and less uncertain than B_0, pushing the intervention nearer the previous frontier.

Although A_1 and B_1 are superior to A_0 and B_0, they are not preferable to N_0. While this implies that the new intervention should be chosen, data collected during the first time period suggest results from the new intervention are more uncertain than expected. At the end of the first time period, N has failed to push the portfolio frontier outwards, the observed technology shock is not sustained, and the intervention settles at N_1. Even though the three interventions hold different locations in the E, U space, they are all on or close to the frontier. Therefore, A and B may be near equivalents to N. In such situations, health systems decision makers must weigh the pros and cons of their available portfolio of interventions. For instance, they may choose A over B and N because the former offers the lowest variability in results, even if this choice results in efficiency levels being lower than possible with the latter two options. Conversely, in more time-constrained situations, as with the COVID-19 outbreak, efficiency levels rather than variability may be the metric of success.

As Fig. 3e shows, HTA decisions made at launch may not reflect the changing dynamic of continual disturbance frequently observed within health systems. For instance, new interventions may fail to perform as well as expected, stimulate improvements in competitor treatments, and shift from being dominant at launch to being equivalent during use. In drawing these conclusions, we must also acknowledge that existing technologies do not always respond positively to the launch of seemingly superior alternatives. For instance, C improves in cost effectiveness compared with C_1, but this is only achieved by reformulating the product in a way that produces greater variability in its results. As another comparison, intervention D remains stuck at D_0 because it cannot evolve. In sum, C and D would not be interventions of choice when N is launched, and continuous economic evaluation confirms that the choice not to use them remains correct.

Fig. 3f shows the effects of continual small reactions and modifications. The four interventions shown are stimulated by evolving technical processes and, crucially, the evolution of factor prices exogenous to the health sector. For instance, A and C improve their rate of return during period one without increasing the uncertainty of their results. In contrast, intervention B only achieves an improvement in its expected return by adopting methods that increase the variability of its outcomes, while D experiences a drop. At points A_1, B_1, C_1, and D_1, the interventions are located differently to their original positions. Sensing the ability to catch up, the manufacturers of intervention B modulate their production processes to move their product nearer to the position of A_1. This move affects the relative position of C, the producers of which try to gain ground by lowering its price, thus ending up at C_2. Sensing a shifting market dynamic, the makers of A push their product to be more cost effective, moving from A_1 to A_2. In doing so, they outperform their existing competitors, shift the portfolio frontier outwards, and create gains by accepting more uncertainty in their results. At the other end of the spectrum, intervention D performs increasingly worse over time, with the consequence that it should probably be excluded from future purchasing plans when it reaches D_2.

4.4. Decision makers

Fig. 4 presents a more sophisticated form of our PFA that again includes the MER point in the diagram. The minimum expected return is located where any asset would give zero variation in its rate of return and offer the minimum return expected by health system purchasers. This is an important benchmark because MER is a corner solution that indicates the minimum return, zero uncertainty position acceptable to health systems purchasers. MER can be used as an anchor point when forming a new efficiency frontier, which is tangent with the most efficient intervention available, shown by A. This frontier differs from other efficiency frontiers because its position is partly determined by decision makers themselves, whereas the previous frontiers were located solely by the position of the observed interventions. In response, we label this new frontier the "decision-makers' frontier" (DMF). Given their ability to set MER via modulations in the economy, DMF is a frontier that health systems can control, subject to the location of existing technologies in the analytic space.

The DMF is an economic frontier that may rarely be observed in practice but may be useful in shaping abstract models of decision making. Its key features may be interpreted as follows. First, the frontier's connection with the vertical axis represents an intervention with no uncertainty in its expected returns. No variability in results should always be preferred to variability, which makes the MER an important guide to optimal choice. Next, the tangency of DMF with point A represents an uncertain return that is equivalent to MER, which gives a higher pay-off to compensate for greater variability. In other words, the slope of the DMF is equivalent to the rate of exchange between expected returns and uncertainty that would be observed if health systems operated at their DMF frontier. Finally, all points on the DMF (except A and the MER) are hypothetical and represent the most efficient combinations of μ and σ that could feasibly exist, given MER and A. The DMF acts as a frontier against which to judge the relative worth of interventions available to health systems purchasers.

PFA has practical applications for health systems decision-makers. Currently, HTA information is generated pre-launch usually by agencies outside the management structures of the organizations that purchase and supply interventions to patients. Given its structure, PFA allows decision-makers to set floor and ceiling values for effectiveness and risk. Control is also given over the MER and the decision-makers' frontier. As PFA evaluates risk and returns during the dynamic process of use, it suggests that bodies running the day-to-day operations of health systems should be in charge of implementing this method in practice. For instance, PFA should be operated by NHS England rather than NICE in the UK, which would give the former greater control of efficiencies,

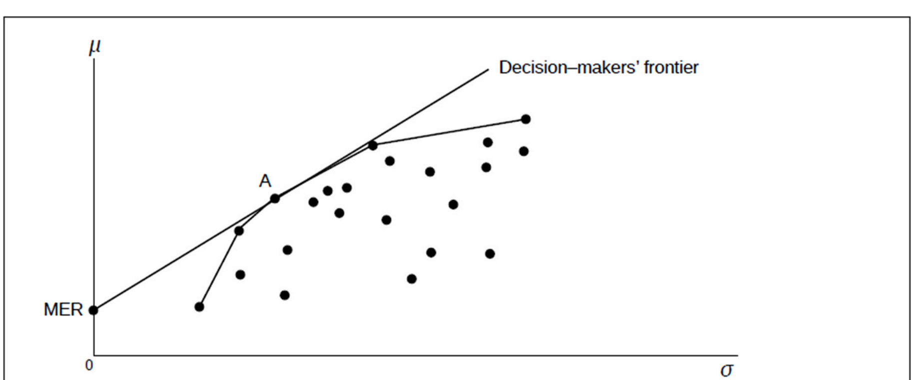

Fig. 4. Decision-makers' frontier (DMF) with expected value (μ) and risk (σ).

especially in relation to high-cost drugs (Baines, 2018b). To enable such applications, researchers need to identify suitable forms of RWD useable in different health systems that facilitates the successful implementation of PFA.

The implementation of PFA will present new challenges for researchers and health systems funders. First, there will be the issue of identifying which interventions best suit this approach. For instance, products that are not dominant or dominated in the cost-effectiveness plane may be suitable candidates for post launch analysis. However, cost-effective interventions may also require continuous evaluation because their competitors may change methods of delivery and prices in order to compete with their new rivals. Alternatively, areas of medicine where innovative products are rare, but healthcare professionals make constant, but small improvements in their performance may be good areas for implementing PFA. In the short term, decision-makers may choose to focus on specific groups of interventions as a means of building experiencing of HTA using a portfolio approach. For instance, attention will need to be paid to how information about expected values and risk going beyond using these "average" values and to explore the role of opportunity cost in portfolios. In pre-launch HTA, the recommendation of adoption usually involves strict criteria for use and assumptions about cost effectiveness thresholds that are rarely acknowledged once adoption occurs. With PFA, the way the product is used in conjunction with decision-makers economic criteria will constantly inform choices about which interventions to use. Therefore, PFA makes HTA a "living" process that health system decision-makers can continuously control.

Finally, the possible use of Value of Information (VoI) studies in the PFA is useful to discuss (Claxton and Sculpher, 2006). VoI analysis provides an analytic framework for establishing the value of acquiring additional information to inform a decision problem. If decision-maker aim to maximise health outcomes subject to a budget constraint, then the choice to adopt or reimburse an alternative should be based on expected cost, expected outcomes and the cost effectiveness threshold. With perfect information, decision-makers can select interventions that maximise the net benefits of adoption for a particular health threshold. The expected value of perfect information is simply the difference between the expected value of decisions made with perfect information and the decisions made on the basis of existing evidence. When further research appears to be worthwhile, the approach indicates where evidence about particular parameters will be most valuable and suggests the type of research design which might be most important. Therefore, we suggest that the possibility of using VoI studies to support PFA should be further explored.

5. Conclusion

In suggesting the use of PFA, we are aware of the robustness of the current HTA methods employed worldwide, so full justification for preferring PFA is required. First, we believe our novel approach could supplement existing health economics and HTA methods, especially given the challenges of the COVID-19 crisis. Even when comprehensive modelling studies are undertaken, mainstream evaluations often fail to capture the dynamics of the evolutionary economic systems determining health systems outcomes. Currently, most formal assessments of health economics evidence are performed outside of the health systems they are designed to serve, usually by independent HTA bodies. As a result, their recommendations are usually based upon once-over evaluations, which may have diminishing validity over time as their findings are challenged by continual disturbances slowly operating within health systems.

As noted, techniques of production change over time as technical knowledge and factor prices alter (Salter, 1960). In health systems, these drivers generate continuous streams of new working methods. For instance, healthcare professionals will continually learn how to generate incremental benefits from existing techniques until an unexpected technology shock disrupts the technical and economic dynamics of their production processes. In most instances, slow but continuous change will generate the most significant economic gains. Unfortunately, current methods rarely focus on the technical changes made by market incumbents in response to new interventions and modulating factor prices.

The operation of the complex economic processes that shape health systems outcomes is rarely observed for economic evaluations. As once-over, pre-adoption evaluations cannot capture the cumulative benefits of small improvements in production methods, PFA could be a useful addition to current health economics and HTA methods. In proposing PFA as a form of post-adoption analysis, we wish to refocus attention on analyzing the effects of the underlying dynamics that drive health systems. Indeed, measuring the impact of new and existing technologies during their use could be informative because, often, the full benefits of innovation emerge over time because of continued learning from use. Only with the gift of complete foresight can pre-adoption evaluations truly predict the economic consequences of allowing new technologies to enter the market. Given that we live in a second-best world, post-adoption methods such as PFA could become useful additions to existing approaches to health systems economics and HTA.

PFA may be a useful tool for post-adoption decision making in health systems in five main ways: (1) the PFA diagram represents combinations of μ and σ that can be attained with available interventions, (2) the approach can help analysts separate efficient from inefficient interventions, (3) the framework could help decision makers select the combinations of μ and σ that best suit the needs of their health systems, (4) the approach can help determine which interventions are the most suitable for particular patient groups, and (5) its ability to use real-world data means that PFA can reflect the continual disturbances that shape the economics of healthcare provision.

Widespread adoption of PFA has the potential to affect the mainstream working procedures of HTA in the following ways. First, the use of PFA implies that only potentially "dominant" technologies should be subject to full-scale pre-adoption assessments. Consequently, interventions that are likely to be equivocal in their cost effectiveness should be evaluated in use and withdrawn from purchasing portfolios if they perform poorly. Next, if HTA bodies only perform pre-adoption assessments on technologies that shock existing production processes, the job of evaluating all other interventions could be a regular activity of health systems themselves. Used in a feedback loop, the methodology could be employed by health systems in all countries (regardless of their structure and public/private mix) to improve efficiency and reduce uncertainty. Moreover, as it does not require inputs from a formal HTA agency, PFA could be performed by analysts working within (or for) health systems themselves, who could employ local management structures to encourage providers to improve the results being monitored. Finally, new sources of real-world data and methods of analytics will be required to operationalize PFA. Although this will be costly, cash-limited health systems have little option but to actively manage their growing portfolios of interventions if they wish to successfully stay within budget limits. Therefore, with its dynamic nature, PFA may be an appropriate support tool for health systems decision makers concerned with the returns and risks of their spending.

Credit author statement

Darrin Baines: Conceptualization; Formal analysis; Methodology; Writing – original draft, Marta Disegna: Methodology; Visualization; Writing – original draft, Christopher Hartwell: Conceptualization; Project administration; Writing – review & editing.

References

Akehurst, R.L., Abadie, E., Renaudin, N., Sarkozy, F., 2017. Variation in health technology assessment and reimbursement processes in Europe. Jan 1 Value Health 20 (1), 67–76.

Alonso, J.M., Clifton, J., Díaz-Fuentes, D., 2015. The impact of new public management on efficiency: an analysis of Madrid's hospitals. Health Pol. 119 (3), 333–340.
Baines D. The Frontier Framework (and its eight Frontier Archetypes): a new conceptual approach to representing staff and patient well-being in health systems. Soc. Sci. Med.. 2018 May 4;208(July):98-106.
Baines, D., 2018b. High-cost medicines in England. Jan 1 InEquitable Access to High-Cost Pharmaceuticals. Academic Press, pp. 23–34.
Birch, S., Donaldson, C., 1987. Applications of cost-benefit analysis to health care: departures from welfare economic theory. Sep. 1 J. Health Econ. 6 (3), 211–225.
Bradford, W.D., Kleit, A.N., Krousel-Wood, M.A., Re, R.N., 2001. Stochastic frontier estimation of cost models within the hospital. Rev. Econ. Stat. 83 (2), 302–309.
Briggs, A., Sculpher, M., Buxton, M., 1994. Uncertainty in the economic evaluation of health care technologies: the role of sensitivity analysis. Health Econ. 3 (2), 95–104.
Brouwer, W.B., Culyer, A.J., van Exel, N.J., Rutten, F.F., 2008. Welfarism vs. extra-welfarism. Mar 1 J. Health Econ. 27 (2), 325–338.
Caro, J.J., Nord, E., Siebert, U., McGuire, A., McGregor, M., Henry, D., et al., 2010. The efficiency frontier approach to economic evaluation of health-care interventions. Health Econ. 19 (10), 1117–1127.
Claxton, K.P., Sculpher, M.J., 2006. Using value of information analysis to prioritise health research. Nov Pharmacoeconomics 24 (11), 1055–1068.
Elster J. Explaining Technical Change: A Case Study in the Philosophy of Science. CUP Archive; 1983 Jun 9.
Emanuel, E.J., Persad, G., Upshur, R., Thome, B., Parker, M., Glickman, A., et al., 2020. Fair allocation of scarce medical resources in the time of Covid-19. Epub 2020 March 23 N. Engl. J. Med.. https://doi.org/10.1056/NEJMsb2005114.
Fernandes, N., 2020. Economic Effects of Coronavirus Outbreak (COVID-19) on the World Economy. Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3557504. (Accessed 30 April 2020).
Gafni, A., Walter, S., Birch, S., 2013. Uncertainty and the decision maker: assessing and managing the risk of undesirable outcomes. Nov Health Econ. 22 (11), 1287–1294.
Gold, M., 1996. Panel on cost-effectiveness in health and medicine. Dec 1 Med. Care DS197–D199.
Hollingsworth, B., 2008. The measurement of efficiency and productivity of health care delivery. Health Econ. 17 (10), 1107–1128.
Legido-Quigley, H., Asgari, N., Teo, Y.Y., Leung, G.M., Oshitani, H., Fukuda, K., et al., 2020. Are high-performing health systems resilient against the COVID-19 epidemic? Lancet 395 (10227), 848–850.
Makady, A., Ten Ham, R., de Boer, A., Hillege, H., Klungel, O., Goettsch, W., 2017. Policies for use of real-world data in health technology assessment (HTA): a comparative study of six HTA agencies. Apr 1 Value Health 20 (4), 520–532.
Markowitz, H., 1952. Portfolio selection. J. Finance 7 (1), 77–91.
Markowitz, H.M., 1999. The early history of portfolio theory: 1600–1960. Jul 1 Financ. Anal. J. 55 (4), 5–16.
Mukherjee K. Integrating Technology, Innovation and Policy: COVID-19 and HTA. Health Policy and Technology. 2021 Feb 4.
Palmer, S., Raftery, J., 1999. Opportunity cost. Jun 5 BMJ 318 (7197), 1551–1552.
Paulden, M., 2020. Calculating and interpreting ICERs and net benefit. Aug Pharmacoeconomics 38, 785–807.
Rome BN, Avorn J. Drug evaluation during the Covid-19 pandemic. N. Engl. J. Med.. Epub 2020 April 14. DOI: 10.1056/NEJMp2009457.
Salter, W.E.G., 1960. Productivity and Technical Change. Cambridge University Press, Cambridge.
Sculpher, M., Claxton, K., Pearson, S.D., 2017 Feb 1. Developing a value framework: the need to reflect the opportunity costs of funding decisions. Value Health 20 (2), 234–239.
Sendi, P., Gafni, A., Birch, S., 2002. Opportunity costs and uncertainty in the economic evaluation of health care interventions. Jan Health Econ. 11 (1), 23–31.
Sendi, P., Al, M.J., Gafni, A., Birch, S., 2003. Optimizing a portfolio of health care programs in the presence of uncertainty and constrained resources. Dec 1 Soc. Sci. Med. 57 (11), 2207–2215.
Sun, D., Ahn, H., Lievens, T., Zeng, W., 2017. Evaluation of the performance of national health systems in 2004-2011: an analysis of 173 countries. PloS One 12 (3), e0173346.
Thokala, P., Carlson, J.J., Drummond, M., 2020. HTA'd in the USA: a comparison of ICER in the United States with NICE in England and wales. Sep J. Manag. Care Spec. Pharm. 26 (9), 1162–1170.