Limits to measuring the bias of adapted patients: a response to Eyal

Polly Mitchel

King's College London, School of Education, Communication and Society, London, UK
Corresponding author: E-mail: polly.mitchell@kcl.ac.uk

I. INTRODUCTION

Nir Eyal identifies an oversight in much discussion of hedonic adaptation in patients.[1] He convincingly argues that quality of life assessments made by people who have previously had a condition and have been cured of it are crucial in informing evaluations of disease-related quality of life. Eyal proposes that the assessments of cured patients can be used to identify the maladaptive aspects of current patient's assessments. There is a tendency to oversimplify the complex and intractable problem that adaptation presents, and to advocate oversimplified solutions to it.[2] Eyal's proposal is a valuable and disruptive contribution to the literature, and his characterization of the problem helps to expose nuance that is too often overlooked.

In this short response to Eyal's paper, I will endeavor to show that things are more complex still. Although I welcome recognition of the voices of cured patients in decision-making about resource allocation, I resist the idea that their assessments provide a straightforwardly 'better' measure of disease-related quality of life than those of current patients. Cured patients' assessments of quality of life can valuably contribute to resource allocation decision-making in some contexts, but their potential is more limited than Eyal's proposal suggests.

---

1   Nir Eyal, *Measuring Health-State Utility via Cured Patients*, in DISABILITY, HEALTH, LAW, AND BIOETHICS (I. Glenn Cohen, Carmel Shachar, Anita Silvers and Michael Ashley Stein eds., 2020).
2   In the interests of space, I will not spell out the problem of patient adaptation—I direct those unfamiliar with the problem, and indeed those familiar with it, to Eyal's lucid discussion thereof.

First, I argue that there is reason to think that Eyal's proposed 'bias of adapted patients' (BAP) measure has rather more limited application than he suggests. Eyal argues that BAP can be calculated for some groups of patients and then used to identify and offset maladaptation in other relevantly similar groups of patients. However, such projection of BAP values encounters difficulties securing external validity. I argue that this presents particular problems for the projection of BAP values onto groups of patients with incurable diseases.

I go on to question Eyal's claim that quality of life assessments in healthcare decision-making should reflect objective contributors to and detractors from people's well-being. I suggest that we should instead try to find and use assessments that are fit for purpose, recognizing that there is no single right answer to this epistemically complex, socially significant, value laden set of issues. The limited applicability of the BAP measure is far less threatening if we see it as part of our toolkit, rather than an outright solution.

II. PROJECTING THE BAP

Eyal argues that cured patients offer a means of understanding the extent to which patient adaptation is maladaptive, representing delusion, or cognitive bias. Central to his argument is the claim that cured patients are in an epistemically superior position to both current patients and nonpatients. He emphasizes that people who once had a disease or condition but were subsequently cured of it have experience of the health state in question, making them unlikely to exhibit the biases attributed to the general public, who may exaggerate the negative impact of diseases on quality of life.[3] Moreover, cured patients are unlikely to exhibit the biases attributed to patients, who may downplay the negative impact of the disease, perhaps having an investment in maintaining a degree of delusion about their quality of life.[4] This hypothesis that cured patients avoid these biases is borne out by the results of studies which show patients retrospectively evaluating their disease-related quality of life as worse once they have been cured.[5, 6]

Eyal proposes that by dividing average cured patient assessments of quality of life by average patient assessments, we generate a value that represents the maladaptive BAP. BAP values can be calculated for different groups, based, for example, on a shared disease or condition, shared health state or physical functioning, or common demographic factors. BAP is intended to be calculated for some diseases or patient groups and projected onto others: a BAP value that has been calculated for one patient group can be used to modify current patients' assessments of disease-related quality of life in a different patient group, producing a representation of their quality of life without maladaptive bias. Projection is justified when the group of patients onto whom BAP will be projected is 'relevantly similar' to the group of patients for who BAP

---

3   Peter A. Ubel, George Loewenstein and Christopher Jepson, *Whose quality of life? A commentary exploring discrepancies between health state evaluations of patients and the general public*, 12 Qual. Life Res., 599, (2003).

4   Id.

5   E.E.M Adang, G. Kootstra, G.L. Engel, J.P. van Hooff and H.L.G.J. Merckelbach, *Do retrospective and prospective quality of life assessments differ for pancreas-kidney transplant recipients?* 11(1) Transpl. Int., 11 (1998).

6   Dylan M. Smith, Ryan L. Sherriff, Laura Damschroder, George Lowewenstein and Peter A. Ubel, *Misremembering Colostomies? Former Patients Give Lower Utility Ratings Than Do Current Patients*, 25 Health Psych., 688 (2006).

was originally calculated. The suggestion, then, is that the BAP can be assumed to be similar across similar groups of patients. Eyal suggests that such projection can occur not only when BAP has not yet been calculated for a group of patients, but also when it is not possible to calculate BAP for a group of patients. This latter category includes currently incurable diseases, such as HIV and complete spinal cord injury, and also genetic conditions such as achondroplasia and Down Syndrome—groups where there are no former patients.

Eyal argues that the grounds for determining similarity will emerge through the measurement of BAP in various groups of patients. He postulates that as BAP is calculated for different groups of patients, clusters of groups generating similar BAP values will materialize. BAP clusters might comprise groups of patients with a similar degree of functional impairment, length or trajectory of disease, socioeconomic status, demographic characteristics, or other factors. The same BAP value can then be taken to apply to other patients and groups of patients who have these same characteristics.

I will argue that the projection of BAP values from one group of patients to other 'relevantly similar' groups of patients suffers from problems of external validity. External validity is an assessment of the applicability of the conclusions of a study outside the study context.[7] The BAP measure is externally valid if, and to the extent that, the BAP value that is generated by the measurement of quality of life in a specific set of patients can be used to calculate the maladaptive bias of other groups of patients. External validity is essential if BAP values are to be generalized outside of the original study population.[8]

The BAP measure will fail to exhibit external validity if the maladaptive bias in the second, projected population is subject to different causal factors from those to which the first, measured population is subject.[9] Some of the causal factors which may make maladaptation more or less likely are age and patient demographics, features of the disease—whether it is acute or chronic, what the prognosis is—and features of its treatment—whether treatment is risky or safe, whether it is a long or short process. If the factors which cause maladaptive bias in the projected population are different from the factors which cause maladaptive bias in the measured population, then the two populations have different causal structures with respect to maladaptive bias. Calculating BAP in one patient group may, as Eyal suggests, tell us the extent to which current patients exhibit maladaptive preferences. It thereby secures a causal link between the surveyed patient group and the resultant BAP value. However this does not thereby reveal the causal structure which describes this mechanism—although it tells us the 'extent' of maladaptation, it neither reveals the correct description of the causally efficacious factors(s) which lead a given group of patients to generate a corresponding BAP value, nor those differences in the patient group that would have resulted in a different BAP value. This means that, if a projected population differs from

7 Marion K. Slack and JoLaine R. Draugalis, *Establishing the internal and external validity of experimental studies*, 58(22) Am. J. Health Syst. Pharm., 2173 (2001).
8 Allan Steckler and Kenneth R. McLeroy, *The importance of external validity*, 98(1) Am. J. Pub. Health, 9 (2008).
9 Nancy Cartwright, *Evidence, external validity and explanatory relevance, in* Philosophy of Science Matters: The Philosophy of Peter Achinstein (Gregory J. Morgan ed., 2011).

the measured population in an unidentified respect, which is causally significant with respect to maladaptation, projection of the BAP will not be warranted.

How, then, can we be justified in exporting a causal claim from one context to another? Justification may proceed by showing that the two populations share all the factors that are causally relevant to maladaptive bias. But the success of this strategy depends on prior knowledge of the causal factors in the study population, which is the exact information that is lacking. Measuring BAP in many different groups of patients can provide indications of causally relevant factors, but it can never exhaustively identify the causal factors.

By way of illustration, consider two patient groups discussed by Eyal, which exhibit similar BAP values: patients who have had a colostomy (compared with people who have had their colostomy reversed) and patients with acute kidney disease who are on dialysis (compared with people who have had a kidney transplant and no longer need dialysis). That the two groups generate the same BAP values need not imply that they have the same causal structure, nor that the similarity of their BAP values is caused by them sharing some causal factor. Projecting this BAP value onto other patient groups, even if they appear to share characteristics with these groups, would be epistemically reckless if not backed up by evidence which isolates and substantiates a proposed causal factor. Such evidence would need to involve the measurement of BAP in patient groups lacking that factor, showing clearly distinct levels of bias. Apparent similarity of patient groups is not enough. This need not entirely preclude the projection of BAP values beyond the exact conditions of their measurement. But it does mean that this practice should proceed with caution, with recognition of the epistemic limitations of projection and the provisionality of any projected values. Variables such as the tightness of the clustering of BAP values, the type, and extent of the observed differences between different patient groups with similar BAP values will be relevant to assessing the justifiability of projection.

This does not call into question the value of the BAP as a measure of maladaptive bias per se. However, it does throw doubt on Eyal's tentative proposal that BAP values might be projected onto patients in whom BAP cannot be measured because no cure exists for their condition. Projecting a BAP value onto incurable patients involves making the assumption that the incurability of their condition is not a significant causal factor with respect to their maladaptive bias. But the evidence that can be produced to support this assumption is severely limited. Isolating curability as a causal factor would involve measuring BAP in incurable patient groups—but this is ruled out by the very fact of the incurability of their condition. Even if the group is similar to other patient groups for whom BAP can be calculated in 'some' respects—in terms of demographics or functional impairment, for instance—if the incurability of their condition is a significant causal factor with respect to maladaptive bias, they will not be relevantly similar for the purposes of BAP projection. Incurability might increase or decrease the proportion of bias that is maladaptive, and this need not be consistent across different incurable conditions. For example, whether an incurable disease or disability is congenital or acquired might affect its causal implications for maladaptive bias. Incurability might also interact with other causal factors in unknown ways. In the absence of evidence as to whether, and to what extent, curability is a causal factor, using the BAP values of cured patients to calculate maladaptive bias in incurable patients is

not warranted. There may be evidence that patients in the incurable disease group do exhibit maladaptation, but this does not alone justify the assumption that the degree of maladaptation that they exhibit can be determined or accurately estimated using a BAP value calculated in another, necessarily curable, patient group.

This brief discussion of external validity highlights some of the epistemic limitations of the BAP, and cautions that it should be used with prudence. To characterize a patient's self-reported experience as erroneous, and to modify her assessment based on the self-reports of others, is a morally significant act. For it implies that some other person, or group of people, understands what life is like for someone like her better than she does herself. This is not to say that such judgments are not sometimes justified. Eyal argues that projection of BAP values is not an attempt to discriminate against patients, but merely an endeavor to correct mistaken assessments of quality of life. But without robust evidence that such assessments are indeed mistaken, well-meaning attempts to rectify fallible quality of life reports will easily slide toward unjustified denials of self-determination.

III. LOOKING FOR THE RIGHT ANSWERS

Eyal's paper ends with the assertion that what we want from quality of life assessments for healthcare resource allocation is 'primarily to get the answers right'. Cured patients' epistemically superior perspective on disease-related quality of life—so the argument goes—justifies us in using their assessments, even though this involves rejecting current patients' assessments. For Eyal, measures of quality of life that are used in resource allocation decisions should seek to capture 'objective' contributors to and detractors from people's well-being, and to this end, we should use the assessments of cured patients. Cured patients undoubtedly have a distinctive perspective on the health-related disadvantages of the disease or condition with which they were afflicted, and the benefits of treatment and cure. However, I will conclude with a few thoughts which suggest that it would be misguided to treat them as objectively representing disease-related quality of life.

First, insofar as cured patients' assessments of quality of life can sometimes be characterized as the 'right' answers for healthcare decision-making, this will be limited to certain informational and decision-making contexts. Whether cured patients are in an epistemically superior position or not depends on what it is that we want to know. Measures of disease-related quality of life can be used for a wide variety of purposes in relation to healthcare and healthcare resource allocation, and not all of these will benefit from replacing patient assessments with cured patient assessments. Eyal focuses on healthcare resource distribution that involves prioritizing different conditions based on the relative health-related disadvantaged that they confer on patients. The worse a health state is, he suggests, the more funding should be directed toward its prevention, treatment, and cure. The assessments of cured patients are perhaps most useful in contexts which involve broad comparisons between the relative health disadvantages of different conditions, where a single summary value is needed to represent the quality of life associated with each. But in other contexts, the quality of life assessments of current patients will be much more valuable. When, for instance, assessments of quality of life are used to guide clinical decision-making, the specific details of symptoms, functioning and affect that a patient is experiencing will be more instructive than a

single summary score of overall quality.[10] Or, if measures of quality of life are needed to better understand the trajectory of a chronic or terminal disease developing a rich and detailed picture might involve comparing assessments taken from current patients at a series of time points across their disease progression.

Modifying the assessments of current patients using a BAP value would be counterproductive in these more localized contexts. Good decision-making in such contexts is just as much a part of delivering justified resource allocation and appropriate medical care as high-level funding allocation decisions. This suggests that the value of the BAP measure as a tool for modifying the quality of life assessments made by current patients is limited to certain decision-making contexts. And its appropriateness in those contexts is partly determined by our pragmatic ends and priorities.

But moreover, even within specific contexts, there are significant epistemic limitations to the measurement of quality of life which make it difficult to warrant claims that any given assessment captures objective contributors to well-being. Measures of disease-related quality of life paint a relatively limited picture of how life is going for someone with a given condition. Quality of life measures can only ever capture a snapshot of how life is going right now or represent an indication of how things were going during a specified time period via a summary score. Eyal discusses retrospective quality of life assessments made by cured patients, which constitute a single characterization of quality of life for the duration of ill health, perhaps lasting for a period of years. The sense in which they represent quality of life is thus inevitably impressionistic and perspectival. Measures of quality of life capture 'something' about how life is going, at some point in time, but unavoidably fall short of representing the rich and shifting texture of our inner lives. Moreover, different quality of life measures model quality of life differently, some focusing on the satisfaction of human needs, others on psychological factors, others on health and functioning.[11] These differences in part reflect pragmatic limitations on the length and scope of reliable survey tools, but they are also grounded in substantive, normative disagreement about the definition of quality of life. Different measures of quality of life will produce different, often mutually exclusive, assessments of quality of life. Add to this a number of other complicating factors: the quality of life values that are used in healthcare decision-making are usually aggregated across groups of people; furthermore, identifying and alleviating cognitive biases in the measurement of quality of life is, as Eyal notes, extremely challenging. All this amounts to an epistemically complex picture, where the sense in which a single value can be said to represent the quality of life of any individual is partial, provisional and perspectival—and, moreover, contestable.

Where does this leave us? The search for the 'right answer', with respect to the measurement of quality of life may be a futile pursuit. Although the quality of life assessments of cured patients can sometimes help us toward the best solution for resource allocation decisions by mitigating against maladaptive bias, I resist the sug-

10 Irene J. Higginson and Alison J. Carr, *Using quality of life measures in the clinical setting*, 322 BMJ, 1297 (2001).
11 Ann Bowling, Measuring Health: A Review of Subjective Health, Well-Being and Quality of Life Measurement Scales, 4th ed. (2017).

gestion that they enable us to measure objective contributors to disease-related quality of life. Reports of quality of life are necessarily anchored to particular perspectives, and identifying the 'correct' quality of life value to use in decision-making is a normatively and politically laden affair. We should be seeking measures and assessments of quality of life that are context appropriate, fit for purpose, and that make good use of the limited available evidence. The diverse, and sometimes contradictory, assessments of quality of life generated by different people and groups generate evidence which can contribute to epistemically modest and balanced judgments, rather than decisive answers. The assessments of cured patients and the estimates of maladaptive bias delivered by the BAP measure are necessary, but partial and defeasible, elements of this open-ended process.