ies Report

Institute for Employment Studies

Recommendations for an economic evaluation of health coaching

Helen Gray

NHS Leadership Academy

better conversation better health health coaching

ies institute for employment studies

16 April 2019

Report 520

Institute for Employment Studies

IES is an independent, apolitical, international centre of research and consultancy in public employment policy and HR management. It works closely with employers in all sectors, government departments, agencies, professional bodies and associations. IES is a focus of knowledge and practical experience in employment and training policy, the operation of labour markets, and HR planning and development. IES is a not-for-profit organisation.

Acknowledgements

The author is indebted to Dr Penny Newman and Dr Andrew McDowell for their invaluable advice and guidance throughout the study. I would also like to thank Alison Carter and Megan Edwards at the Institute for Employment Studies for all their input and support.

Institute for Employment Studies
City Gate
185 Dyke Road
Brighton BN3 1TL
UK

Telephone: +44 (0)1273 763400
Email: askIES@employment-studies.co.uk
Website: www.employment-studies.co.uk

IES project code: 00938-5255

Contents

Executive Summary

Introduction

Health coaching has been shown to be effective in a number of international studies, but there is limited evidence on its impact within the NHS. As a result, little is known about whether an investment in health coaching offers a good return on investment relative to other types of health intervention. This report proposes an approach to address this evidence gap.

Requirements for a credible economic evaluation

To estimate the costs and benefits of health coaching it is necessary to:

1. Have a clear and consistent definition of health coaching which is adhered to across all sites involved in testing its effectiveness.
2. Have a comprehensive understanding of the range of outcomes that health coaching is likely to effect, including any unintended consequences. This includes documenting outcomes which cannot be easily measured.
3. Be able to distinguish between patients who have received health coaching and those who have not.
4. Have a large enough number of research participants to reach conclusions about the impact of health coaching with a high degree of certainty. Power calculations can indicate whether this is likely to be the case, given the design of the study and the expected impact of health coaching on the primary outcome measures.
5. Take into account the potential for self-selection. If patients and/or staff participating in the study have a choice over whether to participate the impact of health coaching may appear to be higher than if it was extended to less willing participants.
6. Identify suitable outcome metrics. It is easier to identify relevant outcomes when focusing on a single health condition, but this would not indicate whether health coaching is effective across a range of conditions. Generic outcome measures make it possible to estimate the impact of health coaching across diverse conditions.

Data sources

Any study of health coaching is likely to make some use of existing data sources. These could be used to verify that the treatment and comparison groups are similar prior to health coaching and to explore whether the treatment group are representative of the wider population of patients with a particular health condition. Using existing data sources can reduce the burden of data collection and avoid reliance on recall. It may also increase

the proportion of study participants for whom outcomes can be observed compared to response rates for a bespoke longitudinal survey.

The scale of data collection required to carry out an economic evaluation of health coaching depends on the design of the study and the availability of, and access to, existing data sources. There are advantages to designing data collection mechanisms with the explicit intent of estimating the impact of health coaching. However, with limited resources it is likely to be advisable to concentrate the bespoke data collection on information which is not available, or is of poor quality, in existing data sources.

Key outcome measures

Table 1.1 summarises the main outcome measures relevant to a study of the economic impact of health coaching. The inclusion of clinical measures and the precise list of relevant outcomes depend on whether the study focuses on a particular health condition, or on the impact of health coaching across a more diverse selection of patients.

The wide range of measures reflects the fact that the potential costs and benefits of health coaching accrue to many different parties, including patients, employers, NHS staff and the exchequer. An economic evaluation would need to devise a credible way of attaching a monetary value to the impact of health coaching on these outcomes and to document any impacts, or potential impacts, which could not be valued. This would include considering how any omissions might affect the findings of the study.

Table 1.1: Summary of key outcome measures for an economic evaluation of health coaching

Type of outcome	Measure
Clinical	Depends on health condition e.g. for diabetes: • HbA1c (blood glucose level); • GAD-7 (General Anxiety Disorder scale) and PHQ-8 (Patient Health Questionnaire-8) for anxiety and depression; or • 14-item Hospital Anxiety and Depression scale (HADS).
Generic	• PAM (Patient Activation Measure) score • EQ5D, capturing physical and mental health • WEMWBS (7- or 14-item Warwick-Edinburgh Mental Well-being Scale) on wellbeing
Service use	Measures reflecting the use of a range of different health services over a particular timeframe e.g. over the 12-months after starting to receive health coaching, such as: • Number of visits to A&E; • Number of inpatient appointments; • Number of outpatient appointments; • Attendance at appointments; • Medication required.
Staff outcomes	• Job-satisfaction • Employee well-being • Absence rate (percentage of working days lost due to sickness or other absence) • Rate of voluntary labour turnover (percentage of employees who left or resigned voluntarily)
Other outcomes	• Employment rate • Rate of sickness absence • Earnings • Percentage claiming benefits • Amount of benefits/tax credits received • Costs of statutory sick pay • Costs of social care • Costs of medication • Patient satisfaction with health services used Non-experimental outcomes: • Experiences of patients and staff.

Recommendations

A large-scale randomised control trial (RCT) of a single health condition could provide robust evidence on the impact of health coaching in the NHS. However, it would be some time before impacts are apparent. Also, a study of a single condition would not indicate

whether health coaching is effective for different types of health conditions. A quasi-experimental study using existing data could potentially consider the impact of health coaching on a wider range of patients at an earlier point in time. Therefore, the key recommendations to produce credible evidence on the cost effectiveness of health coaching are to undertake the following:

- An RCT to estimate the impact of health coaching on type 2 diabetes. This would:
  - Focus on the primary outcomes of HbA1c, GAD-7 and PHQ-8 or HADS alongside a range of secondary outcomes.
  - Be largely based on data collected in the course of the trial from management information and baseline and follow-up surveys of participants.
  - Explore the feasibility of using of existing data from NHS Digital and other sources to explore the representativeness of trial participants and wider economic impacts.
- A quasi-experimental study to estimate the impact of health coaching across patients with a range of different health conditions within areas which use PAM. This would be based on analysis of existing data, using a similar range of sources to the RCT.
- An economic evaluation of health coaching, drawing on the evidence of impact from the RCT and quasi-experimental studies.

1 Purpose of the report

Much of the existing evidence on the cost-effectiveness of health coaching comes from studies undertaken overseas. For example, a large-scale Swedish study explored whether health coaching could be used to reduce the demand for emergency treatment. This randomised control trial of over 12,000 patients demonstrated that health coaching reduced the rate of hospitalisation by 12 percent (Edgren et al. 2016). A US study of patients with diabetes, hypertension, or hyperlipidemia found that health coaching improved control of hemoglobin A1c and LDL levels, but not blood pressure, compared with usual care (Willard-Grace et al. 2014). Again, this evidence was based on a randomised control trial of over 400 patients.

Whilst there have also been multiple studies which have suggested that health coaching is also likely to have positive effects in the UK context, these have tended to be small in scale, or have lacked a robust comparison group. One such example is a study of health coaching on an elderly care rehabilitation ward at the Royal Hampshire County Hospital, Winchester, which was based on analysis of data supplied by 46 participants. The small sample size meant that whilst the findings suggested that health coaching improved patient outcomes, the observed changes were not statistically significant (Kibble et al. 2016).

The small scale of many past UK studies of health coaching, and/or the methods used to assess impact means that there is a lack of conclusive evidence on the impact of health coaching within the NHS. This evidence is vital to inform future funding decisions as it is important to know whether extending the use of health coaching is likely to result in savings for the NHS, or for the exchequer or society. There is also a need to consider how the return on an investment in health coaching might compare to the returns from a similar investment in other services, so that resources can be deployed to maximum effect.

Having a robust and defensible estimate of impact is an essential pre-requisite for an economic evaluation. Therefore, the main objective of the current feasibility study is to outline possible approaches to estimating the impact of health coaching. The focus is on approaches which use a counterfactual i.e. they seek to estimate the outcomes that those who receive a health coaching intervention would have achieved if they had not been subject to this intervention and then compare these counterfactual outcomes with the outcomes that they actually attained. This gives an estimate of the impact of health coaching.

Discussions so far have suggested that there are two potential approaches to estimating the impact of health coaching which are potentially viable:

7. An experimental approach (a randomised control trial), including bespoke data collection.

8. Quasi-experimental analysis based on existing data;

The requirement to engage in data collection for a randomised control trial (RCT) would increase costs but has the potential to explore a wider range of outcomes and to better control for differences between treatment and comparison groups. There are likely to be advantages to making use of secondary data sources alongside bespoke data collection, even in the case of an RCT.

The paper will reviews the relative advantages and disadvantages of the two possible approaches and set out a proposed design to estimate the impact of health coaching using each method. It will provide a specification which can be used to commission research on the impact of health coaching, or to make an application for funding from the National Institute of Health Research (NIHR) for an experimental study.

2 Challenges to carrying out an economic evaluation of health coaching

2.1 Introduction

This section set out the main challenges to carrying out an economic evaluation of health coaching, drawing on the previous literature, as well as findings from the feasibility study. It sets out the basic requirements to carry out an economic evaluation and considers how these apply in the case of health coaching in particular. It also considers how health coaching presents particular challenges when seeking to carry out an economic evaluation.

2.2 Defining health coaching

Wolever et al. (2013) noted that there have been differences in the way in which health coaching is defined in some previous studies. A clear and consistent definition is needed to ensure that interventions which take place on different sites or which are applied to different types of health conditions are sufficiently similar to reach conclusions on the overall effectiveness of health coaching. Wolever et al. (2013) propose a definition of health coaching which draws on the common features of previous studies, namely:

> *'a patient-centred process that is based upon behaviour change theory and is delivered by health professionals with diverse backgrounds. The actual coaching process entails goal-setting determined by the patient, encourages self-discovery in addition to content education, and incorporates mechanisms for developing accountability in health behaviours.'*
>
> Wolever et al. (2013)

For the current study to produce convincing evidence on the effectiveness of health coaching it is important that it adheres to the established definition. RCTs of service delivery often consider how closely service delivery matches what is intended at the outset (known as fidelity).

2.3 Mapping the full range of outcomes health coaching is likely to effect

An economic evaluation of health coaching is only likely to be robust and defensible if it takes into account the full range of outcomes that health coaching is designed to effect, as well as any effects which are unintended or unforeseen at the outset (both positive and negative). Health coaching might potentially affect a wide range of outcomes, both for individual patients, staff delivering the coaching, the NHS, society more generally and the

exchequer. The breadth of the potential impacts from health coaching means that it may be difficult to obtain reliable information on all of the outcomes of interest. With any estimate of impact there is a risk of focusing on things that can be easily measured, rather than those which are important, but hard to capture. If important impacts are omitted from the economic evaluation, the benefits (or costs) of health coaching may be under- or over-stated. It may not be possible to estimate the impact of health coaching on all potential outcomes, but an economic evaluation should note all of the likely direct and indirect outcomes, even if some cannot be measured. This provides an insight into whether costs and benefits are likely to be under- or overstated.

Previous studies of health coaching in the US and elsewhere, as well as existing UK case studies, provide some insights into potential outcomes from health coaching when the approach is applied to a range of different health conditions. This existing evidence can therefore be used to scope the range of outcomes which should be considered in an economic evaluation of health coaching. Chapter 4 provides further details on each of these measures.

In the case of an RCT, it is necessary to develop a protocol which states the primary outcomes which are expected to result from the intervention, how these outcomes will be measured and the size of the expected impact. This protocol must be fixed before the trial commences. Without a protocol there is a risk that outcome measures are adjusted during the course of the trial to focus on those which show a stronger positive impact from the intervention. Whilst it is necessary to identify the primary outcome measures prior to the start of the trial, it is nevertheless possible to estimate impacts on a wider range of outcomes so that these impacts can also be taken into account when estimating the costs and benefits of health coaching.

2.4 Distinguishing between those who have and have not received health coaching

To estimate the impact of health coaching it is necessary to be able to distinguish between those who have received health coaching and those who have not. It is necessary to know how consistently health professionals use the approach and to also have access to information on a comparison group of similar patients who have not experienced health coaching. For example, even when health coaching is used for a particular group of patients, there may be some health care professionals who adhere less strictly to the principles of health coaching, or do not use the approach with some patients. If some patients are assumed to have received health coaching when in practice it was not used, or was not used consistently, the impact of health coaching may be understated.[1] Likewise, if some patients in the comparison group did in fact experience health coaching, the impact of health coaching may be understated. In the most extreme cases it may be impossible to observe whether a patient has received the health coaching approach or not.

[1] This is sometimes thought of as the patient experiencing a lower than intended 'dose' of the intervention.

2.5 Adequate sample sizes

If few individuals receive health coaching there is a risk that, even if it is effective, it may not be possible to conclude with certainty that it has had the expected effect. This particular problem has affected the ability to obtain conclusive evidence on the effectiveness of health coaching in previous studies e.g. using health coaching in elderly care rehabilitation services (Kibble et al. 2014). Whilst it may be possible to improve the ability to detect the impact of an intervention in circumstances where few individuals are treated by increasing the size of the comparison group, this is not always possible.

Given the importance of having an adequate sample of both treated and untreated individuals, it can be helpful to estimate the size of the sample required to detect a statistically significant impact from the intervention on the outcomes of interest before deciding whether to evaluate impact. Power calculations are used to estimate appropriate sample sizes, but a number of factors determine the usefulness of such calculations:

- Whether well-informed estimates of the counterfactual and the expected impact of the treatment are available. If little is known about expected impacts and the counterfactual for the group who will be targeted for health coaching, a power calculation could result in the wrong conclusions being drawn about required sample sizes.
- Power calculations typically assume that eligible individuals will be assigned to the treatment at random. They are less likely to provide an accurate indication of required sample sizes in circumstances where this is not the case.
- In some instances it is apparent from the outset that sample sizes are unlikely to be sufficient to be able to detect a statistically significant impact of the expected size. For example, this might be the case if capacity constraints mean that an intervention can only be offered to a small number of individuals. In these circumstances, power calculations would be of limited value.

2.6 Potential self-selection

There are a number of ways in which self-selection might affect the estimated effectiveness of health coaching. Patients may vary in the extent to which they respond to a health coaching approach. This might be due to the way in which the approach is implemented by those delivering the coaching, or the individual patients themselves. It is also possible that staff who lack the skills to be effective health coaches choose not to do the training, so that pilot studies are based on delivery by the most able health coaches.[2]

The fact only a subset of patients may choose to actively engage in health coaching, or only a subset of staff may be capable of delivering it, means that it may be difficult to reach general conclusions about the effectiveness of health coaching if it was rolled out

[2] This risk can be reduced by focusing on the impact of health coaching for those offered the approach, rather than those choosing to participate.

more widely. For example, findings based on outcomes for engaged patients working with staff committed to delivery may overstate the impact of health coaching if the approach was extended to a wider range of (potentially less engaged) patients and staff.

One of the advantages of an RCT is that because patients are randomly allocated either to health coaching (**the treatment group**), or to a **control group**, the average outcomes experienced by the control group provide an estimate of the outcomes that those who are subject to the health coaching would have attained if they had not received this type of support. This is known as **the counterfactual**. Provided the number of patients allocated to each group is sufficiently large, they should be well-matched on all characteristics likely to determine outcomes and so deducting outcomes for the control group from those for the treatment group should provide a credible estimate of impact. However, whilst there would be no reason to expect that those assigned to the control group would be less likely to achieve any beneficial outcomes from health coaching than those assigned to the treatment group, there is still a risk that patients and staff who choose to participate in the trial have a stronger level of interest in health coaching and are more likely to benefit from it than those who choose not to participate. It is therefore important to consider the potential implications of this possibility when interpreting the findings of the analysis.

With a quasi-experimental design, outcomes for individuals who are subject to the health-coaching approach are compared to those for a comparison group of individuals who are not subject to health coaching. For this to provide a robust estimate of the counterfactual, it is important to have access to detailed information on the full range of characteristics which determine whether a patient receives health coaching, as well as any characteristics which have a bearing on the outcomes that they experience. If it is difficult to observe the factors which mean that a patient is more likely to engage with health coaching, there is a greater risk that the treatment and comparison groups are poorly matched and outcomes for the comparison group are less likely to provide a robust estimate of the counterfactual. For example, it is possible that staff who provide health services to the comparison group are less likely to achieve the same outcomes as those who engage in health coaching irrespective of the approach used.

2.7 Finding suitable outcome metrics

The fact that health coaching appears suited to treating a wide range of health conditions suggests that potentially the numbers of individuals who could be offered a health coaching approach could be substantial. In theory, this reduces the risk that any impact from health coaching would not be detectable, due to small sample sizes. However, in practice, this depends on having common outcome metrics which are suited to capturing the impact of health coaching across a range of different conditions. As the approach can be applied to individuals with very different types of health problems, including diabetes, heart disease, cancer patients and those recovering from strokes, the outcome measures which are most likely to capture the impacts of health coaching on each condition are likely to differ.

The fact that the most suitable outcome measures are likely to differ depending on the health condition under examination means that it is likely to be necessary to focus on a single condition which affects a large number of people to ensure that it is possible to detect any impact from health coaching on a relevant outcome. However, it may be difficult to reach conclusions on the overall effectiveness of health coaching from a study of a single condition, or a series of studies on different conditions with different outcomes measures. In particular, it may be difficult to establish whether health coaching is more or less effective for particular types of conditions, or equally effective for all.

3 Data

3.1 Using existing data

One of the main advantages of making use of existing data sources in a study of health coaching is that it can be used to verify that the treatment and comparison groups are similar in the period before they receive health coaching. Where historical information is available from existing data sources, this can be used to reduce reliance on recall. When using quasi-experimental methods, it is particularly important to have access to detailed information on the treatment and comparison group before health coaching commences to ensure that outcomes for the comparison group are likely to provide a reliable estimate of the counterfactual.

In a randomised control trial, there is less risk that the treatment and comparison groups are poorly matched prior to any health coaching, as individuals are not able to choose whether they receive health coaching or not. However, access to a wider database of individuals, in addition to those who put themselves forward for a trial does provide an insight into whether participants are likely to be representative of the wider population of individuals with a particular health condition. As mentioned in Section 2.6 even with an RCT it is possible that trial participants have more favourable views of health coaching than those with the same health problem who choose not to participate. If attitudes towards health coaching determine its effectiveness, the findings from an RCT based on an engaged subset of patients may be very different to the likely impact of health coaching on the wider population of patients with that type of health problem. Having detailed background information on all those with a particular type of health condition, as well as on those who participate in the trial can help to identify any potential selection issues and can aid the interpretation of findings from the trial.

A further advantage of making use of existing data sources in a study of health coaching is that it reduces the reliance on treatment and comparison groups cooperating in data collection. It can be particularly difficult to collect information on outcomes from the comparison group, given that they will not have experienced health coaching and are therefore unlikely to experience any personal benefit from participation in a trial or other study. Even for those who receive health coaching it may be difficult to maintain contact for sufficient time to observe the full range of expected outcomes. Low response rates can threaten the robustness of the findings and the practical difficulties of maintaining contact with individuals and collecting information on a wide range of characteristics and outcomes mean that primary data collection can be expensive. It is therefore generally advisable to make use of existing data sources wherever possible, even when some bespoke data collection is necessary.

A final advantage to making use of existing data sources is that it may be feasible to produce quantitative evidence on the effectiveness of health coaching at an earlier point in time than would be possible with a study reliant on bespoke data collection. Depending on the period of time over which outcomes from health coaching are likely to be experienced, it could take months or years to observe outcomes from an RCT. If it is possible to identify individuals who have experienced health coaching and the outcomes that they have experienced in existing data sources, a quasi-experimental design based on analysis of secondary data might provide earlier insights into the effectiveness of health coaching.

3.2 Bespoke data collection

The main advantage of bespoke data collection is that the data collection tool can be designed to ensure that it provides detailed information on the characteristics which are likely to determine whether someone actively participates in health coaching, aspects of the delivery of health coaching which shape its effectiveness, and the full range of benefits that can result. Bespoke data collection may also be required to estimate the costs of service delivery, or the costs that participants might incur, over and above the costs that they would experience from participating in alternative types of treatment.

Even when bespoke data collection is possible, as the previous section suggested, there may be practical constraints on the amount of time that participants are prepared to spend answering questions, or the reliability of the information that they are likely to be able to provide. This means that even in the case of an RCT, where participants are aware that they are participating in a trial, it may be beneficial to make some use of existing data sources to reduce respondent burden, reduce problems with non-response or improve the accuracy of the information supplied. A pragmatic approach is often to develop bespoke methods of data collection to capture information which is hard to observe in existing data sources, or where existing sources only provide poor proxies. The bespoke items can then focus on collecting vital information which is not available from other sources.

4 Key outcome measures

4.1 Introduction

The costs and benefits of health coaching are potentially widespread and so it is necessary to decide on the scope of the study before relevant outcome measures can be identified. For example, the study could focus on estimating the impact of health coaching for eligible patients, or it could look at the impact on the NHS more generally, including outcomes related to staffing and service usage.

If the focus of the economic evaluation is on the costs and benefits of health coaching to the exchequer, it is necessary to have a detailed understanding of the full range of likely outcomes, including any unintended consequences and positive and negative spillovers for other patients or groups of staff. Being able to produce an economic evaluation depends on having already mapped the full range of possible outcomes through qualitative work. If important outcomes are not considered or cannot be measured, this will reduce the credibility of any economic evaluation.

The following sections outline the key outcome measures that would need to be considered in an economic evaluation of health coaching, including a range of clinical and non-clinical measures and outcomes for staff and the exchequer. Whilst estimating the impact of health coaching on all of these measures may be desirable in terms of capturing its full economic impact, it is likely to be costly. It may be difficult to obtain the required funding until there is clear evidence of impact. In the first instance, it may be prudent to focus on devising a robust and defensible estimate of impact on the key outcome measures for a particular patient group. Once this has been obtained, it may be feasible to obtain funding to explore wider economic impacts.

Finally, it is worth noting that whilst economic evaluations seek to map the range of potential outcomes from an intervention, it can still be possible to produce a credible economic evaluation when information is lacking on some aspects of costs or benefits. Known gaps in information should be noted, and the potential implications for the findings considered, but an economic evaluation could still be informative as long as there is good reason to believe that any omissions would have limited impact on the main conclusions of the study. Related to this point, it is often necessary to make some assumptions about the value of particular benefits or costs during the course of an economic evaluation and to explore the implications of varying these assumptions. This provides an insight into how any areas of potential uncertainty might affect the estimated ratio of costs and benefits.

4.2 Key outcomes from health coaching

The following subsections set out the key outcomes that health coaching is likely to affect and which would need to be considered within a full economic evaluation of health coaching. These are summarised in Table 4.1 at the end of this section.

4.2.1 Clinical outcomes

Health coaching is designed to have a direct effect on the physical and/or mental health of patients. For those with physical health conditions, any impact from health coaching might be observed as an improvement in an objective measure which captures changes in a particular health condition. For example HbA1c is used to measure blood glucose levels for patients with diabetes. A reduction in HbA1c is associated with a reduction in the risk of complications from diabetes. As well as the obvious benefit to the patient from improved health, a reduction in the risk of complications would in turn reduce the likely future costs to the exchequer of funding treatment for the patient. It is therefore possible to attach a value to any benefit that health coaching might have on the clinical outcomes experienced by patients with a given health condition, both for the patients themselves, and for society.

Similarly, there are standard clinical measures which are designed to capture changes in mental health conditions and which have been validated. These include the seven-item General Anxiety Disorder scale (GAD-7), and the eight-item Patient Health Questionnaire-8 (PHQ-8) which cover anxiety and depression, and the 14-item Hospital Anxiety and Depression scale (HADS). These scales are widely used and again, improvements in scores on these scales can be linked to benefits both to the patient and in terms of the likelihood that they will require further treatment. [3]

4.2.2 Generic outcomes

The value of considering generic outcome measures in a study of health coaching is that they are relevant to a range of different conditions. They can therefore be used to assess the effectiveness of health coaching for different types of health problems. In some cases there also established ways of attaching value to a given change in a scale. Set against this, the most widely-used measures are proprietary and it is necessary to pay to administer the scale.

The generic measure which is most likely to be useful to an economic evaluation of health coaching is the Patient Activation Measure (PAM) score. This is a 100-point scale (derived from a series of 13 questions) which captures the extent to which an individual feels able to manage their health, taking into account their knowledge, skills and

[3] Sforzo et al. (2017) identify nine outcomes which have been commonly considered in past studies of health and wellbeing coaching. These are Body Weight (or BMI), Systolic blood pressure (SBP), Low-density lipoproteins (LDL), Hemoglobin A1C, Health risk appraisal (HRA), Pain, Psychological factors, Exercise behaviors and Nutrition behaviors. The relevance of these outcomes to a study of health coaching obviously depends on the specific health condition that the coaching is being used to address.

confidence. It is relevant to a wide range of different types of health conditions. Patients can be divided into four groups based on their level of activation.

Changes in PAM can be linked to changes in the use of health services, making it possible to attach a value to any impact that health coaching has on PAM. For example, a study of patient activation amongst 33,163 patients in Minnesota found that the predicted costs of health care for patients with low activation scores were higher than for patients with high activation levels (Hibbard et al. 2013). A study of over 2,000 high-cost patients in the US found that moving from one level of PAM to the next level was associated with a reduction in ongoing health-care costs of 8.3 percent (Lindsay et al. 2018). PAM is licenced by Insignia Health, so there is a charge for its use. However, it is already used by some NHS sites.

Another licenced scale which is widely used in health studies is EQ5D. This is an indicator of health across five physical and mental dimensions, including mobility, self-care, usual activities, pain/discomfort and anxiety/depression. The scores on each of the five dimensions can be combined into a 5-digit score recording health status on each dimension. This can in turn be converted to a value on a single index which can be used to compare the health of individuals with different types of health problems. As with PAM, it is necessary to pay to use the electronic version of the scale in larger-scale health studies. It is licenced by EuroQol.

In addition to the generic measures which are directly health-related, there may be value in considering a measure such as wellbeing. Whilst the link with financial benefits would be indirect, wellbeing is an important outcome which health coaching could be expected to effect and changes in this might be an early predictor of the impact of health coaching on other outcomes which could be more easily valued, such as entering or re-entering work. It would therefore be useful to consider including measures of wellbeing in a more comprehensive study of the impact of health coaching. The Warwick-Edinburgh Mental Well-being Scale is used to capture wellbeing in health studies. There are short (7-item) and long (14-item) versions. It is necessary to register to use the scale and the standard licence only allows use by not-for-profit organisations.

4.2.3 Service use

If health coaching has a positive impact on patients, it is likely that it would reduce demand for NHS provision, at least over the longer term. The impact on this particular type of outcome over a shorter period of time may be more mixed, for example, if the health coaching means that patients are more engaged with their treatment and consequently make greater use of the services which are available to them.

Irrespective of the expected impact of health coaching on the demand for health services over different periods, it seems probable that any impact that health coaching has on the use of NHS services would account for a large proportion of the potential costs and benefits resulting from the approach. Outcome measures could include the number of visits to A&E and the number of inpatient and outpatient appointments, which could potentially be observed from the Hospital Episode Statistics datasets, available from NHS Digital. They might also reflect attendance at appointments and the need for medication.

4.2.4 Staff outcomes

As well as the obvious potential impact from health coaching on patients, health coaching might be expected to have an impact on a range of outcomes for staff involved in the delivery of health services. The outcomes affected could potentially include job satisfaction, wellbeing, absence and voluntary labour turnover, with obvious costs and benefits associated with increases or decreases in each of these outcome measures. The staff affected could include both those directly involved in delivering health coaching, but also other NHS staff who might indirectly benefit from a reduced workload or improved job satisfaction from working with more engaged patients.

4.2.5 Other outcomes

There are a range of other outcomes for patients, employers and the exchequer which may result from health coaching. For example, if health coaching is effective, it may increase the likelihood that patients are able to sustain, or enter work. It may reduce benefit and /or tax credit payments and increase tax receipts for the exchequer. Demand for social care may also be reduced, with resulting savings for local authorities. There may also be savings if the need for medication and statutory sick pay is reduced.

If the study involved bespoke data collection, it may also be useful to explore the impact of health coaching on patient views of the health services that they have used, for both the treatment and control groups. This might provide an insight into whether patient satisfaction varied depending on whether they were offered health coaching. Whilst this would not be essential for an economic evaluation, evidence on whether health coaching was well-received by patients may inform future decisions about increasing the use of health coaching within the NHS.

As well as collecting information on patient satisfaction with NHS services, both for those who experience health coaching and a control group of those who do not, a wider study of health coaching might also explore in detail the experiences of patients who receive health coaching and staff involved in delivery. As these measures would only be relevant to the treatment group and could not therefore be used to estimate the impact of health coaching, they are known as non-experimental outcomes. Qualitative research with staff and patients, or a larger-scale survey, could be used to explore the experiences of patients and staff and might provide an insight into ways in which health coaching could be improved for either group. Whilst past case studies of health coaching have addressed this to an extent, there may be value in producing evidence which can be directly linked to an experimental study.

Table 4.1 Summary of key outcome measures for an economic evaluation of health coaching

Type of outcome	Measure
Clinical	Depends on health condition e.g. for diabetes:

	• HbA1c (blood glucose level); • GAD-7 (General Anxiety Disorder scale) and PHQ-8 (Patient Health Questionnaire-8) for anxiety and depression; or • 14-item Hospital Anxiety and Depression scale (HADS).
Generic	• PAM (Patient Activation Measure) score • EQ5D, capturing physical and mental health • WEMWBS (7- or 14-item Warwick-Edinburgh Mental Well-being Scale) on wellbeing
Service use	Measures reflecting the use of a range of different health services over a particular timeframe e.g. over the 12-months after starting to receive health coaching, such as: • Number of visits to A&E; • Number of inpatient appointments; • Number of outpatient appointments; • Attendance at appointments; • Medication required.
Staff outcomes	• Job-satisfaction • Employee well-being • Absence rate (percentage of working days lost due to sickness or other absence) • Rate of voluntary labour turnover (percentage of employees who left or resigned voluntarily)
Other outcomes	• Employment rate • Rate of sickness absence • Earnings • Percentage claiming benefits • Amount of benefits/tax credits received • Costs of statutory sick pay • Costs of social care • Costs of medication • Patient satisfaction with health services used Non-experimental outcomes: • Experiences of patients and staff.

4.3 Scope of the outcome measures

The previous section set out the key outcomes that health coaching might be expected to effect, with a particular focus on those which are likely to have financial implications. In an economic evaluation, it is important to identify all the possible sources of costs and benefits, even if it is not feasible to estimate the impact of the intervention on every one, or to attach a monetary value to all possible impacts. An economic evaluation would need to build on the list of potential outcomes set out above. It would also need to identify credible ways of attaching a monetary value to any estimated impact on these outcomes,

for example, by drawing on previous studies which have valued the impact of a given change in PAM scores.

It is not necessary to be able to observe the impact of health coaching on every possible outcome, but it would be important to seek to estimate its impact on the outcomes which are expected to result in the most substantial costs and benefits. It is also necessary to ensure that any assumptions made about the value of costs and benefits are defensible and to document any potential impacts which have not been taken into account, or cannot be valued, including how they might affect the conclusions of the study.

4.4 Generic versus specific outcome measures

The main advantage that generic outcome measures, such as the PAM score or EQ5D, have over more specific measures of mental or physical health is that they are suited to capturing changes in different types of health conditions. As they are relevant to different patient groups, they can be used to consider the impact of health coaching across a range of health conditions. This means that a study of health coaching using these measures could potentially explore how effects vary depending on the nature of the health condition. It may also be possible to estimate impacts for a larger sample of patients, increasing the likelihood that any findings are statistically significant.

Set against this, whilst changes in generic health measures are likely to capture improvements in health, there may be scepticism that the value of a given change is the same regardless of the nature of the health condition. Aside from this, it is possible that the effectiveness of health coaching varies between patients with different health conditions or other characteristics. If this is the case, the composition of the group of patients who participate in health coaching will determine average effect sizes. An economic evaluation based on generic health measures would need to consider how varying the sample of participating patients would affect the findings.

Using measures which are specific to a particular health condition would make it necessary to focus on a more narrowly defined patient group. However, having a narrower focus in the study could improve the robustness of the findings and make it possible to generalise about the impact of health coaching on the wider population of patients with a particular type of health condition.

Given the relative advantages and disadvantages of focusing on either generic or specific outcome measures, it may be beneficial to carry out two separate studies, with one estimating the impact of health coaching on a particular health condition and the other looking at a wider range of patients. The evidence from the latter may be less robust, but it may be possible to carry out such a study with existing data.

5 Proposed approach to estimating impact

This section outlines the two proposed approaches to estimating impact – namely an RCT and a quasi-experimental design. It begins by setting out the essential elements common to either approach. It then outlines a proposed design for a study using each method.

5.1 Basic requirements

This section briefly explains requirements to address the challenges to carrying out an economic evaluation of health coaching, previously set out in Chapter 2. These basic requirements apply regardless of the methods used to estimate impact.

5.1.1 Having a clear definition of health coaching

Regardless of whether the approach to estimating impact uses experimental or quasi-experimental methods it will be important to ensure that the underlying concept to be tested reflects the key features of health coaching, for example, based on the definition proposed by Wolever et al. (2013).

5.1.2 Eligibility criteria

As well as having clarity and consistency in the intervention being tested, it is important to have clear and observable eligibility criteria. If there is an element of discretion in whether to offer an individual patient health coaching, it can be difficult to identify a well-matched comparison group. For example, with limited resources health workers may decide to only offer coaching to patients who they believe are most likely to benefit from the approach. It would be difficult to identify a group of comparators who were equally likely to benefit without information on the factors which determined the decision of the health worker.

5.1.3 A clear statement of intended effects

When estimating impacts on a large number of outcomes, there is a risk of 'false positives', where some outcomes appear to be affected by health coaching when in fact the association arises purely by chance. Clearly identifying the intended effects at the outset reduces the risk of accusations that the primary outcome measures have been selected at a later stage because they demonstrate the desired effect. In the case of an RCT, the key outcome measures would be stated in a trial protocol which would ideally be published before participants were recruited to the trial.

Having a clear idea of the primary effects expected from health coaching at the outset is also necessary to conduct power calculations to ascertain whether it is likely to be possible to detect any effects of the expected order of magnitude with the expected

sample size. In the case of an RCT, power calculations can be used to decide the number of trial participants required to detect an effect of the expected size over the expected counterfactual level of outcomes.

5.1.4 Identifying factors likely to determine costs and benefits

To carry out an economic evaluation of health coaching, it is necessary to map all the factors which might have a bearing on costs, as well as the full range of potential benefits. This should include both monetary and non-monetary costs and benefits, irrespective of whether they can all be observed.

5.1.5 Identifying credible ways of valuing costs and benefits

Having mapped the potential range of costs and benefits, any economic evaluation must consider how a monetary value of each should be calculated. It is important that this is transparent and credible and that it considers how the findings of the analysis vary if assumptions about these expected values differ.

5.1.6 A means of distinguishing between the treatment and comparison groups

Regardless of the methods used to estimate the impact of health coaching it is vital to know which patients have been exposed to health coaching and which have not. If the comparison group included individuals who had received health coaching (**known as contamination**), the estimated impact of health coaching may be understated. Having more detailed information on participation in health coaching is likely to be useful in exploring whether impacts vary between patients depending on whether or not they choose to actively engage with health coaching, as well as variations in the amount of support and the consistency of approach that is available to them.

5.1.7 Information on characteristics and outcomes of treatment and comparison groups

For an impact evaluation of health coaching to be credible, it is necessary to observe the full range of factors which are likely to determine both whether a patient participates in health coaching and the outcomes that they experience. This is because if there are systematic differences between the treatment and comparison groups, outcomes for the comparison group are likely to provide a poor estimate of the counterfactual. This can result in the impact estimates being biased and giving a poor indication of the likely impact of health coaching if it was extended more widely.

5.1.8 Sufficient information for the study to be replicated

It is good practice to ensure that sufficient information is provided on the nature of the intervention and the design and conduct of the evaluation to allow the study to be replicated.

5.2 Experimental design

This section provides an outline for an experimental study to test the impact of health coaching. The total cost to the NHS of treating diabetes each year is £10billion and it affects 3.9 million people (Diabetes UK, 2017: 4). Type 2 diabetes is a growing problem which is costly to the NHS given the requirements for ongoing treatment to deal with complications. The link to patient decisions on diet and exercise mean that the health coaching approach potentially offers a cost-effective way to reduce future expenditure on health care. There is also an established clinical measure (HbA1c) which can be used to monitor changes in the condition over time and which is known to correlate with the likelihood of needing treatment for complications from diabetes.

5.2.1 Eligibility criteria

A trial of health coaching for those with diabetes could explore the impact of health coaching across all patients with type II diabetes, or for those at a particular stage. For example, the impact of health coaching might vary depending on whether the patient had recently been diagnosed, or whether the patient had already started to experience complications. If funding were available for a large-scale study, it might be possible to look at the impact of health coaching for different subgroups of patients, taking into account the severity of the condition at the time of randomisation, but with more limited resources it might be preferable to consider the impact of coaching for those who had recently been diagnosed with the disease. The feasibility of focusing solely on the recently diagnosed would depend on the rate of diagnosis and whether this was likely to result in a flow of patients onto the trial which was sufficient to detect the expected impact.

5.2.2 Randomisation

To test the impact of health coaching using a RCT it is necessary to have a means of assigning those with type II diabetes to either the treatment or control group. Type 2 diabetes is likely to be diagnosed by the GP following urine and blood tests. The GP then discusses treatment options and the patient is invited to attend check-ups every three months initially. The check-ups are with either the GP or a diabetes nurse. The HbA1C test is repeated at each appointment. Once the condition is stable, the check-ups reduce in frequency to once every six months. There is a more extensive check-up once a year, to check the feet, eyes, blood pressure, cholesterol and kidneys.

As the diabetes treatment and check-ups are delivered by GPs and diabetes nurses, the most feasible option for randomisation is to randomly assign some GP practices to be trained in health coaching. Others would be assigned to a control group and would not receive the training. With this type of clustered design it would be important to have detailed information on the practices and the characteristics of the patients that they serve to ensure that the treatment and control groups are well-matched. It would also be important to have a database of trial participants to ensure that any randomised patients who move to another practice can be excluded from the analysis.

Depending on the future plans for the roll-out of health coaching, there may be some value in using a stepped wedge approach. One of the difficulties with an RCT is that there may be ethical issues in denying patients in the control group access to health coaching until the end of the trial. With a stepped wedge approach, health coaching would be extended to a larger number of clusters over the lifetime of the trial until all participating GP practices are using the health coaching approach. This can make it possible to produce robust evidence on the impact of a health intervention in circumstances where there are plans to roll-out the approach across all practices within an area within the lifetime of the trial. However, the fact that practices would start using the health coaching approach at different points in time with a stepped wedge approach complicates the process of calculating the sample size required to detect any impact from health coaching (Hemming et al. 2015). Careful analysis would be required to ensure that the trial was designed in a way that maximised the likelihood of detecting any impact from health coaching.

At the initial appointment the GP would need to inform the patient about the trial. They would be asked to give the patient an information sheet and seek their agreement to participate in the research, including allowing researchers access to their HbA1C results and contact details. With limited time for appointments, it would not be feasible to ask the GP or nurse to collect baseline data beyond the initial HbA1C level. The GP or nurse could be asked to record this at subsequent appointments for trial participants. One potential problem with collecting the HbA1C through the regular appointment with the GP or nurse is that it relies on patients attending check-ups to the expected timetable. If they do not attend, or attend less frequently that intended, there will be variation between patients in when the HbA1C level is observed.

Practices which sign up to RCTs receive payments for participating in research and so there are some existing incentives for practices to take part in the study. Nevertheless, those assigned to the treatment group would need to spend time attending training in health coaching techniques, whilst those assigned to the control group would not experience any benefit from trial participation, other than the payment for taking part in research. It would therefore be important to consider ways of encouraging practices to participate in the trial to ensure numbers are adequate to identify any impact from health coaching. The clustering at practice level increases sample size requirements due to the fact that there are likely to be similarities in patients visiting particular practices. For example, it is known that diabetes is more common in certain parts of the country.

5.2.3 Data collection arrangements

To reduce the burden on GPs or nurses the bulk of data collection would need to be carried out via a survey administered by an external fieldwork company. Once the patient has agreed to participate in the trial, the fieldwork company could contact the patient to collect baseline information, perhaps via a telephone survey. The baseline survey would include information on the characteristics and circumstances of trial participants at the time of randomisation, including whether they had other health problems and were on any medication. It could also collect information on work and health histories. This information would be used to check the treatment and control groups were well matched on

characteristics likely to determine the outcomes they might experience from health coaching.

In addition to collecting baseline information on trial participants and the date of starting on the trial, it would be useful to put in place systems to collect information on the health coaching received. To reduce the burden on GPs and nurses this should be a very simple dataset, perhaps extracted from other sources, and would cover the number and timing of appointments. Ideally, it would also record participation in particular types of activities, including those specific to health-coaching, and details of any medication prescribed. As previously mentioned, the HbA1C level would need to be recorded at each appointment and this entered into the management information database.

As well as data collected from patients during the course of health-coaching, an RCT would ideally involve longitudinal data collection at specific points in time after the patient started to receive health coaching. At a minimum, this would include a census of trial participants at a point in time when the impact of health coaching is expected to be apparent. For the purposes of delivering an RCT, the health coaching approach would be time-limited (this is usually necessary anyway due to resource constraints). In the case of health-coaching, it would be advisable to observe outcomes a fixed number of months after randomisation, and possibly also sometime later, to explore whether effects evolve over time. It is important that outcomes are observed at a similar point in time for both treatment and control groups. The advantage of observing outcomes at an early point is that contact details supplied at baseline (by both the treatment and comparison groups) can become out-of-date, making it difficult to maintain response rates over time. The survey could include a module on views of the treatment received, and for the treatment group, questions on their experiences of the intervention.

As well as collecting information from patients at baseline and as part of a later follow-up survey, there may be value in obtaining information on the patient from existing datasets. These might include Hospital Episode Statistics and even information on benefit claims, earnings and employment. This could reduce the length and complexity of the survey and improve the quality and accuracy of the data.

5.2.4 Outcome measures

Whilst HbA1c levels would be the main clinical indicator of the impact of health coaching on diabetes, it would be advisable to collect information on a wider range of aspects of patient health within an RCT, both to ensure that the treatment and control groups are well-matched and to capture any impact of health coaching on other outcomes. Diabetes patients are supposed to have access to mental health assessments and support and so a measure of mental health, such as GAD-7, PHQ-8 and HADS may also be beneficial (Diabetes UK 2018: 3). It can be challenging to administer these measures in the context of a telephone survey, so some studies omit the item on whether the patient has thought they would be better off dead or hurting themself to reduce the likelihood of distress. It may also be helpful to use a generic measure, such as PAM, as a primary outcome.

The list of secondary outcomes could be drawn from those suggested in section 4.2. It may also be helpful to use the management information to explore whether health

coaching results in better attendance at appointments. A greater demand for appointments by patients would increase resourcing requirements, but having patients who were more engaged in their treatment may have longer-term benefits if it reduces the risk that they experience health complications. Information on changes in the need for medication over time would also be informative in valuing the impact of health coaching.

5.2.5 Ethics

It is important to consider the ethical implications of asking patients to take part in an RCT. In the case of health coaching, case studies have suggested that the intervention is unlikely to be harmful to participants. It is also unlikely to impose a substantial additional burden on participants compared with the standard approach to diabetes monitoring and treatment.

The greatest risk of distress or harm is likely to stem from the data collection requirements. It is important to keep data collection activities to a minimum to ensure that both treatment and control groups are not over-burdened. It is also important to test the proposed questions, or to use questions tested in past studies, to ensure that they do not cause undue distress. They must also be administered in a way which is sensitive to the risk of harm.

A further ethical consideration is that as practices will benefit financially from the participation of patients in the trial, it is important to ensure that patients are made aware that they are not required to participate. It will be important to brief GPs and nurses on the need to make this clear to patients and to ensure that patients are aware that they can withdraw from the trial if they wish.

With a randomised control trial it is necessary to specify at the outset the circumstances under which the trial might be stopped. This would include monitoring outcomes for the treatment and control groups and setting a threshold for any negative impacts which would result in the trial being stopped. Equally, it might include stopping the trial early if there was clear evidence of a substantial benefit. A trial protocol would consider the appropriate level of these thresholds.

5.2.6 Registering the trial

As a trial of health coaching involves carrying out research on patients, it requires approval from the Health Research Authority (HRA). There is an established application process and any changes to the research materials require further approval following the initial application.

As GP practices receive payments for participating in trials, it is necessary to supply the National Institute of Health Research with monthly updates on the number of patients randomised to the trial and some basic information on the practice. This information is anonymised and supplied through an online portal.

It is also good practice to produce a protocol for the trial which sets out the primary and secondary outcome measures, the expected impacts and information on ethical considerations. This ensures that initial expectations are transparent. Trial registration is

also good practice, and essential for studies to be accepted for some journals. The registration must take place before the trial starts recruiting patients.

5.3 Quasi-experimental design

Some NHS sites have existing licences to use PAM and so it may be possible to explore the impact of health coaching through the analysis of existing data on patients in areas where a health coaching approach has been used, compared with areas where health coaching is not available. This study would not focus on the impact of health coaching on a single condition, but rather on its impact across a range of different conditions.

To build up a picture of the impact of health coaching across a broad range of outcomes and taking into account the patient's characteristics prior to the start of health coaching, it would be necessary to match the PAM score to information on the patient from other sources, such as those available through NHS Digital. By working with other government departments it may be possible to negotiate access to a wider range of data sources on the patient, including information on benefit receipt, employment and benefit history and earnings.

If it was possible to obtain data on a rich range of characteristics and to observe which patients were likely to have experienced health coaching, it might be possible to use a quasi-experimental method, such as propensity score matching, to estimate outcomes for a comparison group of patients with a similar propensity to make use of health coaching and to achieve similar outcomes to the treatment group. Existing data sources could be used to estimate impacts on many of the other outcome measures listed in section 4.2. This might include NHS workforce data.

6 Recommendations

This report has considered the challenges to conducting an economic evaluation of health coaching and considered the ways in which some of these challenges could be reduced or overcome. It also sets out the data requirements, including identifying likely sources of data and summarising the key outcome measures and suggests primary and secondary outcomes. It identifies the key factors which would need careful thought before an impact evaluation is likely to be possible. The report concludes by outlining the basic requirements for a randomised control trial, focusing on patients newly diagnosed with type 2 diabetes and also suggests a possible design for a quasi-experimental study to explore the impact of health coaching on a wider group of patients.

In summary, the main recommendations are to pursue a two-pronged approach, applying for funding for an RCT to look at the impact of health coaching on type 2 diabetes, focusing on the outcomes of HbA1c, and either GAD-7, PHQ-8 or HADS, as well as a range of secondary outcome measures to capture wider economic impacts. This would ideally involve analysis of administrative data from NHS Digital and other government departments as well as data collected during the course of the trial from management information and baseline and follow-up surveys of participants. The quasi-experimental study would be based solely on the analysis of existing data in areas which already use PAM, but would also seek to link this to other information on the patient from other data sources, as well as information on the local NHS workforce and information on whether a health coaching approach was offered to patients in that particular area.

Having conducted these studies to obtain a quantitative estimate of the impact of health coaching on a range of different outcomes, this information would be used to consider the costs and benefits of health coaching to the exchequer. Section 4.2 outlines the areas where the costs and benefits are expected to be felt and the economic evaluation would seek to attach a value to these potential impacts, also exploring how the conclusions vary depending on assumptions made about these values.

7 Bibliography

Diabetes UK (2017) 'Our Research Strategy 2015-2019'. London: Diabetes UK.

Diabetes UK (2018) 'Prevention, Targets and Therapies for Type 2 diabetes Clinical Studies Group Annual Progress Report'. Diabetes UK.

Edgren, G., Anderson J., Dolk A., Torgerson J., Nyberg S., Skau T., Forsberg B. C., Werr J., Öhlen G. (2016) 'A case management intervention targeted to reduce healthcare consumption for frequent Emergency Department visitors: results from an adaptive randomized trial.' European Journal of Emergency Medicine. 23(5):344-50.

Hemming, K., Haines, T. P., Chilton, P. J., Girling, A. J. and Lilford, R. J. (2015) 'The stepped wedge cluster randomised trial: rationale, design, analysis, and reporting', British Medical Journal. 350:h391.

Hibbard, J.H., Greene, J. and Overton, V. (2013) 'Patients with lower activation associated with higher costs; delivery systems should know their patients' 'scores'.' Health Affairs. 32(2):216-222.

Kibble, S., Gray, D. Prat-Sala, M., Ross, K., Johnson, K., Packer, J. Shire, E. Cross, R. and Harden, B. (2014) 'Recovery coaching in an acute older people rehabiliation ward' BMJ Quality Improvement Reports. 3(1):1-5.

Lindsay, A., Hibbard, J. H., Boothroyd, D. B., Glaseroff A., Asch, S.M. (2018) 'Patient Activation Changes as a Potential Signal for Changes in Health Care Costs: Cohort Study of US High-Cost Patients', Journal of General Internal Medicine. 33(12): 2106-2112.

Sforzo, G. A., Kaye, M. P., Todorova, I. (2017) 'Compendium of the Health and Wellness Coaching Literature', American Journal of Lifestyle Medicine. 12(6):436-447.

Willard-Grace, R., Chen, E. H., Hessler, D., DeVore, D., Prado, C., Bodenheimer, T., and Thom, D. H. (2015) 'Health Coaching by Medical Assistants to Improve Control of Diabetes, Hypertension, and Hyperlipidemia in Low-Income Patients: A Randomized Controlled Trial.' Annals of Family Medicine. 13(2):130-138.

Wolever, R. Q., Simmons, L. A., Sforzo, G. A., Dill, D., Kaye, M., Bechard, E. M., Southard, M. E., Kennedy, M., Vosloo, J., Yang, N. (2013) 'A Systematic Review of the Literature on Health and Wellness Coaching: Defining a Key Behavioral Intervention in Healthcare', Global Advances in Health and Medicine. 2(4): 34-53.