An approach to make general practitioner referrals suitable for artificial intelligence deployment

Evelyn Lesiawan, Bruce Sutherland, Christoph Schumacher, Andrew Cave, Guy Armstrong

ABSTRACT

Outpatient referrals for hospital specialist assessment are an increasing workload that carry significant risk if not attended to in a timely manner.

This viewpoint discusses how decision support (including artificial intelligence and machine learning) may address this problem. Of the many possible approaches, we choose a combination of two that illustrate the breadth of available tools and how they combine to complement each other.

To understand the issues and inform this discussion, a survey of general practitioners' views was conducted (Appendix 2), an audit of declined referrals was undertaken (Appendix 3) and draft decision trees were constructed (Appendix 4).

To have data suitable for automated decision support, the current referral needs to change from free text to a structured format that ensures every patient has a complete minimum dataset. Regarding triaging decisions, at present there is human variability, but the decision support tools will need to be trained on a set of referrals that have an agreed gold-standard decision. In order to maintain patient safety throughout, the process needs to be incremental. We suggest that one way to assure patient safety is to combine simple decision trees with sophisticated contemporary machine learning.

Support for applying artificial intelligence (AI) to healthcare has recently been expressed at ministerial levels.[1,2] In a healthcare sector beset by staff shortages and limited funding, AI is said to have "*the potential for very high return on investment*".[3] This viewpoint examines the potential for decision support, such as AI, to assist with triaging general practitioner (GP) referrals to cardiology outpatients at Health New Zealand – Te Whatu Ora Waitematā. The general principles and approach discussed have the potential to scale and extend more widely to other districts and specialities.

Keeping pace with the number of outpatient referrals is challenging for hospital specialities. The challenges include the sheer volume of referrals, with their year-on-year increase (Figure 1), as well as the need for timeliness in investigating and treating.[4] Some cardiology conditions carry a mortality risk, making prompt assessment important for more than quality of life reasons alone. AI and conventional decision support techniques have the potential to assist, and we examine how current processes could be adapted for such deployment.

Many decision support tools could be applied to this problem. We have chosen two techniques that are particularly illustrative for clinicians seeking to understand the available options. There are many other tools available that could be applied to this problem.

The two techniques used for illustration are human-designed decision trees and contemporary machine learning (ML). They are at opposite ends of the complexity spectrum of available methods. We suggest that, by employing two complementary approaches, the power of contemporary, sophisticated ML is harnessed while ensuring clinical safety through a simpler, more transparent technique, particularly during the initial deployment.

This viewpoint first defines relevant terms, then describes the current process for handling outpatient referrals, followed by a discussion of barriers to implementing decision support. We then describe a possible approach to addressing these issues by combining decision support techniques in a stepwise process. The aim is to maintain patient safety at every stage of the development process, yet culminate in maximising the benefit from contemporary sophisticated decision support tools. Our findings and proposals

are informed by a survey of GP views (Appendix 2) and an analysis of declined referrals (Appendix 3).

This is an account "from the trenches" designed for the non-expert; a comprehensive position paper is available.[3]

Definition of terms

Decision support: decision tree versus ML

A **decision tree** (flow chart) can help GPs provide all relevant information by using structured questions. A decision tree breaks the decision into a series of simple "yes/no" questions, as per the examples in Appendix 4. The transparency of the decision process makes the tree educational for users. By forcing a stepwise assessment, the chart ensures that referrals are graded on a consistent set of criteria, thus reducing the variation that occurs between human triagers.

A **ML** model would learn from a large dataset of past referrals that have an agreed triaging decision. The model can detect patterns too subtle or complex for a simple decision tree—for instance, combinations of symptoms that, while individually mild, tend to lead to referral acceptance when seen together. Unlike the static decision tree, the ML model continues to learn as more referral and outcome data are fed in, making it adaptive and more accurate than rigid criteria. ML could either replace or augment the decision tree.

AI: ML compared with large language model (LLM)

ML refers broadly to algorithms that learn patterns from data to make predictions or classifications. In contrast, **LLMs** are statistical systems trained on vast amounts of text and other unstructured data to generate language that resembles human communication. In medicine, ML might be applied to imaging or laboratory data to predict disease or identify abnormalities. In contrast, LLM can summarise patient records, draft clinical correspondence, provide natural language responses to medical queries or interpret text entries.

ML models offer a significant advantage over traditional approaches because they do not require researchers to fully specify the structure of relationships in advance. Instead, they sift through data and detect subtle or unexpected patterns that humans might never have anticipated, making them powerful tools for uncovering new insights. The trade-off, however, is that these models often operate as a "black box", producing results without making clear why certain connections were drawn. This can lead to spurious or misleading associations being treated as meaningful, including the embedding of existing biases. For this reason, while ML expands the frontier of what can be uncovered, oversight remains essential to interpret results responsibly and to ensure that identified patterns are both accurate and relevant. An ML tool with excellent average performance will still produce a small percentage of incorrect decisions, and these may be consequential for individual patients.[5]

The problem

Overview of outpatient service and staffing

The current referral process entails a GP making an electronic, largely free-text referral. Once registered by the booking and scheduling clerk, the triaging cardiologist either accepts (and prioritises) the referral or declines it. The clerks then act on these decisions. There are no automated decision aids.

Problems with the current triage service

Processing cardiology outpatient referrals consumes a considerable amount of resources. At Waitematā, nine doctors and one nurse do not keep pace with the triaging of referrals, which increase by 1,374 each year (see Figure 1). The number of clerical staff required for this manual process is not quantifiable, as they are shared across departments. Clinical risk to patients increases with delays at every step, initially with referrals waiting to be triaged, then accepted referrals waiting to be seen and finally seen patients awaiting subsequent investigations (Table 1). Automation with appropriate decision aids could expedite the triage of referrals and thereby reduce wait times and associated clinical risk.

However, the current free-text format delivers varying amounts of information. Triaging cardiologists want referrals to consistently contain specific information relevant to the reason for referral. If automated decision aids are to be useful and trustworthy, this minimum dataset for every patient becomes mandatory.

Another issue with the current human-led service is the (understandable) variation in decisions between individual cardiologists (e.g., see variation in referral decline rates in Appendix 3). This is despite regional recommendations on triaging (Appendix 1). Automation has the potential

Figure 1: Number of cardiology outpatient referrals to Health New Zealand – Te Whatu Ora Waitematā, 2014–2023.

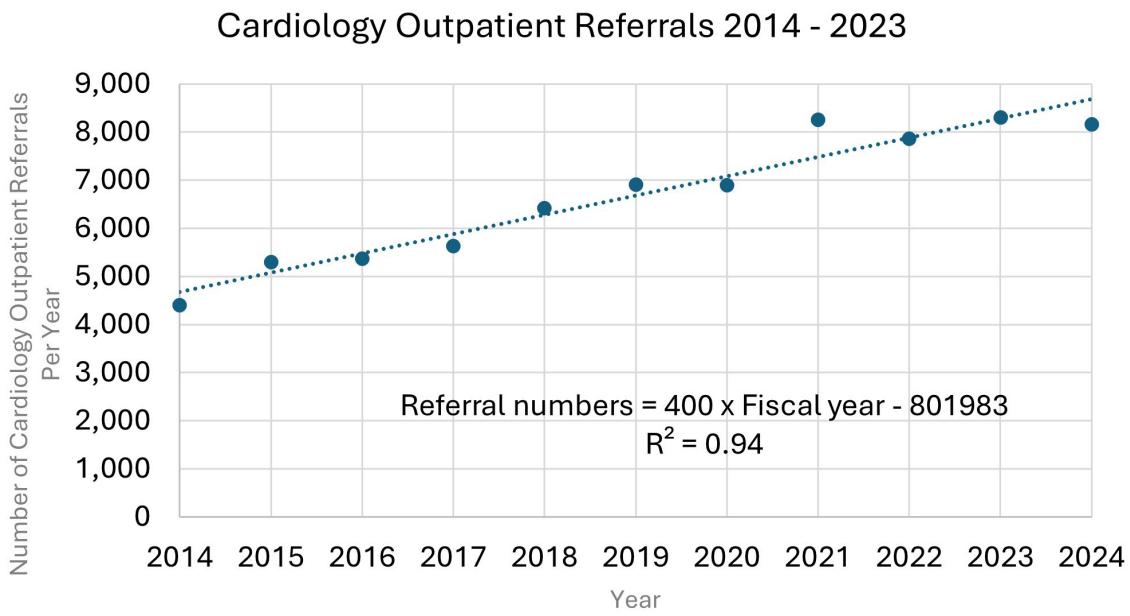

Table 1: Outpatient referral pathway.

Node	Aim	How to achieve aim
A) GP referral	Capture information that is structured and adequate	Information fields filled by clicks and drop-downs, or by voice recognition
	Facilitate GP efficiency	GP dashboard
B) Referral triage	Increase appropriate declines and acceptances	Iterative development of algorithms/decision trees
	Facilitate triager efficiency	Triager dashboard
C) Accepted referral	Increase proportion of FSA done at triage (combined triage and non-contact FSA)	Incentivise triagers
D) Clinic FSA	Reduce need for follow-ups More remote evaluations	Incentivise contingent planning in initial FSA letter Test novel approaches such as telehealth and biomarkers

Possible intervention points to improve the outpatient referral process and prepare the ground for AI: A) Instead of free text, decision trees capture an adequate amount of information in a structured way. Informational aspects are built in to assist GPs in determining what conditions and severity warrant referral. A dashboard containing representations of (and links to) all pertinent information on one screen enhances efficiency and reduces clicks per task. B) Decision trees are iteratively refined until they can reject or accept a proportion of referrals without human intervention. A dashboard allows triaging cardiologists to access relevant cardiac information for the patient more readily. C) Not covered in this study, but once referrals are accepted, a greater proportion of FSA are carried out as a virtual non-contact FSA by the triager providing written advice to the GP. Concluding the FSA at this stage avoids the bulge of referrals moving on to increase wait times for clinic FSA and procedures such as echocardiography and Holter monitoring. D) Not addressed in this study. See text for details.
GP = general practitioner; FSA = first specialist assessment; AI = artificial intelligence.

to reduce variability with its more consistent decisions better reflecting clinical risk. However, training a decision aid is best done with a library of referrals, each with an agreed reference decision.

A solution

Summary

We are concerned that applying black box ML to this dataset may be clinically risky without an intermediate step; for this reason, we suggest using human-designed decision trees. These two decision aids may prove complementary by feeding back to each other. We acknowledge that any decision aid can make mistakes, and the process needs to be structured to minimise the impact on individual patients. Table 2 explains some of the important concepts. Other suitable options, such as "explainable" AI, are not considered here.

Information from GPs that is sufficient and structured

Our illustrative strategy involves a first step of generating, for each patient, a sufficient minimum set of data relevant to the presenting symptom. The draft decision trees in Appendix 4 contain examples of relevant information for some common conditions; they will need to be further developed and improved. Taking chest pain as an example, it is essential to know whether it is exertional and whether it is accompanied by shortness of breath. Such a minimum dataset will help both the current human triagers and any subsequent automation.

This specific information required by cardiologists for each condition could be obtained from GPs using a structured form using tick boxes (still with room for free-text additions). The information would then be digitally incorporated into the hospital information system, relieving clerical staff of the chore of (error-prone) manual entry.

Without the use of structured forms, those patients lacking the minimum dataset of information are at risk of variable and erroneous decisions, whether made by humans or automation. It is not realistic to expect GPs to know what information specialists require. Our GP questionnaire found that only 29% of GPs are confident about the information required for a cardiology referral, and only 50% are confident about which conditions are typically seen by cardiology (Appendix 2).

Only limited information (variables) can be requested from the GP to avoid the process becoming too time-consuming. However, GPs are increasingly using "AI medical scribes", which may facilitate gathering more information without them being overwhelmed. Nevertheless, there will always be substantial amounts of pertinent information missing from a referral. Some information is currently unavailable because it has not yet been obtained (e.g., through a future Holter monitor or echocardiogram), while other details are missing because it would be unreasonable to expect GPs to provide too much information. Furthermore, individual patients will have uncommon factors that are idiosyncratically predictive for them but are not captured by any manageable process. Only a small fraction of the many relevant variables (dimensions) can be captured. These limitations mean that the training of any automation carries a risk to individual patients that needs to be managed carefully.

Agreed end points for training decision support

To train automated decision aids, a gold-standard decision is needed for reference. Currently, this doesn't exist, as human triagers (being human) exhibit variability in their decisions (Appendix 3). Such variation needs to be eliminated to supply a gold standard or reference decision for automation training. This could be achieved by a small number of cardiologists reaching consensus on cases to develop and validate decision trees. Human oversight would need to be an ongoing, iterative process, as new cases challenge the algorithm, clinical practice changes and thresholds alter due to changes in resourcing.[3] Below, we suggest that this onerous human oversight of ML may possibly be devolved to decision trees.

Going straight to ML may be problematic

ML learning in these sparse representations of high-dimensional data carries a risk of unpredictable outcomes, which is an unpalatable risk in clinical medicine.[6] We suggest that success with ML will be more likely if the ground is prepared before deployment. This could be achieved iteratively by first transitioning from free-text electronic referrals to a structured referral form that can support the implementation of decision trees, before considering full-blown black box ML (Appendix 4). The idea of having a transitional state of decision trees inserted between human triage and ML made even more sense when we

became aware that there had been a previous failed attempt at deploying AI to triage cardiology referrals at Waitematā.[7] We have been unable to obtain further information and so do not know the reason for failure.[7]

Human-designed decision trees

Decision trees mimic how a cardiologist makes decisions. By making these thought processes explicit, transparent and fully explainable, decision trees will serve as a helpful bridge between human triaging and subsequent ML.

Decision trees function predictably with small numbers of variables, although at the expense of being biased.[8] However, bias is easy to detect in decision trees, and their transparent nature allows identification and inclusion of the additional variables needed to mitigate the bias.

The act of cardiologists developing decision trees will facilitate the development of consensus reference decisions for subsequent training of ML.

Over time, decision trees may be able to replace humans in the task of ongoing oversight of ML decisions.

In a feedback loop, ML may identify new predictive variables that can then be incorporated into the structured questionnaire, enabling this variable to be obtained for all patients and thereby improving the predictive power of the ML.

Decision support makes mistakes

Like humans, any automated decision support will make mistakes. It may be that, as with self-driving cars, society will be less tolerant of automation error than of human error.[9] Decision trees exemplify the concept of "satisficing", which is finding a good-enough solution when it is not practical to find the optimal solution.[10,11] They may work better than ML when there are many unknown variables (as here).[12] In this setting, their output is more predictable than ML and more readily modified.[3,12] Simple decision trees mimic the decision making of legal judges surprisingly faithfully and may perform similarly for cardiologists, regardless of how sophisticated we view our own decisions as.[12]

Decision trees or ML can only ever be "probably approximately correct"[5] when evaluated on large numbers of patients. That means that decisions on individual patients have the potential to deviate sufficiently to be a clinical safety issue. At the outset, the error bounds for both the approximation and the probability are unknown, hence the need for oversight. We believe the human-designed decision trees provide the necessary safety, at least for the initial stages. Their deployment is carried out in a stepwise fashion, allowing for iterative refinement with minimal clinical risk and helping to maintain clinician trust.[3] Regardless, some misclassification will occur at the conceptual level due to the challenge of crystallising the diversity of human symptoms into binary variables. However, if data collection is digital, misclassification due to data entry errors and data handling will be minimised.[13]

At the opposite end of the spectrum are the black box forms of ML, where it is not possible to explicate or understand the rationale behind the algorithm's recommendation, or even which variables were used to predict the recommended outcome.

Sequence of deployment

ML is more powerful than decision trees, but, initially at least, it will be more prone to erratic and deviant results, given that the available information is sparse. ML needs a library of reference decisions before it can be trained.

Structured forms will provide more complete information, which is a necessary condition for any automation. Initial deployment of decision trees would be alongside, but invisible to, the human triagers. The discrepancies between human and algorithmic outcomes will be reviewed, and the algorithm will be refined accordingly. Next, the tree results are made visible to the human triagers for further refinement. Finally, a decision is made on whether some classes of referrals can be accepted or declined solely by the algorithm. There will be an indeterminate group where acceptance or rejection will need human input. Further iterations are performed to minimise the size of this indeterminate group. However, indeterminate presentations should never fall to zero, as there will always be complex and poorly differentiated cases. Forcing these into the algorithm risks misclassification.

Declined referrals should be accompanied by standardised information to the GP on why the referral was rejected, together with suggestions for management. This will assist the 42% of GPs who perceive that a referral has been declined inappropriately, and it functions as a just-in-time education tool for the GP referrer (Appendix 2).[14]

The next step is to train the ML model on the entire content of the referral, including both structured information and free text. This latter may be

Table 2: Elements of automated decision support for GP referrals.

Term	Description
Structured information	Required for whatever decision support method is chosen. Provides a minimum dataset, which makes ML predictions less variable and facilitates more rapid training of AI.
Human-designed decision trees (transparent predictions)	Designing these will facilitate cardiologists' consensus in the development of reference decisions for training ML. Can be improved iteratively as cases are encountered that don't fit the current tree decision. Contain domain knowledge of cardiologists that can frontload the training of MI. Can replace human oversight of ML decisions. Interacts in a continuous feedback loop whereby new predictors from ML are added to the decision tree, which then provides a more complete dataset on every patient. This makes ML predictions more consistent and reliable.
Machine learning (opaque predictions)	Extends the predictive power of decision trees by incorporating free-text information contained in referrals. Identifies new predictors that are then incorporated in the structured questionnaire, thus improving the minimum dataset obtained on each patient.

GP = general practitioner; ML = machine learning; AI = artificial intelligence.

extracted by an LLM and fed into the ML algorithm. The LLM output is a probabilistic tool (i.e., it may give a different result each time it is fed the same information) and so adds an additional element of unpredictable variation that needs oversight to ensure the safety of individual patients.

It may be most useful for the decision trees and ML to proceed in parallel. If the ML identifies novel variables, these could be fed back and incorporated into the decision trees. As clinical safety is assured, increasing weight can be given to ML, which should eventually overtake the decision tree in its predictive ability.

Limitations

The current Auckland Region eReferral system is not suited to our proposals. We understand there is a project to review and upgrade the software, which will be more suited to deliver the dashboard views (Appendices 2 and 3).

This report is on the current system for referral triaging. However, this is embedded within the larger New Zealand health IT infrastructure, and the need for compatibility will influence which solutions are most appropriate.

The decision tree concept was developed with support and input from the Waitematā cardiology liaison GP, but has not yet been discussed more widely within the GP community. However, the orthopaedic service has successfully implemented decision trees with tick boxes and drop-down selections. Informal discussions with GPs indicate that they have accommodated the increased time required to complete the forms by scheduling a separate appointment specifically for the orthopaedic referral. They see an advantage in the form, as it provides an immediate answer as to whether the patient qualifies for joint replacement surgery and at what priority. This enables real-time discussion between the GP and the patient about the reasons for acceptance or rejection. GP practice software will need to incorporate the decision trees and AI software, which will take time to implement. It may also be helpful to incorporate the existing GP guidelines ("health pathways") into this software, providing more

ready access to advice.[14]

Insufficient information is available for a quantitative cost–benefit analysis. There are alternative approaches to this problem, utilising different automation tools. We view the approach outlined here as particularly illustrative for those unfamiliar with decision support tools.

Conclusion

There are many opportunities for AI to assist healthcare. This viewpoint examines the potential for automated decision support, including AI, to assist in triaging GP referrals. It has the potential to improve efficiency, reduce personnel requirements and provide more consistent decisions when compared with human triagers. We review an approach that is illustrative for those unfamiliar with decision support, while acknowledging that other options will be suitable.

The volume of GP referrals is substantial and poses significant challenges. However, two steps are required to prepare the ground for automation. Firstly, adequate information is crucial, especially for black box ML. A minimum dataset is necessary for every patient, which requires switching to a structured referral form instead of free text. Secondly, a library of historical referrals with reference (gold standard) decisions is needed for the training of automation.

We suggest that human-designed decision trees can complement contemporary black box ML by mitigating the risk of erroneous decisions that may affect the safety of individual patients.

COMPETING INTERESTS
Nil.

AUTHOR INFORMATION
Evelyn Lesiawan: Advanced Physician Trainee, Health New Zealand – Te Whatu Ora, New Zealand.
Bruce Sutherland: General Practitioner, Kawau Bay Health, Warkworth, Auckland, New Zealand.
Christoph Schumacher: Professor of Economics, School of Economics and Finance, Massey Business School, Albany, Auckland, New Zealand.
Andrew Cave: Digital Hospital Implementation Lead, Data & Digital, Integration & Delivery, Health New Zealand – Te Whatu Ora Northern, New Zealand.
Guy Armstrong: Cardiologist, Health New Zealand – Te Whatu Ora Waitematā, Auckland, New Zealand.

CORRESPONDING AUTHOR
Guy Armstrong: Lakeview Cardiology, North Shore Hospital, Private Bag 93 503, Takapuna, Auckland, New Zealand.
E: guy.armstrong@waitematadhb.govt.nz

URL
https://nzmj.org.nz/journal/vol-138-no-1627/an-approach-to-make-general-practitioner-referrals-suitable-for-artificial-intelligence-deployment

REFERENCES
1. Bhamidipati S. Experts cautiously optimistic about expanding AI in education, health sectors [Internet]. RNZ; 2024 Jul 3 [cited 2024 Jul 22]. Available from: https://www.rnz.co.nz/news/national/521217/experts-cautiously-optimistic-about-expanding-ai-in-education-health-sectors
2. Espiner G. AI for school tutoring, instant medical analysis part of NZ's future - Judith Collins [Internet]. RNZ; 2024 Jul 3 [cited 2024 Jul 22]. Available from: https://www.rnz.co.nz/news/in-depth/521123/ai-for-school-tutoring-instant-medical-analysis-part-of-nz-s-future-judith-collins
3. Office of the Prime Minister's Chief Science Advisor. Capturing the benefits of AI in healthcare for Aotearoa New Zealand [Internet]. [cited 2025 Jun 15]. Available from: https://www.pmcsa.ac.nz/artificial-intelligence-2/ai-in-healthcare/
4. Health New Zealand – Te Whatu Ora. Health Delivery Plan [Internet]. 2025 Jul 17 [cited 2025 Oct 1]. Available from: https://www.tewhatuora.govt.nz/publications/health-delivery-plan
5. Valiant L. Probably Approximately Correct. Basic Books; 2013. p. 57-86.
6. Li R, Kumar A, Chen JH. How Chatbots and Large Language Model Artificial Intelligence Systems Will Reshape Modern Medicine: Fountain of Creativity or Pandora's Box? JAMA Intern Med. 2023 Jun 1;183(6):596-597. doi: 10.1001/jamainternmed.2023.1835.
7. NZTech. AI about to speed up hospital processes – big time [Internet]. NZTech. 2019 [cited 2024 Aug 21]. Available from: https://nztech.org.nz/2019/10/31/ai-about-to-speed-up-hospital-processes-big-time/
8. Yarkoni T, Westfall J. Choosing Prediction Over Explanation in Psychology: Lessons From Machine Learning. Perspect Psychol Sci. 2017 Nov;12(6):1100-1122. doi: 10.1177/1745691617693393.
9. Mayer MM, Buchner A, Bell R. Humans, machines, and double standards? The moral evaluation of the actions of autonomous vehicles, anthropomorphized autonomous vehicles, and human drivers in road-accident dilemmas. Front Psychol. 2023 Jan 4;13:1052729. doi: 10.3389/fpsyg.2022.1052729.
10. Simon HA. Rational choice and the structure of the environment. Psychol Rev. 1956;63(2):129-138. doi: 10.1037/h0042769.
11. Artinger FM, Gigerenzer G, Jacobs P. Satisficing: Integrating Two Traditions. Journal of Economic Literature. 2022 Jun;60(2):598-635. doi: 10.1257/jel.20201396.
12. Katsikopoulos KV, Şimşek Ö, Buckmann M, Gigerenzer G. Classification in the Wild: The Science and Art of Transparent Decision Making. The MIT Press; 2021. doi: 10.7551/mitpress/11790.001.0001.
13. Lash T, VanderWeele T, Haneuse S, Rothman K. Chapter 13: Measurement and Measurement Error. In: Modern Epidemiology. 4th ed. Wolters Kluwer; 2020.
14. Wikipedia. Just-in-time learning [Internet]. Wikipedia. 2024 [cited 2024 Jul 23]. Available from: https://en.wikipedia.org/w/index.php?title=Just-in-time_learning&oldid=1221408288

Appendices

Appendix 1

Te Whatu Ora
Health New Zealand

Northern Regional Cardiology Follow-up Guidelines

Version: May 2023 **Review Due**: May 2024

- These are follow-up guidelines only. Individual patient requirements may differ
- Always consider virtual follow-up
- Registrars should discuss follow up arrangements with their supervising SMO
- Letters to GP must clearly state follow up arrangements
- In general patients with hypertension or dyslipidaemia should be followed long term in primary care once condition stable/treated adequately
- Valves – clinical and echo follow up. If frail and comorbid (where valve intervention inappropriate) consider returning care to GP
- Rheumatic valve disease patients - refer to rheumatic fever nurse specialists if available, or ensure appropriate secondary prevention and follow-up as per valve lesions

Valves - follow-up frequency

Pre-Op	Mild	Moderate		Severe without current surgical indication
Aortic stenosis (including bicuspid valve)	5 years	2 years		6-12 months
Aortic regurgitation	Discharge	Normal LV 2 years	Dilated LV 1 year	6 months
Mitral stenosis	5 years	2 years		12 months
Mitral regurgitation *consider TOE for? repair	Discharge – Unless abnormal valve then 5 years	Normal LV size & function - 2 years	Dilated LV* - 1 year	6 -12months*

Post valve replacement/surgical mitral valve repair

- **All should have baseline post-op echo**
- **TAVI: 3 years until evidence of valve degeneration then annually**
- **Mechanical: 3 years**
- **Bioprosthetic: 3 years until evidence of valve degeneration then annually**
- **Surgical mitral valve repair: 1 year then every 3 years**

Post elective PCI and post ACS

If EF>50%, no further revascularization or device therapy planned, no other medical issues than nurse led cardiac rehabilitation clinic only. Otherwise, medical FU, timing to be determined at discharge.

Post CABG

Routine cardiac rehab follow-up 2-4 weeks, with medical FU up 6-12 weeks post op, then discharge unless EF<50% in which case repeat echo after 3 months of max. tolerated medical treatment to inform need for device therapy.

Heart failure

- Initiation/titration of heart failure disease-modifying drugs:
  - refer to (Nurse-led) Heart Function/Heart Failure Clinic, depending on local pathway/criteria
- Resistant/fragile clinical heart failure requiring case-management (especially diuretic regimes):
  - refer to (Nurse-led) Heart Function/Heart Failure Clinic
  - consider if suitable, referral for heart transplant assessment if not improving on medical/device therapy
- Patients < 75yrs with HFrEF (EF < 40%), once Rx optimised:
  - consider yearly FU
- Stable/max. medical treatment, no device therapy planned:
  - discharge to GP and/or referring clinician
- CRT/ICD + HF patients should be followed long term - 1-2 yearly
- If uncertainty re FU plans, please discuss with lead cardiologist

Atrial fibrillation

1. Anticoagulation decision finalised/good rate control and no other cardiac reason for follow-up: Discharge
2. Post DCCV: 4-6 weeks with an ECG, clear plan for future management including eligibility for repeat DCCV, rhythm or rate control communicated to GP
3. On flecainide/sotalol/amiodarone – 1-2 year FU
4. Post ablation FU – d/w EP team

Aortopathy

Genetic Aortopathy

- If suspected or confirmed: comprehensive history and refer to CIDG
- Refer to guidelines for timing and type of serial imaging and treatment for specific conditions

Degenerative aortopathy

Aortic dimension	3.5-3.9cm	4.0-4.4cm	4.5-4.9cm	5.0-5.4cm
Follow-up	Not required (consider follow-up at 5 years)	Repeat at 1 year If no progression, 5 yearly review	Repeat at 1 year If no progression 2 yearly review	Annual imaging
Other imaging modalities	Repeat assessments should use the same imaging modality. Consider CT or MRI if reliable images cannot be obtained by echo. CT or MRI should be performed if being considered for surgical intervention.			

Other

- CRT/ICD patients should be followed long term - 1-2 yearly
- HCM aged <70 every 1-2 years, >70 consider discharge

Northern Region
CARDIAC CLINICAL NETWORK

Appendix 2: outpatient cardiology referrals—a survey of views held by general practitioners (GPs)

Methods

A convenience sample of GPs within Waitematā completed an anonymous Google survey of their views on the current cardiology referral process. Auckland Health Research Ethics Committee approved the project (AH28636). When the authors were fielding phone calls from GPs during their clinical work, they invited each caller to participate in the survey. It was anonymous by design so as to encourage participation; therefore, no information is available on the profile of practices who participated.

Data analysis

This is a descriptive study. Responses are reported as percentages. There was no comparator group, and numbers were too small for statistical analysis.

Results: questions and response rates

1. How many years have you been practicing as a general practitioner?

Twenty-four GPs answered the questionnaire. Years in practice ranged from 1 to 40 years, with most

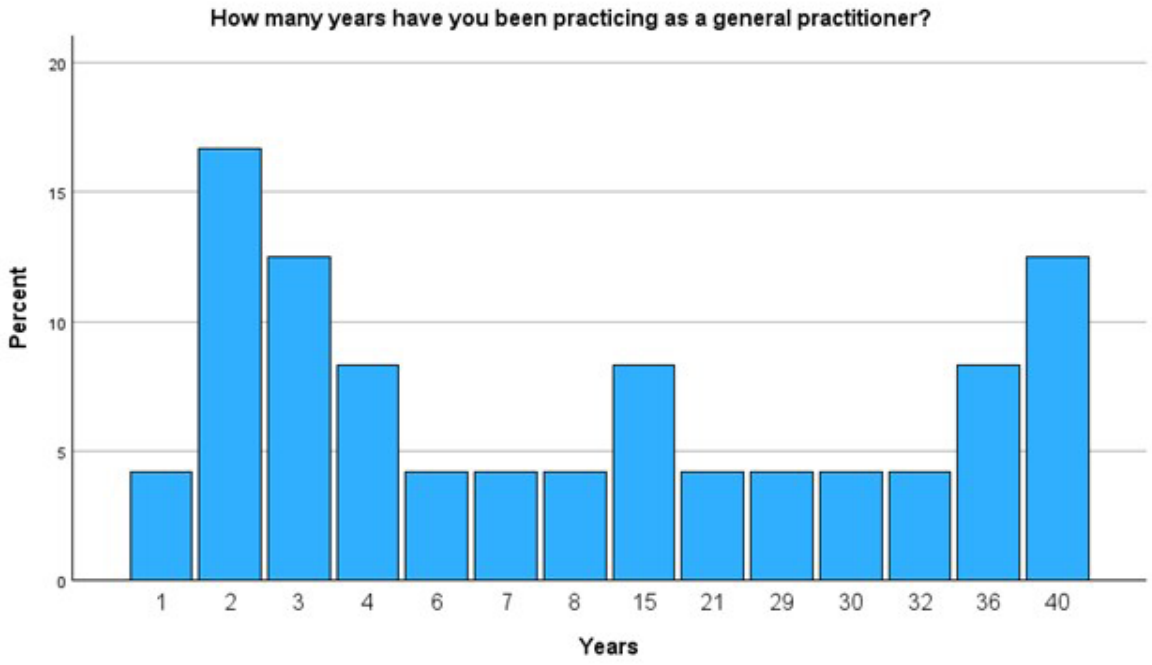

2–4 years and 36–40 years. Most (61%) referred to cardiology monthly, and 39% referred weekly.

2. How often do you refer to cardiology outpatients?

- □ Daily
- □ Weekly
- □ Monthly

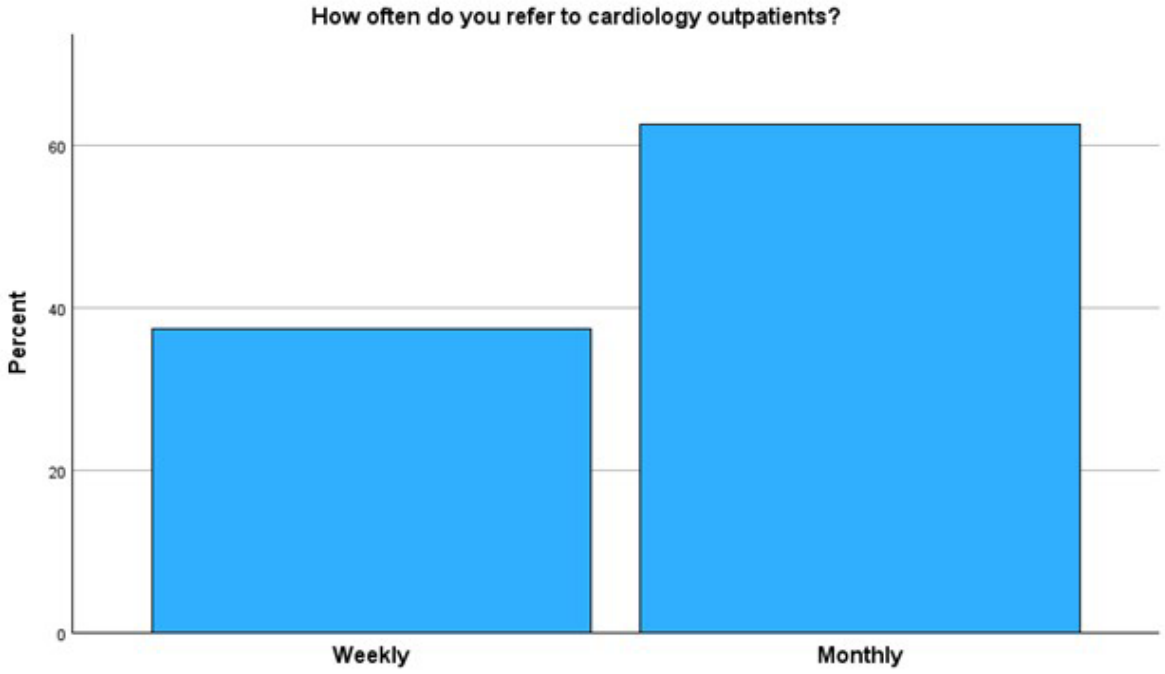

3. Do you feel confident about which conditions and presentations cardiology sees and which it does not?

- □ Yes
- □ No
- □ Sort of

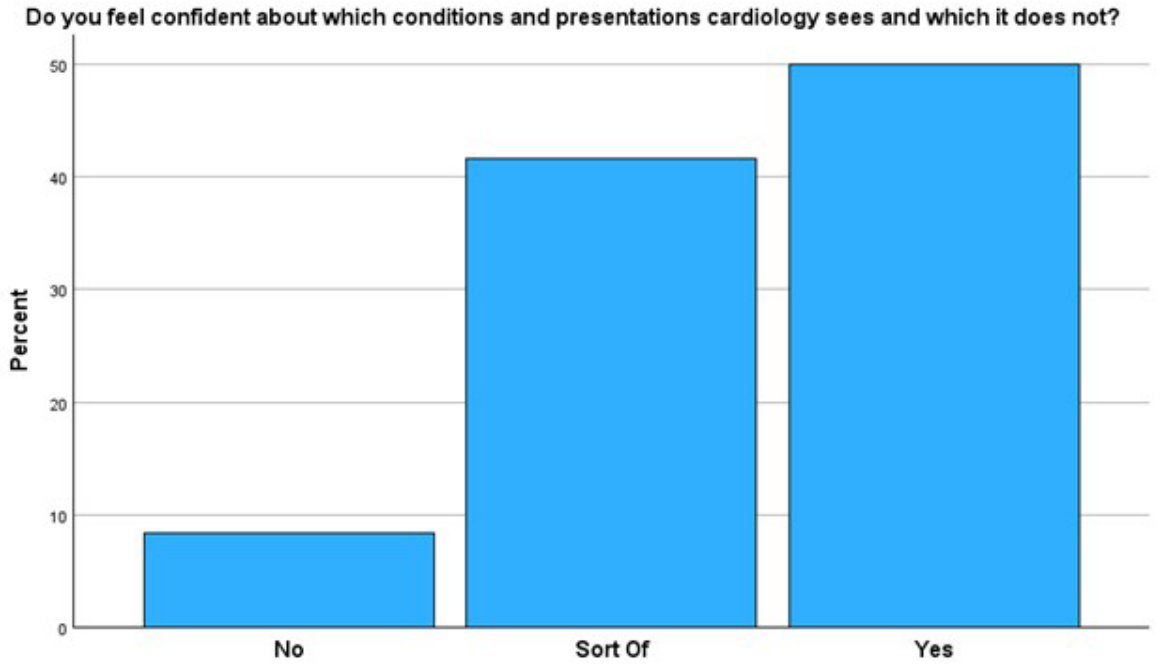

Only half are confident about which conditions and presentations were seen by cardiology.

4. Would it be helpful to have on the front page of the referral portal a list of the conditions we see and those we don't see?

□ Yes
□ No

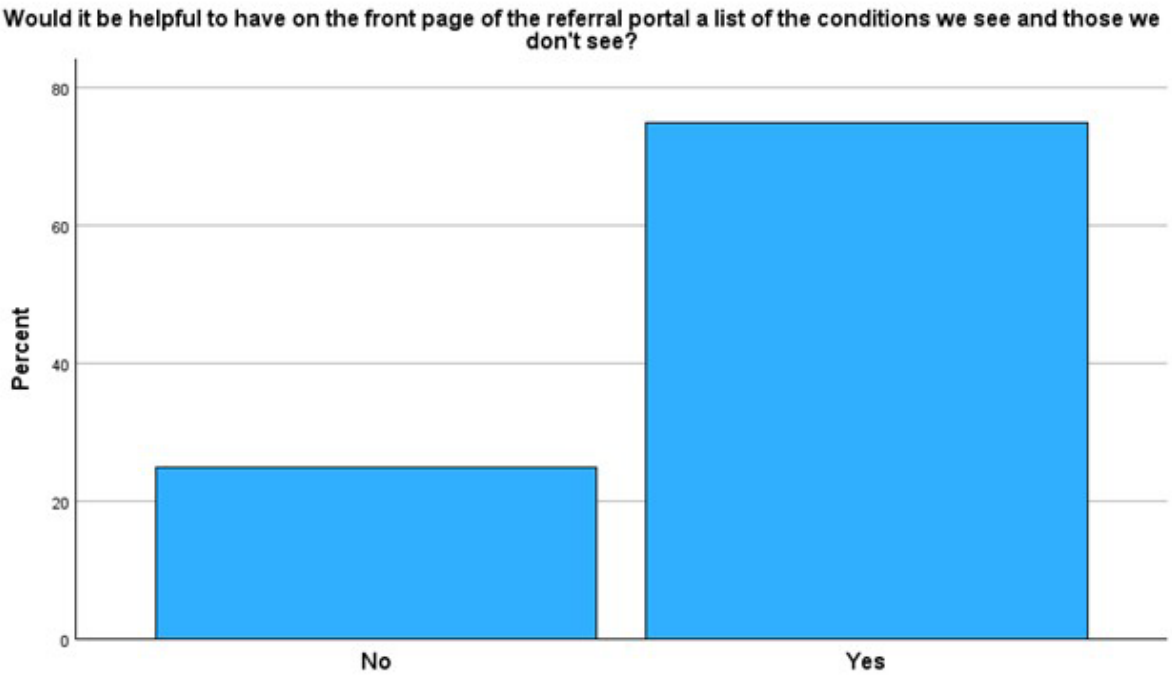

The majority (75%) would like a list on the referral portal of which conditions cardiology accepts and suggestions/links where to refer other conditions.

5. Can you tell the current wait time to be seen in cardiology or have a cardiology procedure? Yes/no

   Is it quick and easy? Yes/no
   Would it help to have this info more readily available? Yes/no

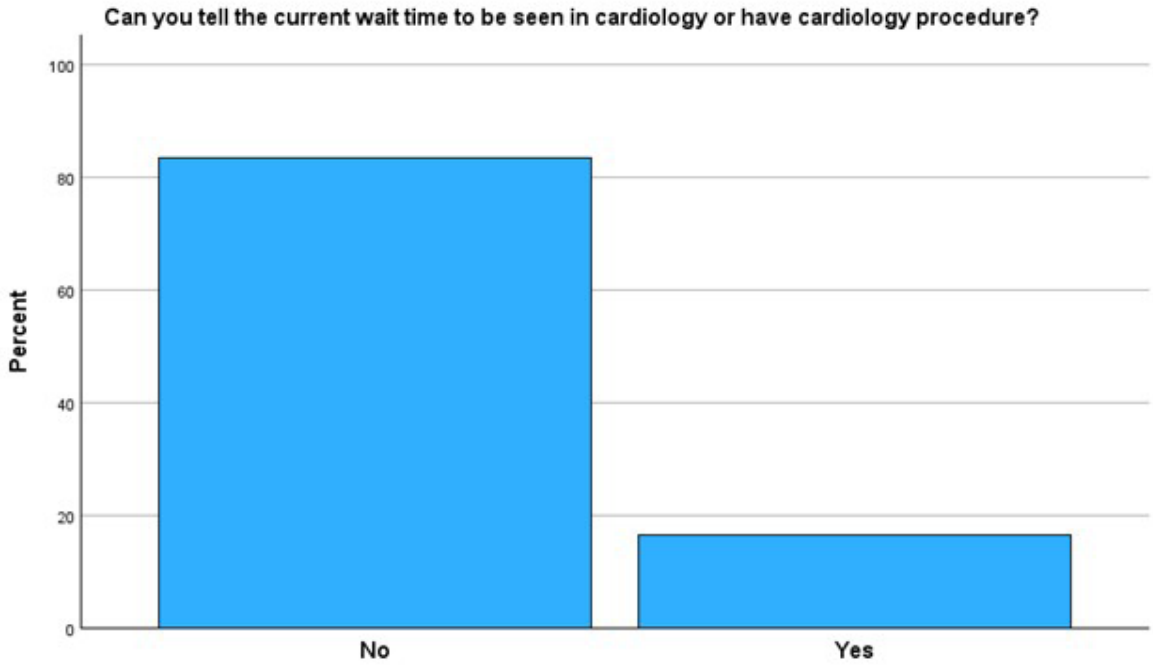
Can you tell the current wait time to be seen in cardiology or have cardiology procedure?
Percent
100
80
60
40
20
0
No
Yes

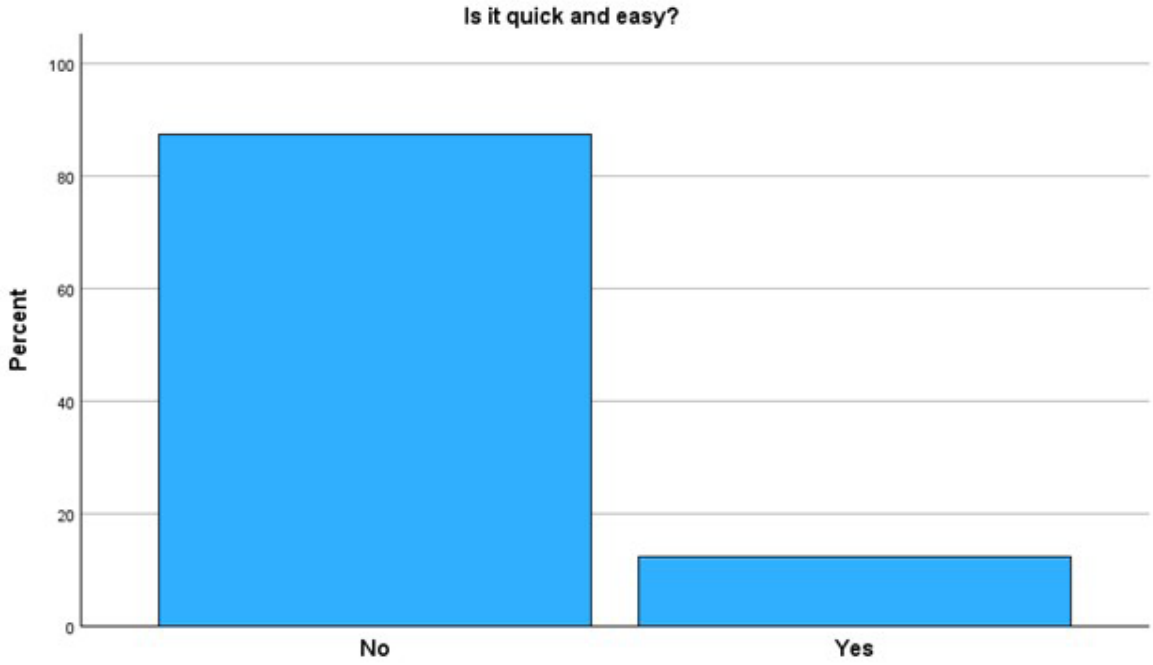
Is it quick and easy?
Percent
100
80
60
40
20
0
No
Yes

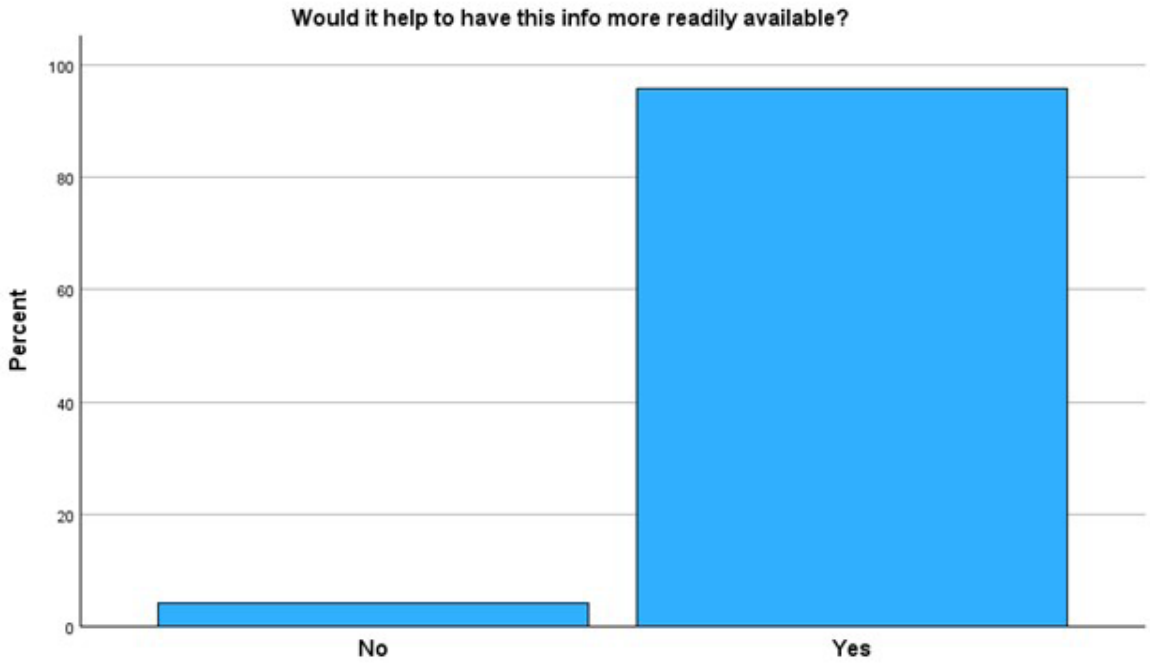

Most (83%) GPs cannot access the current wait times for clinic, with 96% wanting this information to be more readily available.

6. Can you access cardiology letters and reports? Yes/no

   Is it quick and easy? Yes/no
   Would it help to have this info more readily available? Yes/no

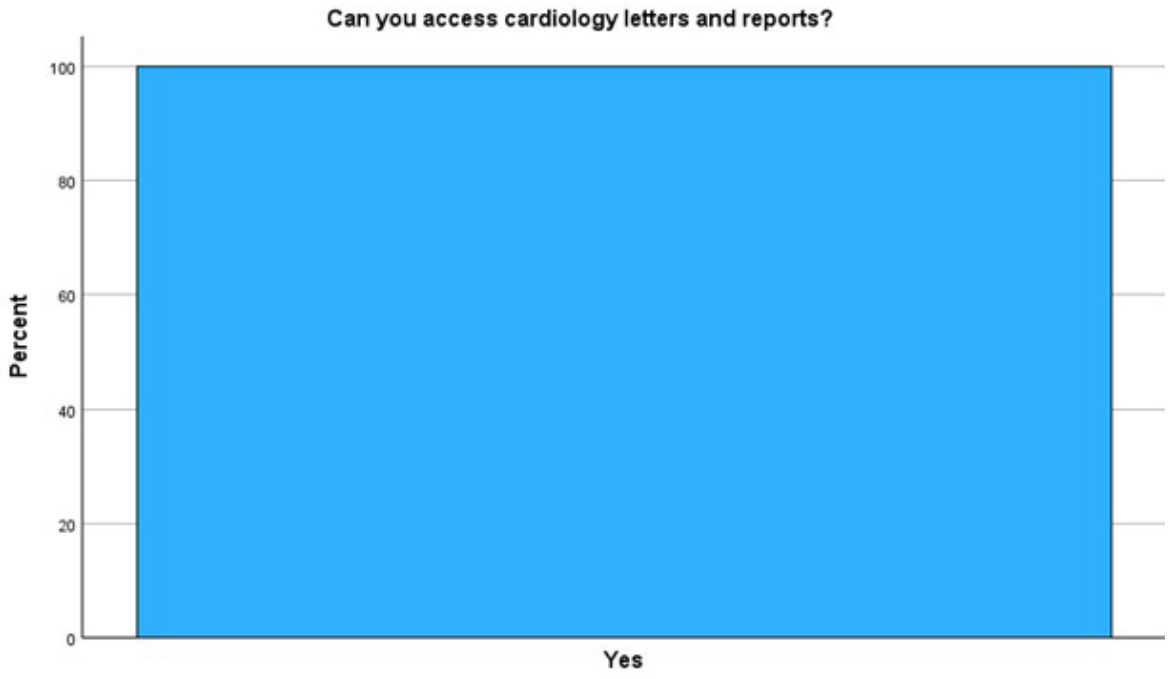

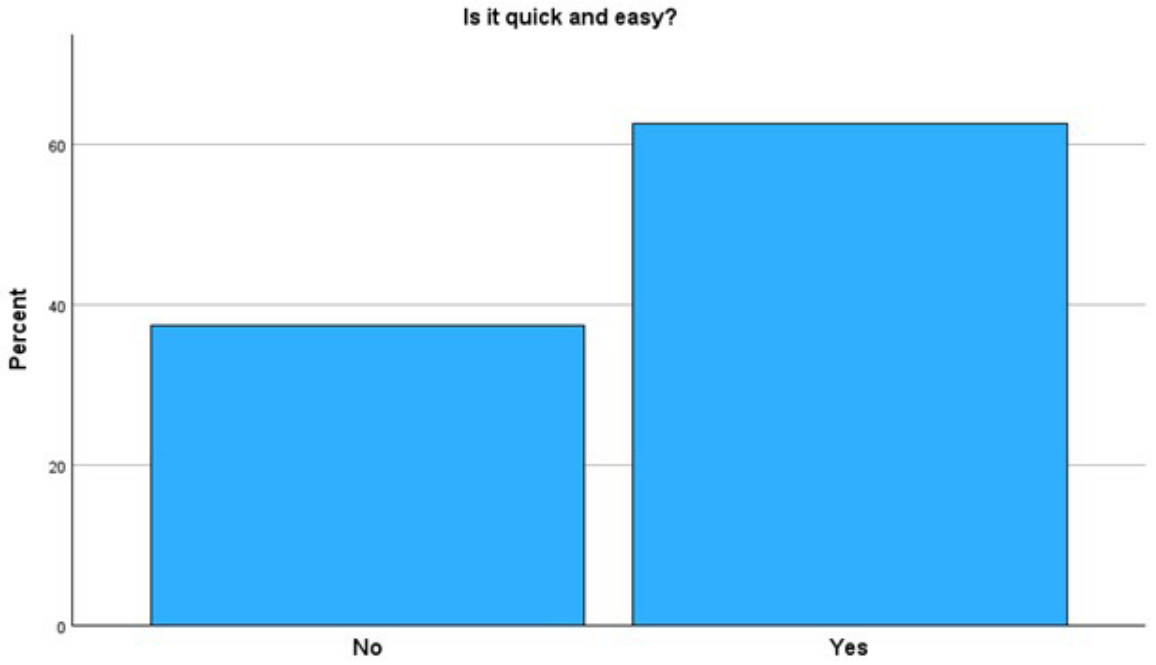

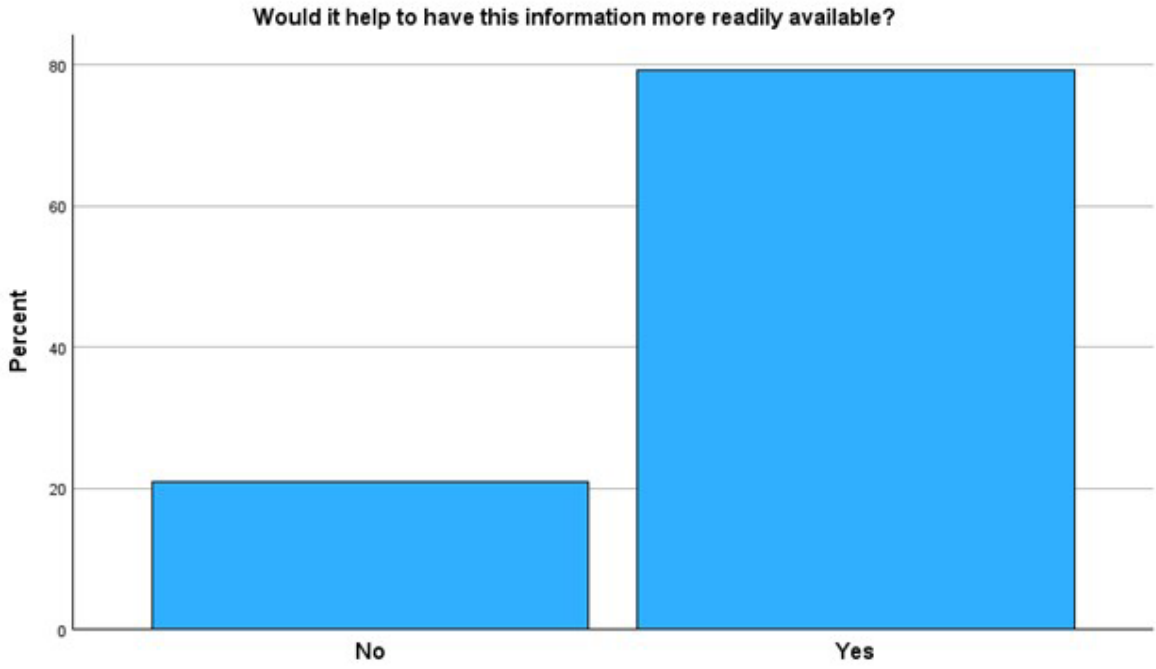

All respondents could access cardiology reports and letters, and it was quick and easy for 61%. Yet 79% felt it could be more accessible.

7. Do you know what information cardiology requires to assess your referrals?

□ Yes
□ No
□ Sort of

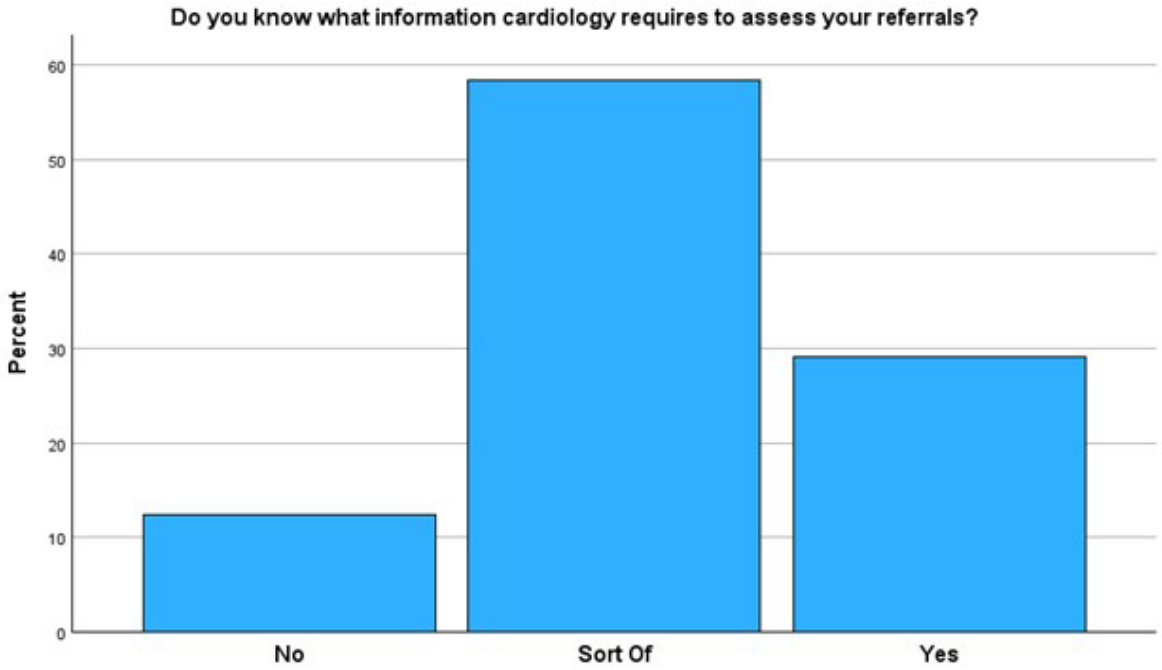

Only 29% felt confident about the information required for cardiology to triage a referral.

8. What resources do you use to assist with management of cardiology patients?

Please tick the resources that you have heard of:
□ Regional cardiology guidelines
□ NZTA guidelines
□ Health pathways
□ Other:

9. Please tick the resources that you use:

□ Regional cardiology guidelines
□ NZTA guidelines
□ Health pathways
□ Other:

10. If you use health pathways for cardiology advice—rate its usefulness:

[Not at all useful] 1, 2, 3, 4, 5 [Extremely useful]

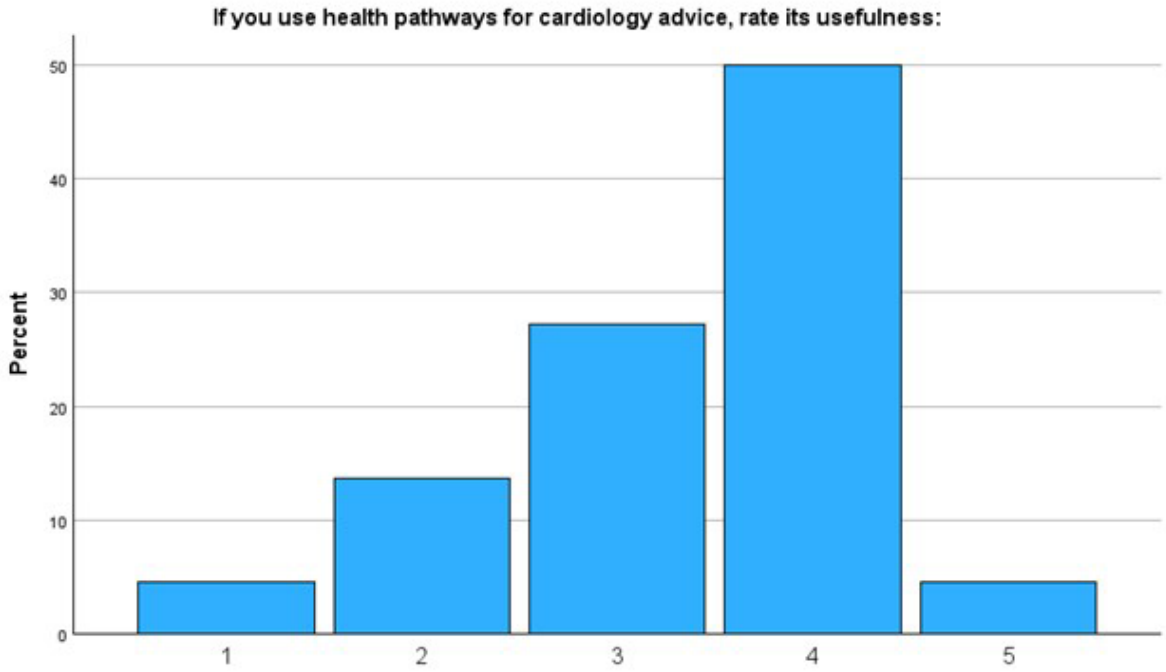

11. How can health pathways be improved?

Prominent on GP respondents' wish lists were shorter wait times and better availability of investigations (45%), quicker specialist advice (21%) and more comprehensive referral guidelines (health pathways) (29%).

12. Do you prefer obtaining the information you need via online resources or sending referrals to cardiology?

When it comes to preferred resources for cardiology issues, GPs were divided. Some relied on referrals to cardiology (33%), others used online resources (25%), and a significant portion used both (38%). Health pathways emerged as the most popular online resource, with 92% of respondents using it. However, not all found it helpful, as indicated by the responses to questions 10 and 11 above.

13. How could the cardiology referrals process be made easier?
14. What else would make it easier for GPs to access the information they need to manage cardiovascular conditions?

GPs' suggestions for improvement included flow charts, more explicit referral guidelines, surveillance information and information on cardiac medication. Notably, many GPs preferred succinct information, highlighting the need for simplicity and clarity in the referral process.

15. Is the cardiology referrals grading/decline consistent, or does it seem variable/random? (Circle)

□ Consistent
□ Variable/random

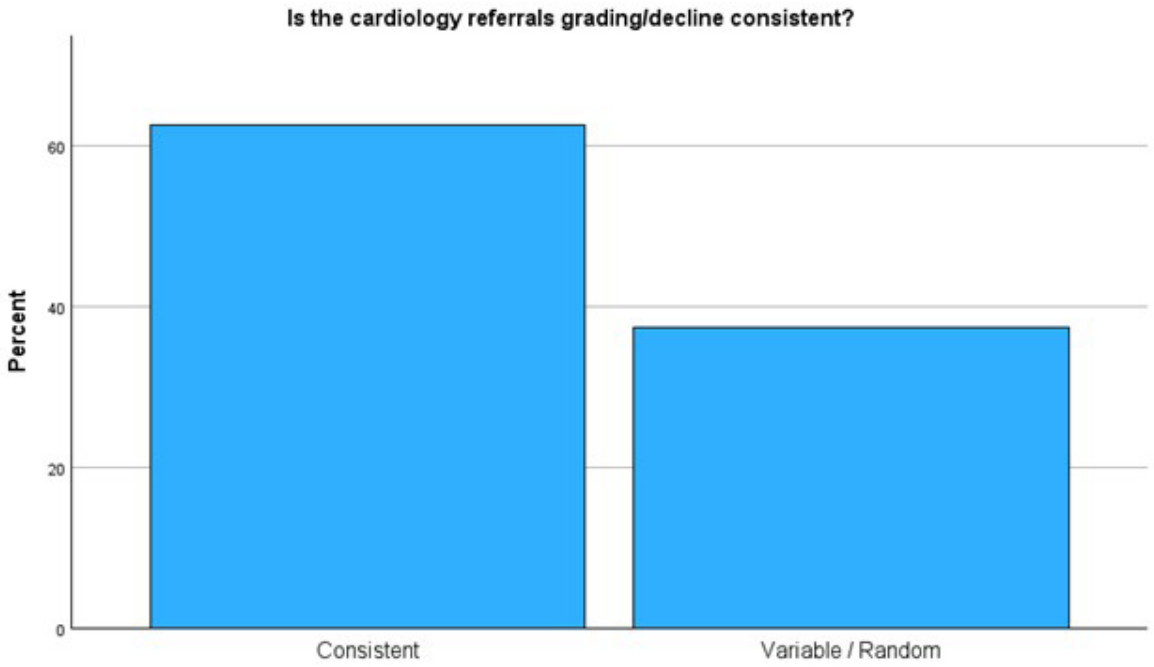

Inconsistent referral grading had been experienced by 38% of respondents.

16. Have you ever had a cardiology referral inappropriately declined? And does this happen often?

Inappropriately declined? Yes/no
Does this happen often? Yes/no

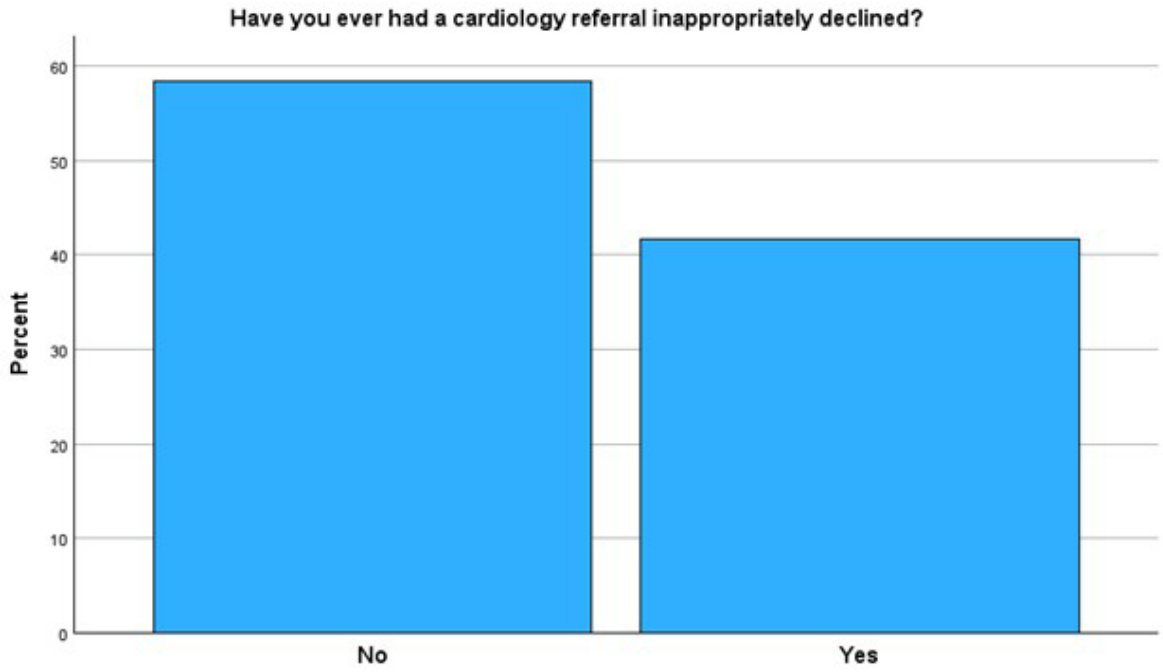

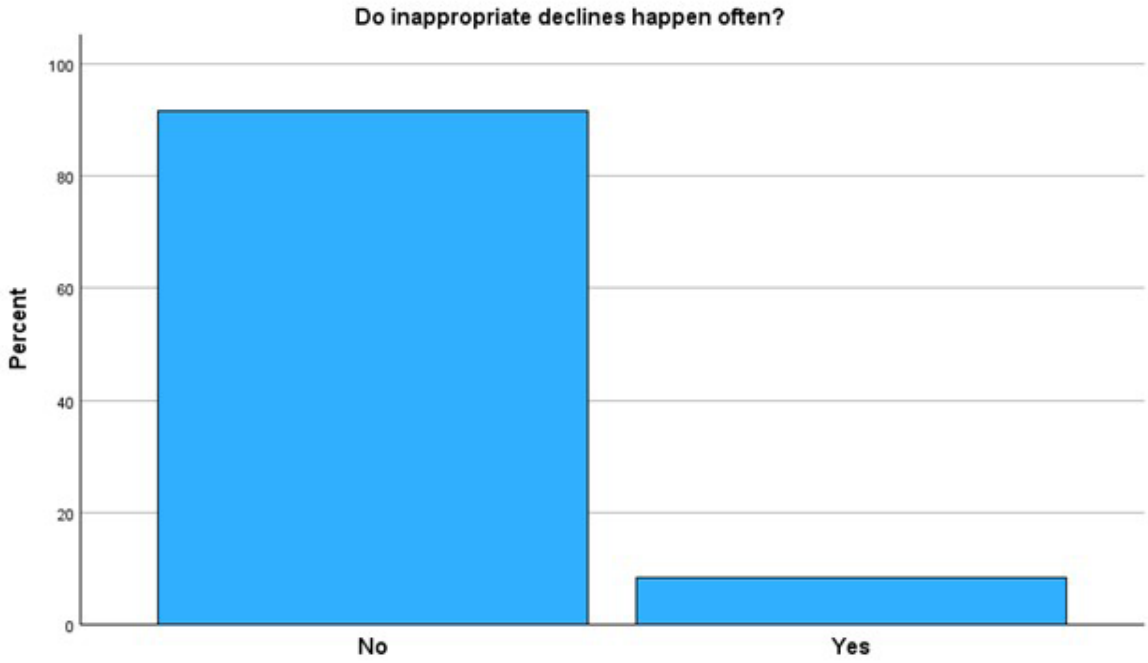

Less than half (42%) had a referral declined inappropriately, but this was infrequent (8%).

Discussion

It is a struggle for triaging cardiologists to keep up with the volume of referrals. The cardiologist shouldn't deal with those that can be handled clerically, and it would help to provide the GP with direct access to relevant information. A GP comment highlighted a perverse incentive in the current system. Sending a referral requesting a cardiologist to provide the result of a Holter monitor test (four mouse clicks) is easier than finding the report online (eight mouse clicks, including a log-in page). Streamlining the system to make it easier for GPs to find information themselves will benefit all parties. We understand that more information is being made accessible to GPs but note that GPs cannot always find available information. One reason is the so-called "friction" in the navigation process. Other reasons need to be explored.

An option to improve accessibility to GPs is to compile relevant links on a single dashboard page, like the Snapshot used for Health New Zealand – Te Whatu Ora Waitematā inpatients (Appendix 2 Figure 1).

Appendix 2 Figure 1: Example of the current Snapshot page for Waitematā inpatients.

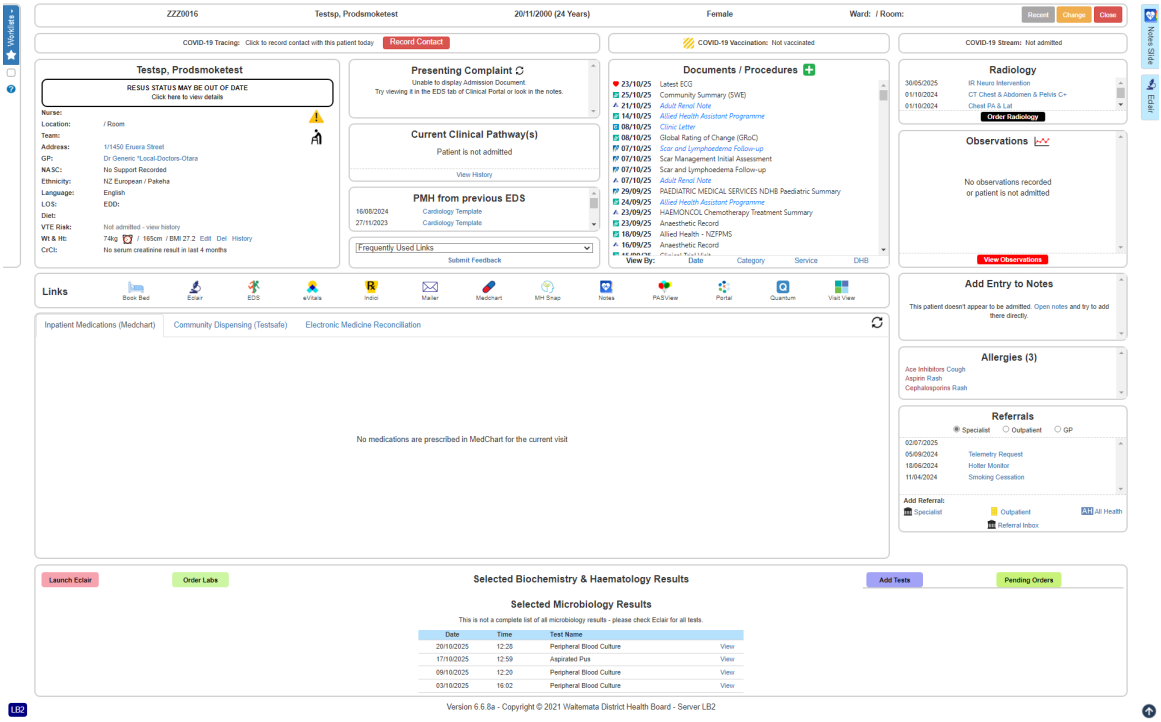

The aim is to reduce the number of declined referrals by providing GP referrers with more information and easier access to that information.

It is suggested that the dashboard have single-click links to the following:

- List of conditions seen by cardiology.
- List of conditions frequently referred to cardiology incorrectly, with links to the correct recipient.
- Current list of approximate wait times for clinic and investigations.
- Regional guidelines on follow-up intervals for surveillance of valvular lesions and dilated aortas (have now been posted).
- Patient-specific information:
  - Pending referrals and appointments.
  - Recent encounters (inpatient/outpatient).
  - Results of recent cardiac investigations.

Limitation

Only 24 GPs were surveyed. Greater numbers would give more confidence in the findings.

Conclusion

The GPs canvassed had positive views on aspects of the referral process but wanted more guidance on who and when to refer. They desired information on waiting times and better access to patient information held by the hospital. Options for improvement are discussed.

Appendix 3: audit of declined outpatient cardiology referrals

Introduction

Some declined referrals may be wasteful when they invoke fruitless time and effort by the referring general practitioner (GP) and triaging cardiologist. This can be a type of inefficiency (or "churn"). We catalogue reasons for referral rejection. Auckland Health Research Ethics Committee approved the project (AH28636).

Methods

Administrative coding proved inaccurate (which suggests caution about deploying artificial intelligence [AI] on such a data source). Between 3 January 2022 and 28 November 2023 (22 months), there were 3,505 referrals with a clerical categorisation of "declined". Reviewing individual patients' records, the total fell to 3,145, and these were manually categorised according to the reasons for referral and rejection.

Results

Clinical reason for referral

Palpitations comprised the largest volume of declined referrals (25%), followed by chest pain (10%) and valvular (10%) (Table 1).

Appendix 3 Table 1: Declined GP referrals; reason for referral.

Disease category	Frequency	Percent
Palpitations/brady/ECG	786	25
Other	646	21
Chest pain	321	10
Murmur/valvular	315	10
AF/flutter/SVT	244	8
Heart failure	241	8
Risk factor management	153	5
Hypertension	152	5
Breathlessness	139	4
Syncope/presyncope	96	3
Dilated aorta	28	0.9
Cardiomegaly	17	0.5
Not a GP referral	7	0.2

GP = general practitioner; ECG = electrocardiogram; AF = atrial fibrillation; SVT = supraventricular tachycardia.

During the time period, there was a total of 28,496 GP referrals. Only 49% were categorised, and we did not manually confirm the category of those that were. Roughly, the clerical reasons for referral were: chest pain 58%, palpitations 19%, breathlessness 8%, murmur 7% and heart failure 5%.

GP referral rates

Figure 1.

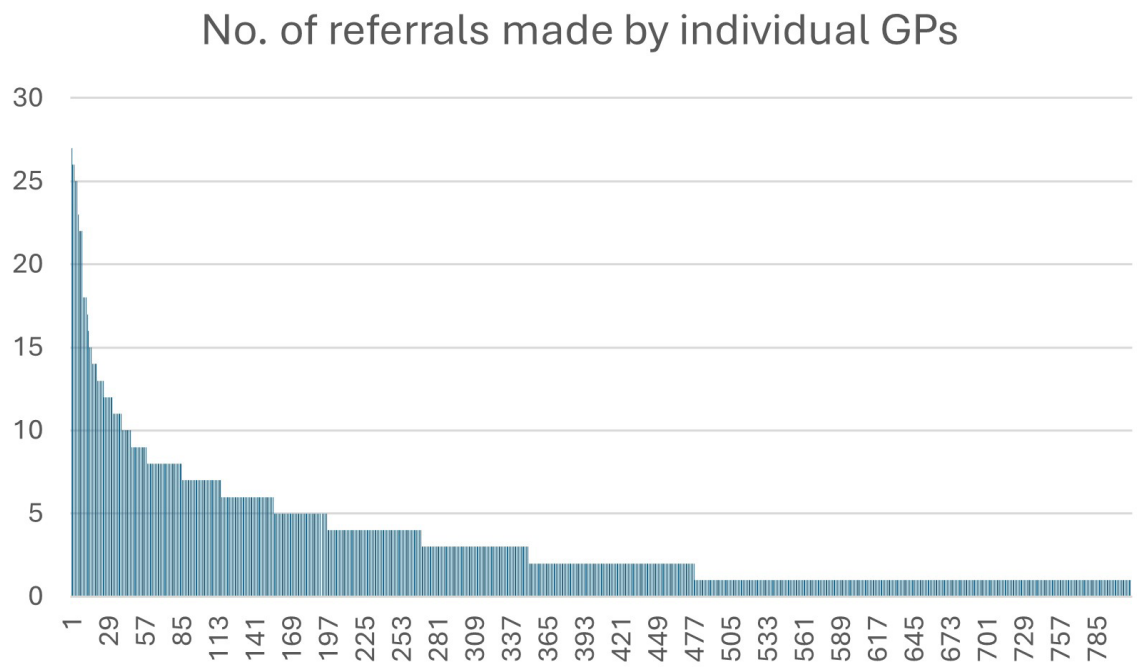

The number of rejected referrals made by individual GPs varied from 27 to one (Figure 1). The median was two, and the mode was one. Little can be inferred from these data without knowing how many patients each GP sees.

Reason for referral rejection

Appendix 3 Table 2: Declined GP referrals; reason for being declined.

Reason for decline	Frequency	Percent	Avoidable
Bespoke advice or interpretation	765	24	
Admin transfer to another service	555	18	Possibly
Relevant booking pending, no change to priority	394	13	Possibly
Generic advice given	297	9	Possibly
More information requested	248	8	Possibly

Appendix 3 Table 2 (continued): Declined GP referrals; reason for being declined.

Seen/admitted since referred	179	6	Possibly
Relevant results communicated	170	5	Possibly
Other	126	4	
Communication for filing	115	4	
Referral declined but procedure ordered	105	3	
Duplicate referral	60	2	Possibly
Referral not declined	33	1	
Advised to refer elsewhere, bespoke	32	1	
Follow-up advice as per regional valve, aorta guideline	33	1	Possibly
Relevant booking pending, priority increase requested	23	0.7	
Patient uncontactable	10	0.3	

The reasons for the referral being declined are in Table 2. The "more info requested" category is usually a request for an ECG. The commonest "bespoke advice" is for the duration of dual antiplatelet therapy, followed by lipid management, other medication, surveillance of aortic dilation, and anticoagulation.

"Admin transfer to another service" is the biggest category of possibly avoidable referrals. Many of these are palpitations transferred to general medicine, and hypertension transferred to the renal service. Such transfers don't relieve the overall demand for hospital services.

Avoidable referrals

We defined "possibly avoidable" referrals as those that would not need to be dealt with by a cardiologist if simple measures were instituted, such as: providing GPs with access to more information such as where to refer certain conditions, thresholds for referral of specific conditions, easily accessible results of investigations and the current appointment status of patient. We deemed 1,936 (62%) of declined referrals potentially avoidable (Figure 2). These included: "sent to wrong service" (24%), "already has appointment" (13%), "does not require review" (10%), "inadequate information" (8%), "seen since referred" (6%), "relevant results sent" (5%) and "duplicate referral" (2%). These are a drag on the efficiency of the triage process, increasing workload for clerical and clinical staff, with an ultimate adverse impact on patient care.

Figure 2.

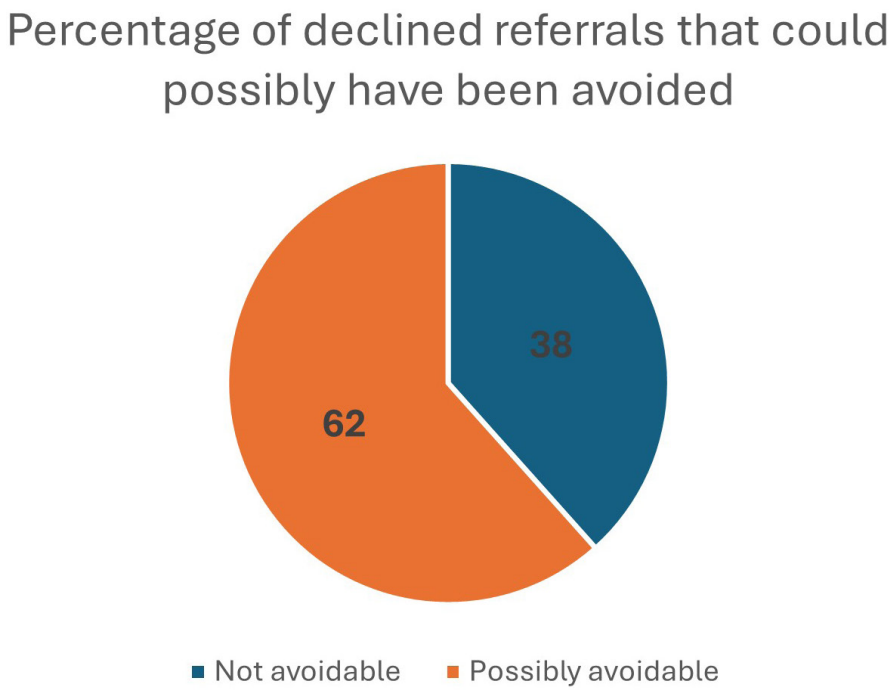

Patient outcomes after a declined referral

Twenty-four patients (0.8%) died within 90 days of a declined referral. Four patients died at −7, 1, 1 and 2 days after referral, suggesting that rejection of the outpatient referral could not have influenced the outcome. Review of the referrals for patients who subsequently died was not informative. The numbers are small and the information contained in the referrals was variable and often succinct, although not different from the referrals for patients who did not die. It is anticipated that changing from free-text referrals to the collection of structured information appropriate to the particular clinical indication will allow informative audit of adverse events. The learnings can then be used to iteratively refine the choice of variables collected and the decision thresholds for rejection.

Hospital admissions as adverse event were not assessed but should be part of the ongoing audit that iteratively improves future collection of structured data.

We did not manually assess re-referrals. The following data from the administrative coding provide ballpark estimates:

- 60% (1,245) of declined referrals do not get re-referred
- 2/3 (826) subsequent re-referrals are accepted

As an observation, if the accepted referrals had adequate information to be accepted the first time, that would be 413 fewer referrals to process on this basis alone. Even declining a referral requires opening several windows in the electronic medical record. Each has a lag time, which is worse during working hours. Therefore, getting the correct information upfront saves a worthwhile amount of time.

Cardiologist variability in rejection rate

Cardiologist decline rates varied widely from 14% to 47%. Variability is an adverse indicator of healthcare quality. Its reduction is the specific goal of recent national initiatives.[1–5]

Six doctors with more than 10 years of clinical cardiologist experience had a decline rate of 26%, and four cardiologists with less than 10 years of experience had a rate of 17% (Figure 4). These numbers are too small for statistical analysis.

The decline rate was unrelated to the cardiologists' triaging volume (Figure 3).

Figure 3.

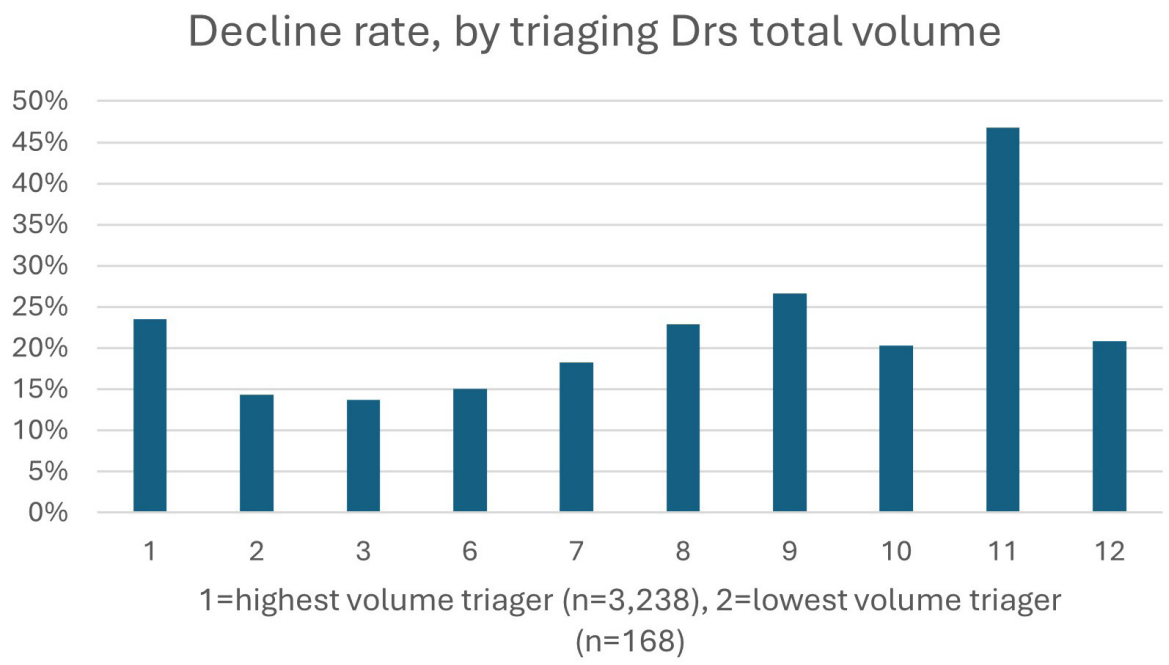

The decline rate is not related to the triaging cardiologists' overall volume.

Figure 4.

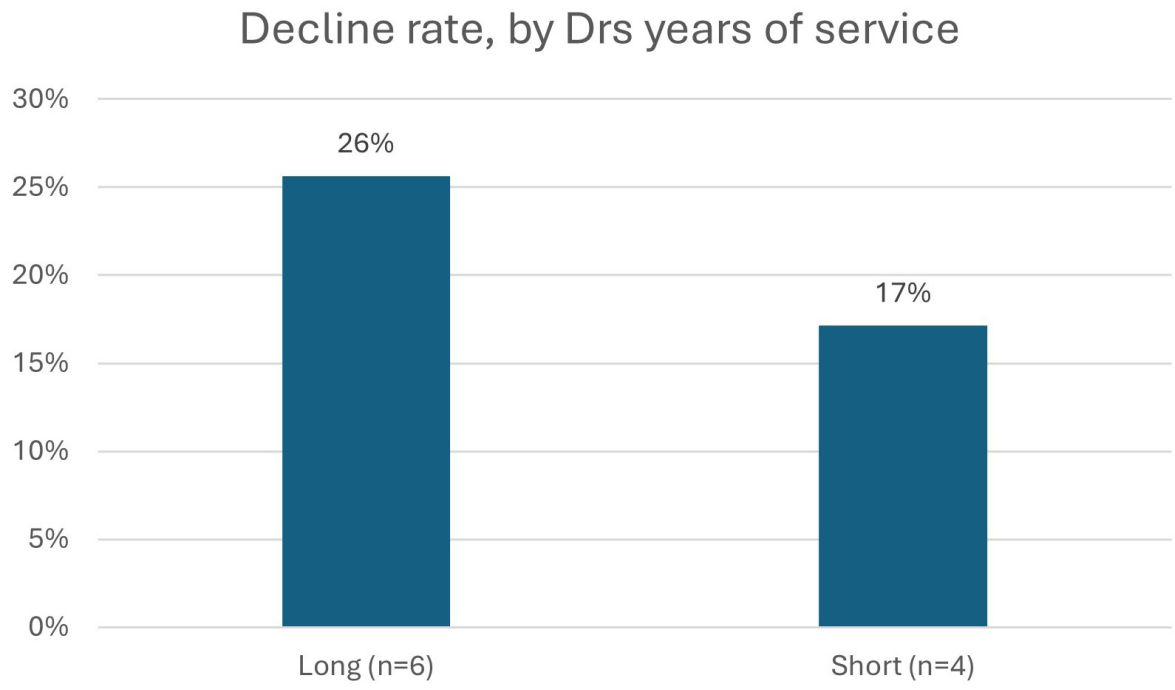

Under-resourcing causes more work via positive feedback

Anecdotal experience has always been that when patient delays increase, there is a concomitant increase in referrals solely due to the delays. The category "relevant booking already pending" comprises a fifth of potentially avoidable referrals and is due to delays in patient investigations. In addition, the category "other" contains a proportion of procedures that have yet to be reported. Similarly, the frequency of "request increased priority" is likely related to waiting times. These are examples of how busyness begets busyness, potentially leading to a spiral of increasing delays and inefficiencies.

Limitations

Some people with functional or non-cardiac symptoms (palpitations, chest pain) may benefit from a normal cardiology work-up and subsequent cardiologist reassurance. We acknowledge this, but the current body of work is driven by an inability to service demand with available resources. Accordingly, the consensus was to design systems that sift out organic heart disease. Being swamped with other patients will delay us seeing patients for whom we may reduce the risk of serious events.

Discussion

This study focussed on declined referrals for two reasons. From a behavioural aspect, removing unnecessary referrals gives the triaging cardiologist more time to devote to necessary referrals. More importantly, the analysis of declined referrals provides a window into the shortfalls and inefficiencies in the overall referral process.

The job of the triaging cardiologist could be facilitated by a "cardiologist triager's dashboard", similar to the Snapshot used for Health New Zealand – Te Whatu Ora Waitematā inpatients. It is a single page containing relevant single-click links to facilitate and speed up the cardiologist's task of triaging referrals.

Such a page should contain single-click links to:

- Pending referrals and appointments.
- Recent encounters (inpatient/outpatient) with emphasis on those that have occurred since the date of referral—it is not uncommon for the patient to have been admitted since the referral was made. At present, this information is not quickly available to the triager, resulting in time wasted.
- Results of recent investigations, including laboratory, radiology and cardiology.

Conclusion

Most declined referrals were possibly avoidable altogether, mainly by providing GPs with ready access to information. There is no signal of harm from a declined referral, and most declined referrals are not re-referred. The decline rates vary between triaging cardiologists, suggesting practice standardisation is needed. Suggestions are made to improve the process.

REFERENCES

1. Institute of Medicine (US) Committee on Quality of Health Care in America. Crossing the Quality Chasm: A New Health System for the 21st Century. Washington (DC): National Academies Press (US); 2001.
2. McGlynn EA, Brook RH. Keeping quality on the policy agenda. Health Aff (Millwood). 2001;20(3):82-90. doi: 10.1377/hlthaff.20.3.82.
3. The Commonwealth Fund. Perspective: Consistency, Continuity, and Coordination-The 3Cs of Seamless Patient Care [Internet]. [cited 2024 Mar 6]. Available from: https://www.commonwealthfund.org/publications/newsletter-article/perspective-consistency-continuity-and-coordination-3cs-seamless
4. Realistic Medicine. Reduce Unwarranted Variation [Internet]. [cited 2024 Mar 6]. Available from: https://realisticmedicine.scot/about/reduce-unwarranted-variation/
5. National Ethics Advisory Committee. National Ethical Standards for Health and Disability Research and Quality Improvement [Internet]. Wellington, New Zealand: Ministry of Health – Manatū Hauora; 2019 [cited 2025 Oct 28]. Available from: https://neac.health.govt.nz/assets/Uploads/NEAC/publications/national-ethical-standards-health-disability-research-quality-improvement-2019-v3.pdf

Appendix 4: draft decision trees for illustrative purposes

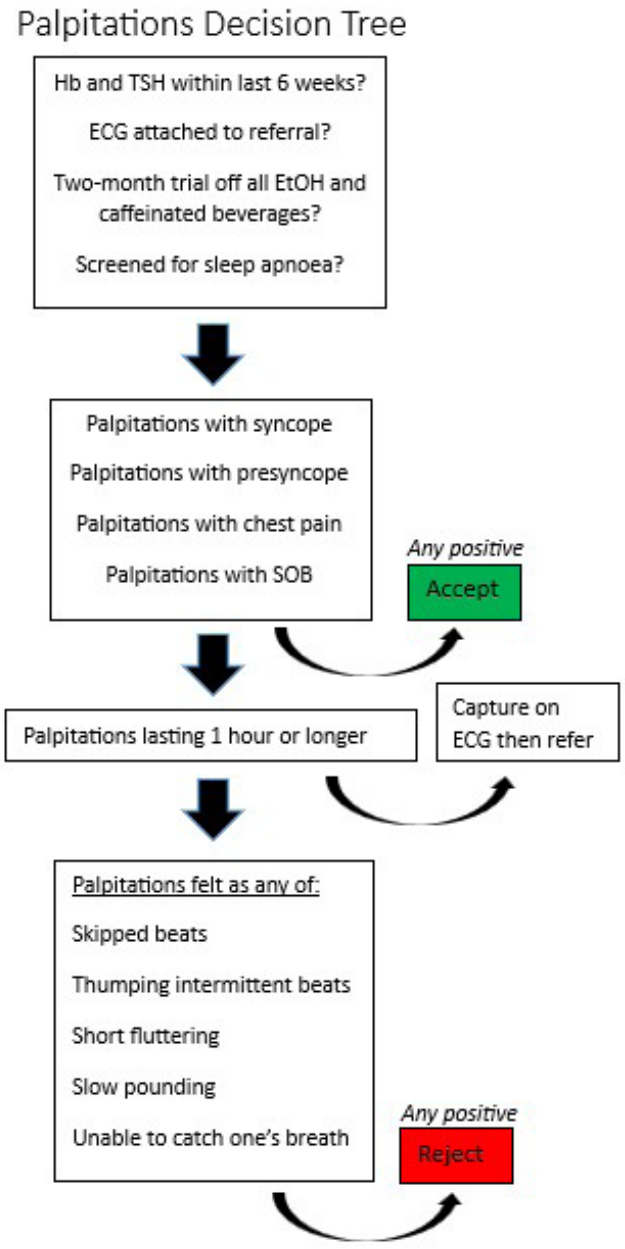

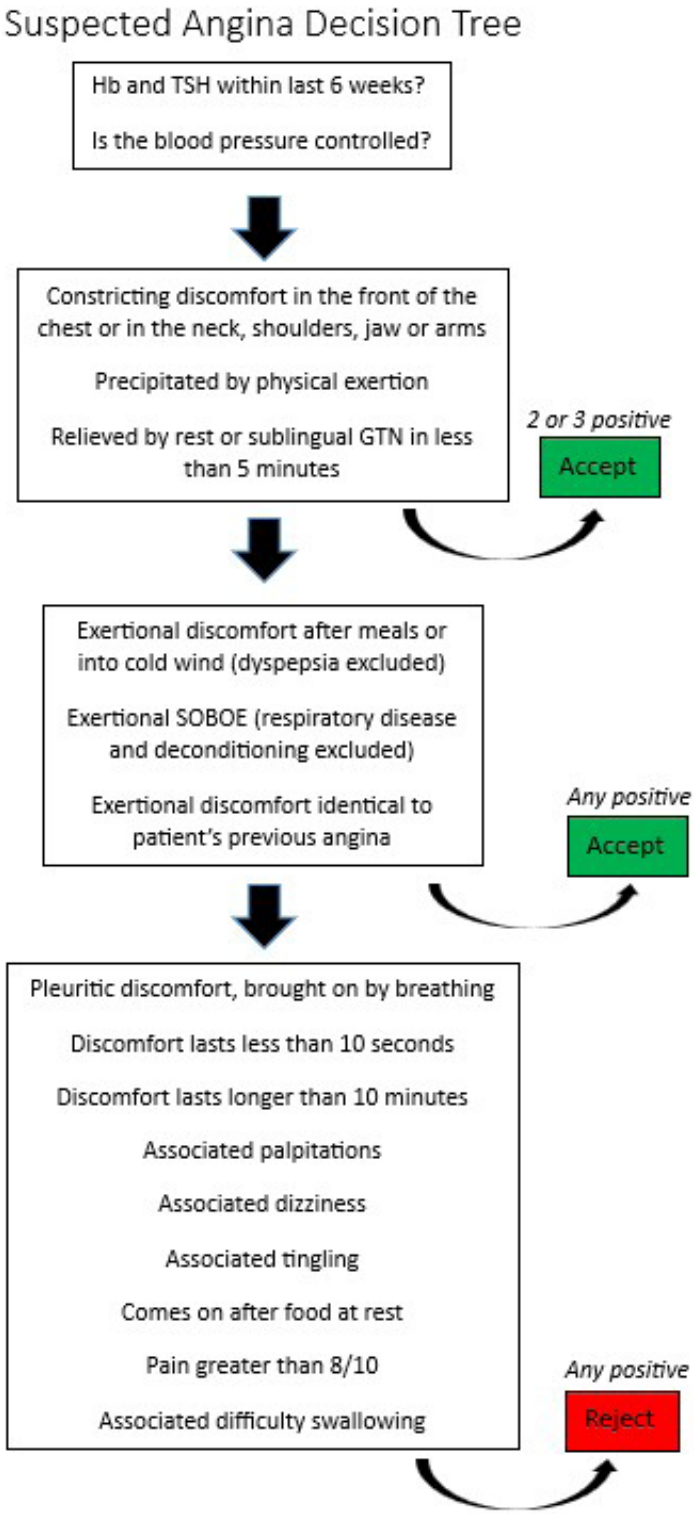
Suspected Angina Decision Tree
Hb and TSH within last 6 weeks?
Is the blood pressure controlled?
Constricting discomfort in the front of the chest or in the neck, shoulders, jaw or arms
Precipitated by physical exertion
Relieved by rest or sublingual GTN in less than 5 minutes
2 or 3 positive
Accept
Exertional discomfort after meals or into cold wind (dyspepsia excluded)
Exertional SOBOE (respiratory disease and deconditioning excluded)
Exertional discomfort identical to patient's previous angina
Any positive
Accept
Pleuritic discomfort, brought on by breathing
Discomfort lasts less than 10 seconds
Discomfort lasts longer than 10 minutes
Associated palpitations
Associated dizziness
Associated tingling
Comes on after food at rest
Pain greater than 8/10
Associated difficulty swallowing
Any positive
Reject

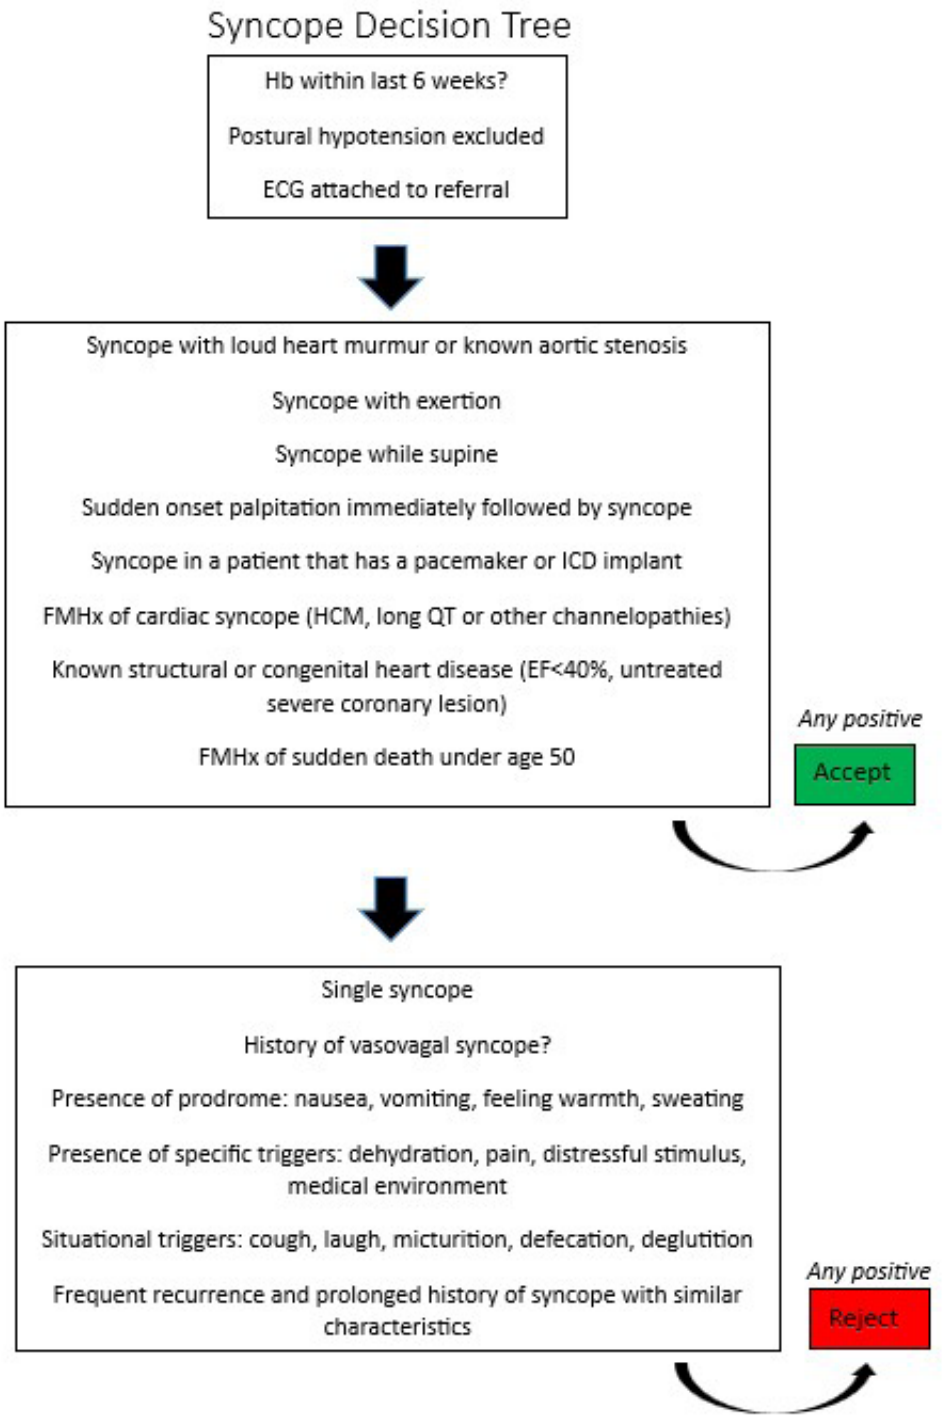
Syncope Decision Tree
Hb within last 6 weeks?
Postural hypotension excluded
ECG attached to referral
Syncope with loud heart murmur or known aortic stenosis
Syncope with exertion
Syncope while supine
Sudden onset palpitation immediately followed by syncope
Syncope in a patient that has a pacemaker or ICD implant
FMHx of cardiac syncope (HCM, long QT or other channelopathies)
Known structural or congenital heart disease (EF<40%, untreated severe coronary lesion)
FMHx of sudden death under age 50
Any positive
Accept
Single syncope
History of vasovagal syncope?
Presence of prodrome: nausea, vomiting, feeling warmth, sweating
Presence of specific triggers: dehydration, pain, distressful stimulus, medical environment
Situational triggers: cough, laugh, micturition, defecation, deglutition
Frequent recurrence and prolonged history of syncope with similar characteristics
Any positive
Reject

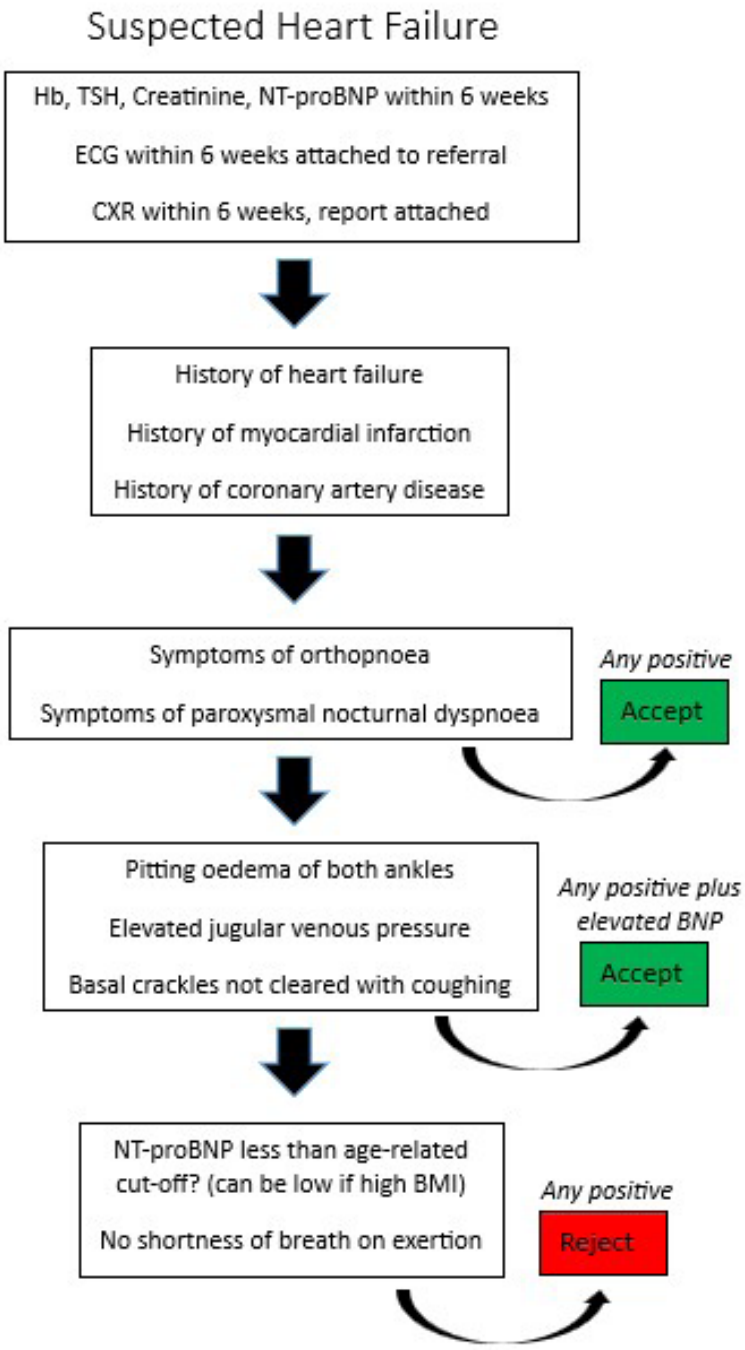
Suspected Heart Failure
Hb, TSH, Creatinine, NT-proBNP within 6 weeks
ECG within 6 weeks attached to referral
CXR within 6 weeks, report attached
History of heart failure
History of myocardial infarction
History of coronary artery disease
Symptoms of orthopnoea
Symptoms of paroxysmal nocturnal dyspnoea
Any positive
Accept
Pitting oedema of both ankles
Elevated jugular venous pressure
Basal crackles not cleared with coughing
Any positive plus elevated BNP
Accept
NT-proBNP less than age-related cut-off? (can be low if high BMI)
No shortness of breath on exertion
Any positive
Reject