This article was downloaded by: [216.73.216.198] On: 18 June 2026, At: 23:23
Publisher: Institute for Operations Research and the Management Sciences (INFORMS)
INFORMS is located in Maryland, USA

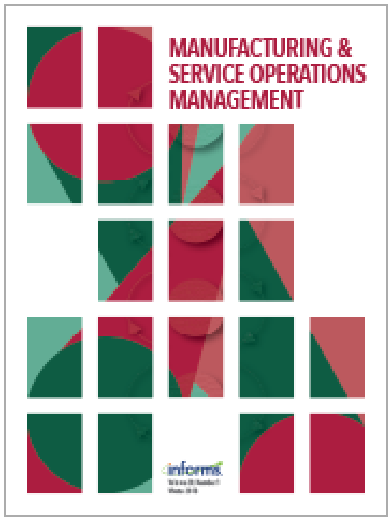

Manufacturing & Service Operations Management

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Frontiers in Operations: Optimal Genetic Testing of Families

Daniel Adelman, Kanix Wang

To cite this article:
Daniel Adelman, Kanix Wang (2024) Frontiers in Operations: Optimal Genetic Testing of Families. Manufacturing & Service Operations Management 26(4):1338-1357. https://doi.org/10.1287/msom.2023.0057

This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License. You are free to download this work and share with others for any purpose, except commercially, and you must attribute this work as "*Manufacturing & Service Operations Management*. Copyright © 2024 The Author(s). https://doi.org/10.1287/msom.2023.0057, used under a Creative Commons Attribution License: https://creativecommons.org/licenses/by-nc/4.0/."

Copyright © 2024 The Author(s)

Please scroll down for article—it is on subsequent pages

With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.
For more information on INFORMS, its publications, membership, or meetings visit http://www.informs.org

https://pubsonline.informs.org/journal/msom

MANUFACTURING & SERVICE OPERATIONS MANAGEMENT
Vol. 26, No. 4, July–August 2024, pp. 1338–1357
ISSN 1523-4614 (print), ISSN 1526-5498 (online)

Frontiers in Operations: Optimal Genetic Testing of Families

Daniel Adelman,[a] **Kanix Wang**[b,*]

[a] Booth School of Business, The University of Chicago, Chicago, Illinois 60637; [b] Lindner College of Business, University of Cincinnati, Cincinnati, Ohio 45221

*Corresponding author

Contact: dan.adelman@chicagobooth.edu, https://orcid.org/0000-0001-5534-4511 (DA); wang2xk@ucmail.uc.edu, https://orcid.org/0000-0003-1355-577X (KW)

Received: January 18, 2023
Revised: May 3, 2023; October 23, 2023; January 24, 2024
Accepted: February 18, 2024
Published Online in Articles in Advance: June 14, 2024

https://doi.org/10.1287/msom.2023.0057

Abstract. ***Problem definition*:** Through the laws of inheritance, knowing an individual's genetic status informs disease risk for family members, but current protocols for deciding whom to genetically test only consider one person at a time rather than design an optimal testing plan for the entire family. ***Methodology/results*:** We develop a Markov decision process framework for maximizing the net benefits of genetic testing that integrates a Bayesian network of genetic statuses, with a functional representation of cost-effectiveness. Our model provides a contingent sequence of family members to test one at a time, that is, a plan that dynamically incorporates new test results, revealed sequentially at random, to decide who next to test. In the general case, we show that optimal stopping follows a structure with two-sided thresholds, previously known only for individual testing. Although the optimal testing sequence, in general, is contingent on the family test results, in the special case of sibling-only tests we can identify this sequence a priori. Our numerical case study, which was conducted in a realistic BRCA1/2 testing setting, demonstrates that an optimal policy significantly improves cost-effectiveness over existing policies. Thus, our framework offers a promising and powerful new approach to genetic testing. ***Managerial implications*:** In an optimal policy, prioritizing testing family members who might otherwise not have been tested can lead to an overall improvement in familial health value, surpassing even the most cost-effective existing protocols. From a management perspective, healthcare organizations and insurance companies can potentially save costs by implementing this approach for such families.

History: This paper has been accepted in the *Manufacturing & Service Operations Management* Frontiers in Operations Initiative.

Funding: D. Adelman is grateful for financial support from Booth School of Business.
Supplemental Material: The online appendix is available at https://doi.org/10.1287/msom.2023.0057.

Keywords: dynamic programming • healthcare management • genetic testing

1. Introduction

The American College of Medical Genetics and Genomics (ACMG) lists 35 conditions with reliable genetic tests that provide medical benefits by allowing patients to prevent or better manage health conditions (Miller et al. 2021). A vast majority of them are transmitted in an autosomal dominant pattern in a family, that is, one mutation is enough to substantially increase disease risk for a family relative. Thus, discovering even one patient's genetic status can strongly inform the rest of the family and their physicians of the underlying risks and the appropriate courses of action, for example, cascade testing where relatives of carriers are recommended to get tested (Robbins 2019). With advances in precision medicine (Phillips et al. 2018, Bergin 2019), a growing consensus is that families, rather than individuals, should be the unit of care when considering genetic testing. Indeed, family genetic testing has been advocated by the American College of Cardiology to "identify and treat those individuals with unrecognized Familial Hypercholesterolemia" (Sturm et al. 2018, p. 677). A United Kingdom-based pilot study further reinforces this perspective, finding that larger families reap the largest per capita benefits after analyzing the genetic sequences of thousands of families (The 100000 Genomes Project Pilot Investigators 2021). Underscoring the urgency for a family-centered approach in genetic testing, a recent article in *The Wall Street Journal* resonates with these findings (Abbott 2023).

Currently, there are several competing guidelines for determining who should undergo genetic testing (Manahan et al. 2019, U.S. Preventive Services Task

Force et al. 2019, Daly et al. 2021). These guidelines are practical as they primarily focus on the individual patient presenting to the clinician. However, none explicitly take into account the informational value of family genetic tests, for example, by recommending to test someone largely because the test results may be valuable information for others in the same family.

To expand this limited reach, health organizations around the world are offering genetic testing beyond the few specific cohorts covered by current guidelines. One way to do that is through elective genetic testing, targeting patients and families not yet recommended by current guidelines (Robbins 2019). Other health organizations cast a wider net, offering free sequencing to all consenting patients, for example, Geisinger MyCode (Carey et al. 2016), University of California San Francisco Data, Discovery, and Diversity (UCSF 3D) (Weiler 2020), and Mayo Clinic Tapestry (Mayo Clinic 2021). However, neither approach is scalable and optimal under resource constraints and testing capacity limits.

Applying a family approach to genetic testing requires the study of all aspects of the dynamic system. These include disseminating genetic information to the rest of the family, assessing the cost-effectiveness under uncertainty, weighing the trade-off between the value of information and health benefits for the tested, and quantifying the future consequences of each test. The optimal strategy may be to test individuals who would otherwise not be tested. For example, one might want to test a grandfather for breast and ovarian cancer genes. In what family situation would this be optimal? How does one balance the informative value of grandfather testing with the health benefits of grandchild testing?

We contribute the first comprehensive modeling framework for optimal genetic testing of families. Our modeling framework takes a family pedigree as input and incorporates information about individual family members' risks. It constructs a contingent sequence of whom to test next in the family by testing one member at a time and using the test result to decide the subsequent action (i.e., to test someone else or to stop testing altogether). It does this to maximize the net reward minus costs, which captures the aggregated social welfare of the family. In this paper, we take the reward to be the expected net quality-adjusted life years (QALYs). Although the QALY concept has been criticized (Rand and Kesselheim 2021), our framework does not depend on using QALYs. Another concept of the reward will do; however, the existent literature on the cost-effectiveness of a test predominantly interprets rewards in QALYs (Center for the Evaluation of Value and Risk in Health 2020).

Our framework integrates three previously unconnected strands of research and makes the following key contributions.

- We formulate a Markov decision process (MDP) that incorporates the pedigree structure of a family with a state space that captures the Mendelian laws of inheritance (Mendel 1866). To our knowledge, our model is the first model to do so. This model is important because in an optimal policy, testing priority may be given to persons who might otherwise not have been tested. This is because their test results may be of more value to the rest of the family in terms of information gain than to the individual being tested.
- A two-sided threshold policy, where individuals are tested when their risk falls between a lower threshold and an upper threshold, is known to be optimal for individual testing (Drakopoulos and Randhawa 2021, Mills and Ziya 2021). However, little is known about the structure of optimal dynamic testing for multiple related individuals. To the best of our knowledge, we are the first to prove that the two-sided threshold policy is also optimal for dynamic sequential testing in families under Bayesian updating. This structural result suggests a simple, practical way to implement optimal family testing. Clinically, it is standard to provide patients with laboratory reports in the form of a simple and intuitive chart, similar to the one shown in Figure 1. The chart uses green color to show risk ranges for optimal testing and includes an individual range marked by a solid line to indicate when testing will be beneficial for the tested individual independently from the family. When the risk falls within the region where the green band and solid line overlap, a test will benefit both the tested individual and the rest of the family. When the risk falls within the green band but outside the solid line, a test will provide mostly informational value for the rest of the family.

Although the optimal testing sequence, in general, is contingent on the family test results, in the special case of sibling testing, we further identify a priori the optimal sequence for testing. In the general case, our model can be viewed as a type of stochastic traveling salesman problem with dynamic decisions (see Toriello et al. 2014), and therefore, the optimal sequence to test is unlikely to exhibit a simple structured optimal policy.

- We base our risk calculations on a Bayesian network of genetic status, a probabilistic graphical model used in most clinical software for genetic diseases (Chen et al. 2004, Terry et al. 2019). Our model updates the posterior distribution of genetic statuses across the family as the test results come in. This kind of Bayesian network has rarely been incorporated into Markov decision processes previously.
- We contribute a functional representation to assess the net QALYs.[1] This is critical to our framework because with it, we can evaluate decisions to stop testing additional family members. In addition, it only uses information readily available in the published medical literature about genetic outcomes, and thus,

Figure 1. (Color online) A Mock-up Chart of Family Testing Results

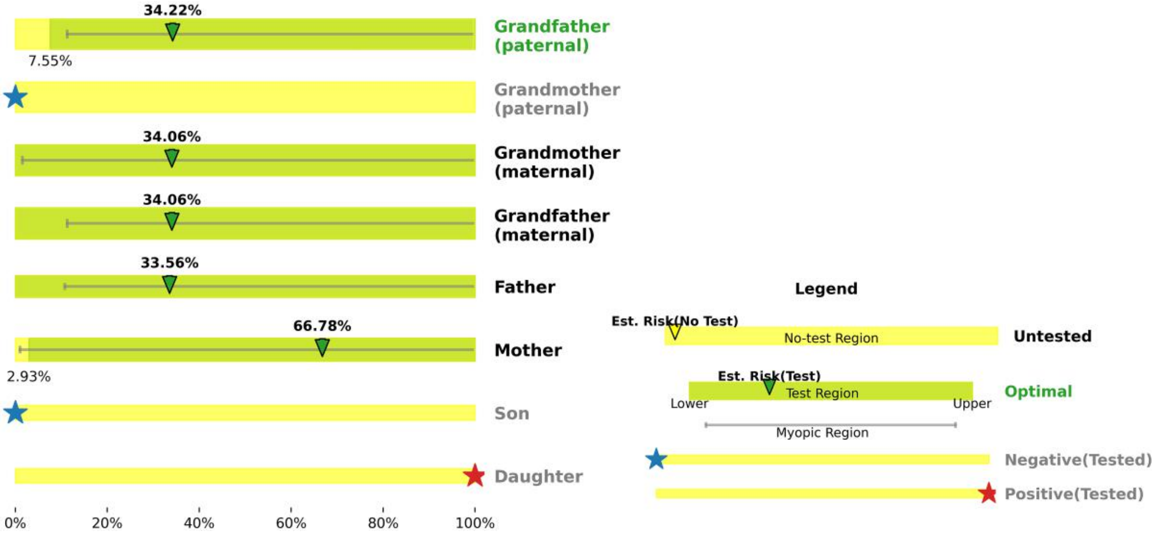

our model allows for this literature to be translated into a decision-analytic framework for family testing protocols, a critical building block for precision medicine. We validate our BRCA1/2 results against Nohdurft et al. (2017) (see Online Appendix C.1), which pioneered the QALY-maximizing age-dependent preventive strategies for BRCA1/2-positive cases (carriers). This means that our reward function implicitly captures the recourse decision of treatments over time. Similarly, the reward function could also implicitly capture the recourse decision to test later in life, but our model only decides whether to test in the here and now (i.e., the current family encounter).

- To demonstrate the need for an optimal framework, even when sophisticated existing protocols are in place, we conduct several numerical studies. We compare our framework with two competing policies: (1) the independent policy, which represents the best current practice, where each individual's net benefits are maximized independently and the results are not shared and (2) a myopic updating policy, which is similar to the independent policy but the results are shared.

Our numerical study shows that an optimal policy can outperform these existing policies arbitrarily. We construct an example where both existing policies fail to test anyone, regardless of the number of siblings, whereas the optimal policy tests to the benefit of the family. An optimal policy might test individuals not recommended by current guidelines or those not intuitively considered for testing.

To further demonstrate these dynamics, we consider three families with contrasting characteristics—a small, high-risk nuclear family; a small, low-risk nuclear family; and a much larger family with complex relationships and diverse risk profiles. In the latter case, we show that an optimal policy has a much wider range of tested individuals than current policies, reflecting the contingent nature of an optimal policy on whether and whom to continue testing. To account for heterogeneity in model parameters, such as family members' ages and testing costs, we also perform sensitivity and robustness analyses.

For realism, we adopt a censored reward function to model the fact that payers would require or a patient may demand a positive genetic test result for major surgeries (e.g., prophylactic mastectomy). Censoring increases the upper threshold; hence, the two-sided stopping policy is more likely to simplify to a one-sided policy.

We demonstrate that an optimal policy can bring in a larger total value for the family than even the most cost-effective existing protocols. However, we defer conclusions about the broader national or population-level impact to a future study. Such inferences are challenging because genetic testing utilization and effectiveness may be affected by cultural and societal factors, and the data currently available on family structures and ages are limited.

1.1. Brief Literature Review

Beyond Nohdurft et al. (2017), Drakopoulos and Randhawa (2021), Mills and Ziya (2021), and the other works referenced, there is not a large literature on medical decision making that directly relates to our work. Our framework designs a sequence of testing decisions by testing one member at a time. An alternative method, known as group testing, involves pooling all members together in a single test tube (El Hajj et al. 2022). Group testing complements sequential testing. For example, group testing can be used as a Bayesian update on prior probabilities before initiating sequential testing. Our optimal value function could serve as the basis for deciding whether to initiate pool testing for a family. If group testing returns a positive result, our framework could help determine how to sequentially test this family, potentially making family testing even more cost effective.

There is a growing literature on MDP models for clinical use, especially in precision medicine. Other examples of MDP models in healthcare include Alagoz et al. (2004), which explored the optimal timing for a living-donor liver transplant; Chhatwal et al. (2010), which used an MDP model to gain insight into how patient demographics influence the optimal breast cancer biopsy decision; and Örmeci et al. (2016), which used dynamic programming to establish guidelines for managing the trade-off between urgent diagnostic needs and nonurgent screening needs. More recently, Lee et al. (2022) identified the optimal intervention policy for postoperative care based on readmission risks and treatment effectiveness. Kong and Zhang (2018) covers a wide range of current decision models. Our MDP model is another example of how such models can aid healthcare professionals in making dynamic, evidence-based decisions.

1.2. Outline

The rest of the paper is structured as follows. Section 2 details the system and the Markov decision process. Section 3 presents the Bayesian network to update risks. Section 4 develops cost and reward functions for cost-effective testing. Section 5 provides structural results on the optimal policy. Section 6 numerically compares the optimal policy's performance with existing policies on BRCA1/2 testing. Section 7 concludes with our main contributions and directions for future research.

2. Model Formulation

In this section, we introduce a Markov decision process for family genetic testing.

2.1. Conceptual Overview

Before delving into the mathematical details of our model, first consider a nuclear family of four—parents and two children—planning to undergo testing for genes A and B: for instance, BRCA1 and BRCA2 genes associated with breast and ovarian cancers. Each family member has known attributes, like age, sex, and medical history. Unbeknownst to them and thus, to the decision maker, the mother and the daughter both carry a mutation in gene A, the father carries two mutations in gene B, and the son carries mutations in both. None of them have undergone prior tests for these genes. Our goal is to develop a model that will construct an optimal sequential testing policy, which dynamically chooses whether to test and if so, whom to test next, given test results seen thus far.

A sensible model should meet several criteria. First, it should represent family relationships, known as a *pedigree*, linking individuals to their parents and siblings. Second, it should have reward/cost functions for testing decisions based on each person's characteristics and genetic risk.

We must also represent each person's actual, although unknown, genetic status, denoted by letter $\mathcal{g}$. Given Mendelian inheritance laws, certain genetic combinations are inadmissible, for example, mutation-free parents will have mutation-free children. Although we do not know the true genetic statuses, we nonetheless can construct a prior distribution based on all available family information updated using Bayes' rule as new test results come in. This problem does not decompose by individual, but reward/cost functions do after individual risks are calculated from the global information. The reward function should reflect the lifetime impact of each possible genetic outcome weighted by its marginal likelihood given all family information available. The cost function should account for both the initial test and any follow-up test needed should a test return positive.

We assume that testing occurs over a short period of time such that the only information changing is the test results. Although we will assume perfect tests[2] (i.e., there is no residual error), only a subset of patients would have undergone testing at any given time. It is vital to track who has been tested and the mutation counts for each gene—denoted by the letter **s**—and then, share these results with family and their doctors. This will suffice to represent our state space, which evolves with incoming results. Clearly, the possible states **s** will depend on the feasible genetic statuses $\mathcal{g}$.

Finally, we need to distinguish between the reward for a person who is selected for testing versus the reward for a person who is not selected when the problem terminates. Those tested learn their true genetic status and can follow medical guidelines. Those untested update their terminal genetic risks based on final observed results and might consider prophylactic treatments, which may or may not require testing. Thus, we need a different reward function for these patients.

In subsequent subsections, we will reference the simple nuclear family as we define notation.

2.2. System Description and Notation

2.2.1. Preliminaries. Let $\mathcal{I}$ denote a family, $\mathcal{J}$ denote a set of genes,[3] and $g_{ij} \in \{0, 1, 2\}$ denote the true genetic status of gene $j \in \mathcal{J}$ for individual $i \in \mathcal{I}$. $\Gamma = \{0, 1, 2\}$ indicates the number of pathogenic mutation copies (i.e., possible genetic testing results). The family genetic matrix is $\mathcal{g} = \{\mathcal{g}_i\}_{\forall i \in \mathcal{I}}$, with $\mathcal{g}_i$ being the $|\mathcal{J}|$-vector for individual i. Mendelian inheritance laws determine which genetic configurations are biologically feasible. For expository clarity, we defer the details of these familial relationships to Section 3 and Online

Appendix D.1. $\mathbb{P}$ represents the probabilities related to genetic statuses, $\mathscr{G}$ represents the sample space for $\boldsymbol{g}$, and $\sigma(\mathscr{G})$ represents its associated σ-algebra. The random vector $\boldsymbol{G}$ is defined on the probability space $(\mathscr{G}, \sigma(\mathscr{G}), \mathbb{P})$. For autosomal dominant genes, health outcomes are determined by the presence of a mutation, not its count. Thus, we introduce the binary carrier status indicator $\theta_{ij} = 0 \Longleftrightarrow g_{ij} = 0$ and 1 otherwise using Θ_{ij} as the random variable. Random variables are typically represented by uppercase letters except for τ, which stands for the random variable for stopping time.

2.2.2. State. Each individual is tested for all genes in J or none. Let $s_{j\gamma} \subseteq \mathcal{I}$ represent individuals with the same test result $\gamma \in \Gamma$ for gene j. Define $\mathbf{s}_j = \{s_{j0}, s_{j1}, s_{j2}\}$ as a partition of tested individuals in $\mathcal{I}$. The state $\mathbf{s}$ is the collection of $\mathbf{s}_j$ for all j represented as $\mathbf{s} = \{\mathbf{s}_1, \ldots, \mathbf{s}_{|\mathcal{J}|}\}$, encapsulating test results for all genes. If individual i is tested, individual i belongs to $\cup_{\gamma \in \Gamma} s_{j\gamma}$ for every $j \in \mathcal{J}$, denoted as $i \in \mathbf{s}$. If not tested, it is written as $i \in \mathcal{I} \setminus \mathbf{s}$. The state random variable is denoted by $\mathbf{S}$, and the next state is $\mathbf{S}'$, with $\mathbf{s}'$ being its realization.

Given the results for the tested individuals, there is a mapping between state $\mathbf{s}$ and the set of feasible $\boldsymbol{g}$: for some $\gamma \in \Gamma, i \in \mathbf{s}_{j\gamma} \Rightarrow g_{ij} = \gamma \;\; \forall \; i \in \mathbf{s}, \; \forall \; j \in \mathcal{J}$. The feasible person-gene matrices corresponding to the state $\mathbf{s}$ are represented as

$$\mathscr{G}(\mathbf{s}) \equiv \{\boldsymbol{g} \in \mathscr{G} : \forall i \in \mathbf{s} \;\; \forall \, j \in \mathcal{J} \; g_{ij} = \gamma \text{ if } i \in s_{j\gamma}\}.$$

The state space is defined as $\mathscr{S} = \{\mathbf{s} : \exists \boldsymbol{g} \in \mathscr{G}(\mathbf{s})\}$. The admissible next states from $\mathbf{s}$ are $\mathscr{S}'(\mathbf{s}, a) \subseteq \mathscr{S}$, with a being the action at state $\mathbf{s}$, to be defined shortly.

Definition 1 (Genetic Risk). We denote the vector $\mathbb{P}_i(\mathbf{s}) = [\mathbb{P}_{i1}(\mathbf{s}) \; \mathbb{P}_{i2}(\mathbf{s}) \; \cdots \; \mathbb{P}_{i|\mathcal{J}|}(\mathbf{s})]$, where the jth element of this vector is the posterior probability that individual i is carrying at least one pathogenic mutation at gene j given $\mathbf{s} \in \mathscr{S}$ as

$$\mathbb{P}_{ij}(\mathbf{s}) \equiv \mathbb{P}(\Theta_{ij} = 1 \,|\, \mathbf{S} = \mathbf{s}) = \mathbb{P}(G_{ij} \in \{1, 2\} \,|\, \mathbf{S} = \mathbf{s}).$$

We use $\mathbb{P}$ to connote that this ties to the probability of $\boldsymbol{G}$. There is no loss of generality by maintaining the posterior distribution of Θ instead of $\boldsymbol{G}$ because knowing $\mathbf{s}$ and the distribution of Θ, one can simply compute the distribution of $\boldsymbol{G}$. We use $\mathbb{E}_i$ to connote the expectation for individual i with respect to $\mathbf{s}$ (i.e., conditioning on i). Under Bayes' law, $\mathbb{E}_i[\mathbb{P}_i(\mathbf{S}') \,|\, \mathbf{s}] = \mathbb{P}_i(\mathbf{s})$.

2.2.3. Action. At each decision epoch, after observing the current state $\mathbf{s}$, the decision maker either selects another untested member in $\mathcal{I} \setminus \mathbf{s}$ for testing or stops testing for the entire family. The symbol $\varnothing$ signifies stopping the test, transitioning to the absorbing state Δ, and obtaining the terminal reward. The symbol C denotes the action of continuing in the absorbing state. For each state $\mathbf{s} \in \mathscr{S} \cup \{\Delta\}$, there is a finite set of available actions $a \in \mathscr{A}_{\mathbf{s}}$,

$$\mathscr{A}_{\mathbf{s}} = \begin{cases} \mathcal{I} \setminus \mathbf{s} \cup \{\varnothing\} & \mathbf{s} \in \mathscr{S} \\ \{C\} & \mathbf{s} = \Delta, \end{cases}$$

where $\mathcal{I} \setminus \mathbf{s}$ means selecting among the untested individuals.

2.2.4. Transition. If the decision maker chooses individual $i \in \mathcal{I} \setminus \mathbf{s}$, the genetic test will reveal $g_{ij} \; \forall j \in \mathcal{J}$, and $\mathbf{s}$ changes accordingly. For every $j \in \mathcal{J}$, if $g_{ij} = \gamma \in \Gamma$, then $s'_{j\gamma} \leftarrow s_{j\gamma} \cup \{i\}$ with probability $\mathbb{P}(\boldsymbol{G} \in \mathscr{G}(\mathbf{s}') \,|\, \boldsymbol{G} \in \mathscr{G}(\mathbf{s})) = \prod_{j \in \mathcal{J}} \mathbb{P}(G_{ij} = g_{ij} \,|\, \mathbf{S} = \mathbf{s})$.

State transition probabilities $p_a(\mathbf{s}, \mathbf{s}')$ represent for each pair of states $(\mathbf{s}, \mathbf{s}')$ and each action $a \in \mathscr{A}_{\mathbf{s}}$ the probability that the next state will be $\mathbf{s}'$ given the current state $\mathbf{s}$ and the current action $a \in \mathscr{A}_{\mathbf{s}}$. In our setting, we have

$$p_a(\mathbf{s}, \mathbf{s}') = \begin{cases} \mathbb{P}(\boldsymbol{G} \in \mathscr{G}(\mathbf{s}') \,|\, \boldsymbol{G} \in \mathscr{G}(\mathbf{s})) & \\ & \text{if } a \in \mathcal{I} \setminus \mathbf{s}, \; \mathbf{s} \in \mathscr{S}, \;\; \mathbf{s}' \in \mathscr{S}'(\mathbf{s}, a) \\ 1 & \text{if } a = \varnothing, \; \mathbf{s} \in \mathscr{S}, \; \mathbf{s}' = \Delta, \text{ or } a = C, \\ & \;\; \mathbf{s} = \mathbf{s}' = \Delta, \\ 0 & \text{o.w.,} \end{cases} \tag{1}$$

where $\mathbb{P}(\boldsymbol{G} \in \mathscr{G}(\mathbf{s}') \,|\, \boldsymbol{G} \in \mathscr{G}(\mathbf{s}))$ is the transition probability when test continues. When testing concludes, the current state transitions to the terminal absorbing state Δ and remains there.

2.2.5. Reward and Cost. Our approach posits that family members' net QALYs are influenced by their beliefs about their genetic status, which we assume are rational, taking into account all available information. Prior to testing, beliefs of higher risk may prompt individuals to adopt more intensive preventive measures. After undergoing genetic tests, their beliefs are updated to either zero or one for each gene examined, that is, $\mathbb{P}_{ij}(\mathbf{S}') \in \{0, 1\} \;\; \forall i \in \mathbf{S}' \;\; \forall j \in \mathcal{J}$, with $\mathbf{S}'$ being the next random state. Consequently, the rational beliefs of untested family members are updated as well.

At each decision epoch, the decision maker evaluates the current state $\mathbf{s}$ and has two options: test an untested family member i or stop testing altogether. Let $c_i(\mathbf{s})$ represent the expected cost associated with testing family member i in state $\mathbf{s}$. Let $r_i(\boldsymbol{p})$ be the expected net QALYs for a tested family member i when the rational belief vector on whether i carries a pathogenic variant at each genes is $\boldsymbol{p} = [p_1, p_2, \ldots, p_{|\mathcal{J}|}] \in [0, 1]^{|\mathcal{J}|}$. Testing incurs an expected cost $c_i(\mathbf{s})$ and yields an expected net reward of $\mathbb{E}_i[r_i(\mathbb{P}_i(\mathbf{S}')) - c_i(\mathbf{s}) \,|\, \mathbf{s}]$. The system then transitions to the next state $\mathbf{s}'$ with probability $p_i(\mathbf{s}, \mathbf{s}')$. Consequently, the rational beliefs of untested family members are updated as well. After updating beliefs, the decision maker may opt to test another individual. Alternatively, if the decision maker chooses to cease testing, the system transitions to the absorbing state Δ, ending the decision-making process.

Let $\overline{r}_i(p)$ represent the corresponding reward when someone is not tested, which may differ from $r_i(p)$ because of some preventative measures being unavailable. To secure insurance coverage for certain procedures or to ensure patient confidence in undergoing invasive treatments, such as prophylactic mastectomy, a conclusive proof of genetic risk is mandatory. In such a case, $\overline{r}_i(p)$ guarantees that each patient who opts for a major surgery receives a definitive genetic test. The remaining untested family members' net QALYs are calculated based on their rational beliefs in the current state $\mathbf{s}'$ (i.e., $\sum_{i\in\mathcal{I}\setminus\mathbf{s}'}\overline{r}_i(\mathbb{P}_i(\mathbf{s}'))$). The absorbing state Δ involves neither reward nor cost. In summary, we have

$$\text{Net Reward per Period} \\ = \begin{cases} \mathbb{E}_i[r_a(\mathbb{P}_a(\mathbf{S}')) - c_a(\mathbf{s})|\mathbf{s}] & \text{if } a \in \mathcal{I}\setminus\mathbf{s},\ \mathbf{s}\in\mathscr{S} \\ \sum\limits_{i\in\mathcal{I}\setminus\mathbf{s}} \overline{r}_i(\mathbb{P}_i(\mathbf{s})) & \text{if } a = \varnothing,\ \mathbf{s}\in\mathscr{S} \\ 0 & \text{if } a = C,\ \mathbf{s} = \Delta, \end{cases} \tag{2}$$

which depends on the state $\mathbf{s}$ and action a taken.

To compare reward with cost, we convert cost to QALYs using a threshold willingness-to-pay (WTP) for health (e.g., in dollars per QALY). Thus, $c_i(\cdot)$, $r_i(\cdot)$, and $\overline{r}_i(\cdot)$ may depend on i's characteristics, such as age, medical/family history, and comorbidities. For more information on the cost-effective analysis with a WTP threshold, please refer to Online Appendix D.2.

2.3. Family Example

Let us consider the aforementioned family of four. We have $\mathcal{I} = \{1,2,3,4\}$, where 1 indicates the mother, 2 indicates the father, 3 indicates the daughter and 4 indicates the son, and two genes $\mathcal{J} = \{1,2\}$. The family genetic matrix $\boldsymbol{g}$ and carrier matrix $\boldsymbol{\theta}$ look like these:

$$\boldsymbol{g} = \begin{pmatrix} 1 & 0 \\ 0 & 2 \\ 1 & 0 \\ 1 & 1 \end{pmatrix}\ \boldsymbol{\theta} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 0 \\ 1 & 1 \end{pmatrix}.$$

For gene $j = 1$, the mother has one pathogenic mutation, the father has none, and both children have one. For $j = 2$, the mother and the daughter have none, the father has two mutations, and the son has one. The carrier matrix shows each member's heightened disease risk: the mother's from $j = 1$, the father's from $j = 2$, the daughter's from $j = 1$, and the son's from both.

For the mother ($i = 1$) testing positive for one mutation in gene $j = 1$ and negative for $j = 2$, the state is represented as $\mathbf{s} = \{\mathbf{s}_1 = \{s_{10} = \{\}, s_{11} = \{1\}, s_{12} = \{\}\}, \mathbf{s}_2 = \{s_{20} = \{1\}, s_{21} = \{\}, s_{22} = \{\}\}\}$, with $\{\}$ denoting an empty set. In this state, actions could be testing $i \in \{2, 3, 4\}$ or could be ending testing.

2.4. Markov Decision Process

We now formally construct a Markov decision process. Denote the set of admissible pairs $(\mathbf{s}, a)$ by $\mathbb{K}$, where $\mathbf{s} \in \mathscr{S} \cup \Delta$ and $a \in \mathscr{A}_{\mathbf{s}}$. The process begins at an initial state $\mathbf{s}^0 \in \mathscr{S}$, and during each period n, there can be at most $|\mathcal{I}\setminus\mathbf{s}^n|$ family members available for testing.

To track the process history, we define the space of admissible histories up to period n as H^n. Specifically, we have $H^1 = \mathscr{S}$ and $H^n = \mathbb{K}^{n-1} \times \mathscr{S} \cup \Delta$. An element h^n of H^n is called an admissible n-history, and it is a vector of the form $(\mathbf{s}^0, a^0, \mathbf{s}^1, a^1, \ldots, \mathbf{s}^n)$. We can also write $h^n = (h^{n-1}, a^{n-1}, \mathbf{s}^n)$ for $n = 1, \ldots, |\mathcal{I}\setminus\mathbf{s}^0|$. We have $\mathbf{s}^n = \Delta$ $\forall\, n \geq |\mathcal{I}\setminus\mathbf{s}^0| + 1$. The random stopping time, τ, is defined as the minimum period such that $\mathbf{s}^{n+1} = \Delta$. Because the last period with nonzero reward in history h^n is τ, we refer to $\mathbf{S}^\tau$, with realization $\mathbf{s}^{\tau'}$, as the terminal state.

The sample space, Ω, consists of the set of admissible histories, $H^{|\mathcal{I}\setminus\mathbf{s}^0|}$, where each sample path, ω, is a sequence of states and actions, such as $\mathbf{s}^0, a^0, \mathbf{s}^1, a^1, \ldots, \mathbf{s}^{|\mathcal{I}\setminus\mathbf{s}^0|}$. The initial state, $\mathbf{s}^0$, is given, and the stopping time, τ, is the period in which testing stops. We define the random variables $\mathbf{S}^n(\omega) = \mathbf{s}^n$, $A^n(\omega) = a^n$ for all n.

We also define a sequence of joint distributions $\{\mu^n\}_{\forall n=0,1,\ldots,\tau}$, where μ^0 is the initial prior distribution of the genetic status of family members $\boldsymbol{G}$ and μ^n is the joint distribution of the genetic status at time n. For any genetic status matrix $\boldsymbol{g} \in \mathscr{G}$ and any state $\mathbf{s}^n$, we use $\mu^n(\boldsymbol{g}|\mathbf{s}^n) = \mathbb{P}(\boldsymbol{G} = \boldsymbol{g}|\mathbf{S}^n = \mathbf{s}^n)$ to denote the probability that $\boldsymbol{G}$ equals $\boldsymbol{g}$ at time n given state $\mathbf{s}^n$. The joint distributions are dynamically updated based on a Bayesian network of genetic status. This network is integral to the MDP's accurate modeling of genetic testing scenarios.

With stationary, positive, bounded rewards and finite state and action spaces, there exists an optimal policy that is Markovian, stationary, and deterministic (Puterman 1994, theorem 4.4.2c). We denote this set of policies as Π. Given that the state of the system evolves without cycles, a policy can be expressed as a single deterministic decision rule. A deterministic decision rule, $\pi(\mathbf{s}) \in \mathscr{A}_{\mathbf{s}}$ $\forall\, \mathbf{s} \in \mathscr{S} \cup \{\Delta\}$, induces a probability measure $P^{\pi}_{\mathbf{s}^0}$ on the $(\Omega, \mathscr{F})$, where $\mathscr{F}$ is the σ-algebra of subsets of Ω such that $P^{\pi}_{\mathbf{s}^0}(\mathbf{S}^n = \mathbf{s}^n | h^{n-1}, a^{n-1}) := p_{a^{n-1}}(\mathbf{s}^{n-1}, \mathbf{s}^n)$, $a^{n-1} = \pi(\mathbf{s}^{n-1})$.

For every initial state $\mathbf{s}^0$, the problem is to find a policy π that maximizes the expected net reward: that is,

$$V^*(\mathbf{s}^0) = \max_{\pi\in\Pi} \mathbf{E}^{\pi}_{\mathbf{s}^0}\left[\sum_{n=0}^{\tau-1}[r_{\pi(\mathbf{S}^n)}(\mathbb{P}_{\pi(\mathbf{S}^n)}(\mathbf{S}^{n+1})) - c_{\pi(\mathbf{S}^n)}(\mathbf{S}^n)] \right. \\ \left. + \sum_{i\in\mathcal{I}\setminus\mathbf{S}^\tau} \overline{r}_i(\mathbb{P}_i(\mathbf{S}^\tau))\right], \tag{3}$$

Figure 2. A Sample Family Pedigree of Eight, Genetic Status Update After g_2 Tested Negative, and a Family Bayesian Network Representation

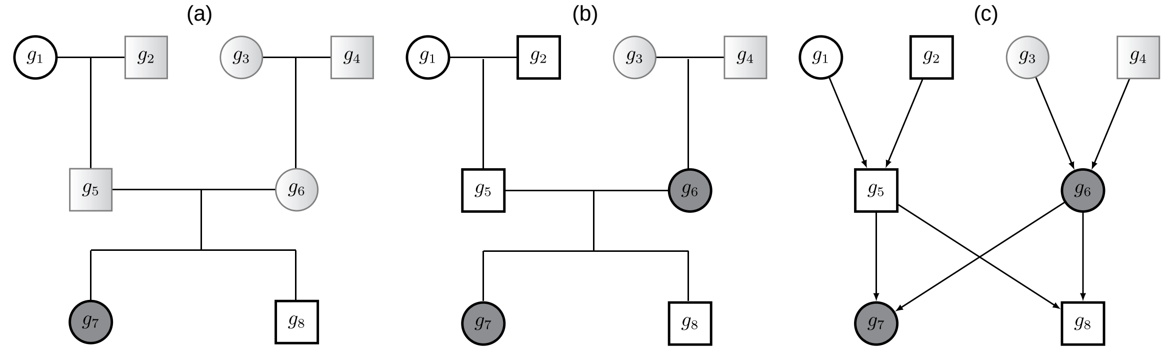

Notes. □ indicates males, and ○ indicates females. Dark gray indicates mutation presence, white indicates mutation absence, and shaded gray indicates status unknown. (a) A sample family pedigree of eight. (b) Genetic status after g_2 tested negative. (c) Family Bayesian network.

where the expectation operator $\mathbf{E}^{\pi}_{\mathbf{s}^0}$ is over all $\omega \in \Omega$ such that the initial state is $\mathbf{s}^0$ and generated by the decision rule π. Given that the probability of obtaining a set of test results is sequence independent (Lemma EC.1 in the Online Appendix), accruing rewards at each epoch is equivalent to collecting all rewards at the end of the horizon (Lemma EC.2 in the Online Appendix). Furthermore, the optimal value function for the MDP, $V^*(\mathbf{s})$, can be found by solving the following equation:

$$V^*(\mathbf{s}) = \max\left\{ \max_{i\in\mathcal{I}\setminus\mathbf{s}} \mathbb{E}_i[r_i(\mathbb{P}_i(\mathbf{S}')) - c_i(\mathbf{s}) + V^*(\mathbf{S}')\,|\,\mathbf{s}],\ \sum_{i\in\mathcal{I}\setminus\mathbf{s}} \bar{r}_i(\mathbb{P}_i(\mathbf{s}))\right\}. \tag{4}$$

The optimal policy in state $\mathbf{s}$ is obtained by taking an action and then, either selecting an untested individual i or ceasing testing, which solves the right-hand side (RHS) of the optimality equation (4). We use the standard backward induction to solve (4).

3. Probabilistic Representation of Family Test Results

Here, we develop a Bayesian network to probabilistically represent a family's genetic test results for the MDP.[4] In addition to capturing the family genetics, this Bayesian network dynamically updates as new results emerge. Notably, the updated marginal probabilities not only reflect the individual risks but also serve as the transition probabilities for the dynamic program (4). We show an efficient method for generating these probabilities and implementing the Bayesian dynamic program. For expository purposes, we will focus on the single-gene setting in this section, i.e., $|\mathcal{J}| = 1$.

3.1. Prior Distribution and Bayesian Network

In biological and medical contexts, to demonstrate the familial relationships between members, families are often depicted as pedigrees (Figure 2). Within these pedigrees, individuals at the top, lacking specified parents, are termed *founders* and are denoted by $\mathcal{F}$, whereas others are labeled *nonfounders*. We use ♂_i and ♀_i as the indices for the father and mother of person i, respectively.

Because there exists a Bayesian network that captures the independencies among the family members, that is, a child's test result is independent of the rest of the family given the parents' results, the joint distribution of the family test results can be factored into a product form (Koller and Friedman 2009, theorem 3.1):

$$\mathbb{P}(G = g) = \prod_{i\in\mathcal{F}} \mathbb{P}(G_i = g_i) \prod_{i\in\mathcal{I}\setminus\mathcal{F}} \mathbb{P}(G_i = g_i\,|\,g_{\text{♂}_i}, g_{\text{♀}_i}). \tag{5}$$

In our setting, $\mathbb{P}(G = g)$ follows the Mendelian laws and incorporates all prior information on the family members where $\mathbb{P}(G_i = g_i)\ \forall\, i \in \mathcal{F}$. Hence, $\mathbb{P}(G_i = g_i)$ are the priors given to us for the untested founders, and $\mathbb{P}(G_i = g_i\,|\,g_{\text{♂}_i}, g_{\text{♀}_i})$ are the conditional probabilities for untested nonfounders given their parents (see Table 1). Here, we assume there is a 50% chance of passing one of the two variants of any gene from a parent to a child.

Furthermore, the prior and posterior distributions on genetic status also depend on information about individuals. Thus, if person i has already been diagnosed with cancer, for instance, this would be incorporated in the joint probability on genetic status for this person and across all family members as well as in the expected reward and cost of testing person i.

3.2. Family Example

In Figure 2, the top four grandparents are the founders of this family (i.e., $\mathcal{F} = \{1, 2, 3, 4\}$). For parental relationships, consider $\text{♀}_5 = 1$, $\text{♂}_5 = 2$, $\text{♀}_6 = 3$, and $\text{♂}_6 = 4$. Currently, 7 has been identified as a carrier of one diseased

Table 1. $\mathbb{P}(G_i = g_i | g_{\male_i}, g_{\female_i})$

$g_{\male_i}, g_{\female_i}$	0, 0	0, 1	0, 2	1, 0	1, 1	1, 2	2, 0	2, 1	2, 2
$\mathbb{P}(G_i = 0 \| g_{\male_i}, g_{\female_i})$	1	0.5		0.5	0.25				
$\mathbb{P}(G_i = 1 \| g_{\male_i}, g_{\female_i})$		0.5	1	0.5	0.5	0.5	1	0.5	
$\mathbb{P}(G_i = 2 \| g_{\male_i}, g_{\female_i})$					0.25	0.5		0.5	1

copy, and 1 and 8 have tested negative, that is, $\mathbf{s} = \{s_0 = \{1, 8\}, s_1 = \{7\}, s_2 = \{\}\}$. This is the exact same family as our previous mock-up example in Figure 1.

3.3. Updating Probabilities with Test Results

Central to the MDP's functionality is the probabilistic transition from state $\mathbf{s}^n$ to $\mathbf{s}^{n+1}$ and the update of genetic risks with new test results. Both are achieved through the Bayesian network. The transition probability from state $\mathbf{s}^n$ to $\mathbf{s}^{n+1}$ is represented as the marginal probability $\mathbb{P}(G_i = g_i | \mathbf{s}^n)$ when individual i is selected in epoch n. We calculate this marginal probability by summing the joint probability $\mu^n(g | \mathbf{s}^n)$ over all possible test results of the family that have $G_i = g_i$:

$$\mathbb{P}(G_i = g' | \mathbf{s}^n) = \sum_{\{g \in \mathcal{G}(\mathbf{s}^n) : g_i = g'\}} \mu^n(g | \mathbf{s}^n) \ \forall\, g' \in \{0, 1, 2\}. \quad (6)$$

After the transition, we update the posterior joint distributions from $\mu^n(\cdot | \mathbf{s}^n)$ to $\mu^{n+1}(\cdot | \mathbf{s}^{n+1})$. The posterior joint distribution $\mu^{n+1}(g | \mathbf{s}^{n+1})$ is given by

$$\mu^{n+1}(g | \mathbf{s}^{n+1}) = \frac{\mathbb{P}(G \in \mathcal{G}(\mathbf{s}^{n+1}) | G = g)\mathbb{P}(G = g)}{\mathbb{P}(G \in \mathcal{G}(\mathbf{s}^{n+1}))} = \begin{cases} \dfrac{\mathbb{P}(G = g)}{\mathbb{P}(G \in \mathcal{G}(\mathbf{s}^{n+1}))} & \text{if } g \in \mathcal{G}(\mathbf{s}^{n+1}) \\ 0 & \text{if } g \notin \mathcal{G}(\mathbf{s}^{n+1}) \end{cases} \quad \forall\, n. \quad (7)$$

The expression $\mathbb{P}(G \in \mathcal{G}(\mathbf{s}^{n+1})) \equiv \mathbb{P}(\mathbf{S}^{n+1} = \mathbf{s}^{n+1})$ denotes the fixed probability of arriving at state $\mathbf{s}^{n+1}$ at the $n+1$th epoch. $\mathbb{P}(G = g)$ is the prior distribution with no testing performed. As detailed in Section 2, G and $\mathcal{G}(\mathbf{s}^n)$ are defined on the probability space $(\mathcal{G}, \sigma(\mathcal{G}), \mathbb{P})$. Because $\mathbb{P}(G \in \mathcal{G}(\mathbf{s}^{n+1}))$ is a computable constant, determining the posterior joint distribution $\mu^{n+1}(g | \mathbf{s}^{n+1})$ equates to reducing the dimension in the posterior distribution corresponding to the recently tested individual i. However, updating the marginal posterior for each family member based on test results is equivalent to the precise inference of the corresponding Bayesian network, a process known to be non-deterministic polynomial-time hard (NP-hard) in general (Cooper 1990).

3.4. Bayesian Dynamic Programming Implementation

Because of the complexity of directly updating the joint distribution, several algorithms have been proposed to efficiently compute marginal distributions (Pearl 1982, Kschischang et al. 2001, Koller and Friedman 2009). We use a tailored version of these algorithms, specifically BRCAPRO, for our Bayesian network because our MDP only requires marginal distributions for state transitions (see Equation (6)). The Bayesian network algorithm passes messages between nodes in the network, ensuring convergence of posterior marginals in two passes with a complexity of $O(2C)$. Here, $C = 3|\mathcal{I}| - 2|\mathcal{F}|$, with three message passings for each parent-child relationship and one for each founder. Thus, we store $3|\mathcal{I}|$ marginals instead of the $3^{|\mathcal{I}|}$ posterior joint distribution. To improve efficiency, we precompute all marginal probabilities, store them as a state transition matrix in a lookup table, and separate them from value function calculations. We solve the dynamic program using backward induction. For $|\mathcal{I}| = N$, a set of N unrelated individuals gives the upper bound on the state space, $(|\Gamma| \cdot |\mathcal{J}|)^N$. To evaluate and determine an optimal policy, the standard backward induction necessitates $\sum_k^N k \cdot (|\Gamma| \cdot |\mathcal{J}|)^{N-k} \cdot (2|\Gamma| \cdot |\mathcal{J}|)$ computations (Puterman 1994, theorem 4.5.1). Therefore, the overall computation complexity is $O(|\Gamma|^N |\mathcal{J}|^N)$. For our numerical study ($|\Gamma| = 2, |\mathcal{J}| = 2$) on a single-core Apple M1 processor, the computation durations were slightly above 1 minute for an 11-member family, 5 minutes for a 12-member family, and 21 minutes for a 13-member family.

4. Cost and Reward Functions

In this section, we define the cost function structure, assume a reward function for rational decision making, and illustrate how to construct a reward function using data from the literature.

4.1. Cost Functions

ACMG standards advise that positive pathogenic results be validated with Sanger sequencing or another method after initial sequencing testing (Rehm et al. 2013). They also strongly emphasize patient counseling after testing (Richards et al. 2015). Thus, costs should include both screening and follow-up.

Definition 2. The expected cost for testing individual i in state $\mathbf{s}$ is

$$c_i(\mathbf{s}) = c_i^{\mathrm{s}} + \left(1 - \prod_{j\in\mathcal{J}}(1 - \mathbb{P}_{ij}(\mathbf{s}))\right)c_i^{\mathrm{f}}, \tag{8}$$

where c_i^{s} is the sequencing screening cost, c_i^{f} is the validation and counseling cost, and $1 - \prod_{j\in\mathcal{J}}(1 - \mathbb{P}_{ij}(\mathbf{s}))$ is the probability that at least one gene returns positive.

4.2. Reward Functions

Prior cost-effectiveness analyses have only considered extreme beliefs of zero or one. We model health benefits from rational behaviors based on individual risk. Our reward function r_i is a function of individual covariates captured by the index i and an individual's rational belief vector $\mathbb{P}_i(\mathbf{s})$. Assuming that testing is preferred when free of cost, monetary or otherwise, we exclude irrational behaviors, like preference for less information.

To construct the reward function, we begin by considering what an oracle would know. Suppose an oracle draws a sample $\boldsymbol{\theta} = \{\boldsymbol{\theta_i}\}_{\forall i\in\mathcal{I}}$ from random variables Θ with $\boldsymbol{G} \sim \mathbb{P}(\boldsymbol{G} = g)$. For each individual i, we define the reward based on their true carrier status $\boldsymbol{\theta_i}$ as $r_i^{\boldsymbol{\theta_i}}$. For the boundary values of $\mathbb{P}_i(\mathbf{s})$, namely zero and one, we turn to data from the literature that quantifies health outcomes and treatment costs following positive or negative test results. For every gene, we identify four data points: taking no action after testing positive, taking no action after testing negative, pursuing interventions after testing positive, and pursuing interventions after testing negative.

We will make several assumptions not necessary for the framework but for interpretability and clinical relevance. One such assumption is that health behaviors are based on rational beliefs.

Assumption 1 (Rational Behaviors). *Assuming rational behaviors, we have*

$$\mathbb{P}_{ij}(\mathbf{s}) = 0 \Longleftrightarrow \theta_{ij} = 0,\ \mathbb{P}_{ij}(\mathbf{s}) = 1 \Longleftrightarrow \theta_{ij} = 1 \ \ \forall\ j \in \mathcal{J}.$$

When the rational belief vector $\mathbb{P}_i(\mathbf{s})$ aligns with true status vector $\boldsymbol{\theta_i}$ (i.e., under Assumption 1), we get $r_i^{\theta_i}(\boldsymbol{\theta_i})$, occurring only for extreme beliefs (zero or one). For $i \notin \mathbf{s}$, $r_i^{\boldsymbol{\theta_i}}(p)$ represents the reward in QALYs if i's rational belief vector is p and i's true status vector is $\boldsymbol{\theta_i}$. We focus on modeling rational behaviors but relaxing this assumption allows for computation of irrational ones.

To maintain the medical ethics of nonmaleficence (do no harm) and to exclude terminal patients, such as those with do not resuscitate orders, we assume nonnegative reward.

Assumption 2 (Nonmaleficence). *Assuming nonnegative reward, we have*

$$r_i^{\boldsymbol{\theta_i}}(\boldsymbol{p}) \ge 0 \ \ \forall\ i\ \ \forall\ \boldsymbol{\theta_i}\ \ \forall\ \boldsymbol{p} \in [0,1]^{|\mathcal{J}|}.$$

We also assume that a rational belief closer to the true status cannot reduce a person's net QALYs.

Assumption 3 (Truth Heals). *Assuming rational belief closer to the truth can not decrease net QALYs, we have*

$$r_i^{\boldsymbol{\theta_i}}(\boldsymbol{p}) \ \le\ r^{\boldsymbol{\theta_i}}(\boldsymbol{p'})\ \ \forall\ \boldsymbol{\theta_i}\ \text{if}\ |\boldsymbol{p} - \boldsymbol{\theta_i}| \ge |\boldsymbol{p'} - \boldsymbol{\theta_i}|.$$

Here, $|\boldsymbol{p} - \boldsymbol{\theta_i}|$ are the absolute values of $\boldsymbol{p} - \boldsymbol{\theta_i}$ elements, and $\ge$ is the component-wise inequality between vectors. That is, $\boldsymbol{p} \ge \boldsymbol{p'}$ means $\boldsymbol{p_j} \ge \boldsymbol{p'_j}$ for every $j \in \mathcal{J}$.

Recall that significant preventative measures, like prophylactic mastectomy, necessitate definitive positive test results, often enforced through genetic testing prerequisites set by payers or patients. For those untested, their net health benefits are capped, granting them only a censored reward, $\bar{r}_i^{\boldsymbol{\theta_i}}$, as opposed to the true $r_i^{\boldsymbol{\theta_i}}$.

Assumption 4 (Censored Reward). *Assuming a single-gene setting, we have*

$$r_i^1(0) \le \bar{r}_i^1(1)\ <\ r_i^1(1) - c_f - c_s \tag{9}$$

$$r_i^0(1)\ <\ \bar{r}_i^0(1)\ \le\ r_i^0(0). \tag{10}$$

In the absence of conclusive genetic results, limited preventative measures yield fewer benefits for a true positive, $\bar{r}_i^1(1)$, compared with post-test options. We define the impact of censoring on the health of a true carrier as $d_i = r_i^1(1) - \bar{r}_i^1(1)$. However, these measures are better than inaction (9). For true negatives, the health gains from prevention are negligible. Censoring helps circumvent the costly treatments. Nevertheless, nonintervention is best for true negatives (10). The term $r_i^\delta = r_i^0(0) + \bar{r}_i^1(1) - \bar{r}_i^0(1) - r_i^1(0)$ represents the total benefit of the censored measures on i. It balances the increase in QALYs for true positive cases against the increased expenses in the true negative cases.

Definition 3. In state $\mathbf{s}$, the expected reward for an tested individual $i \in \mathcal{I} \setminus \mathbf{s}$ is

$$\mathbb{E}_i[r_i(\mathbb{P}_i(\mathbf{S'}))|\mathbf{s}] = \sum_{\boldsymbol{\theta_i}\in\{0,1\}^{|\mathcal{J}|}} r_i^{\boldsymbol{\theta_i}}(\boldsymbol{\theta_i})\prod_{j\in\mathcal{J}}\mathbb{P}_{ij}(\mathbf{s})^{\theta_{ij}}(1-\mathbb{P}_{ij}(\mathbf{s}))^{1-\theta_{ij}}.$$

Here, Θ_i is a random true status vector, and $\mathbb{P}_i(\mathbf{s})$ is the rational belief vector on i's probabilities given the evidence in state $\mathbf{s}$.

Definition 4. In state $\mathbf{s}$, the expected reward for an untested individual $i \in \mathcal{I} \setminus \mathbf{s}$ is

$$\begin{aligned}\bar{r}_i(\mathbb{P}_i(\mathbf{s})) &:= \mathbb{E}_i[\bar{r}_i^{\Theta_i}(\mathbb{P}_i(\mathbf{s}))|\mathbf{s}] \\ &= \sum_{\boldsymbol{\theta_i}\in\{0,1\}^{|\mathcal{J}|}} \bar{r}_i^{\boldsymbol{\theta_i}}(\mathbb{P}_i(\mathbf{s}))\prod_{j\in\mathcal{J}}\mathbb{P}_{ij}(\mathbf{s})^{\theta_{ij}}(1-\mathbb{P}_{ij}(\mathbf{s}))^{1-\theta_{ij}},\end{aligned}$$

where Θ_i is a random true status vector and $\mathbb{P}_i(\mathbf{s})$ is the rational belief vector on i's probabilities given the evidence in state $\mathbf{s}$.

Assumptions 3 and 4 imply that if the testing cost is zero (i.e., $c_i(\cdot)=0$), then it is always better to test than not (i.e., $\mathbb{E}_i[r_i(\mathbb{P}_i(\mathbf{S}'))|\mathbf{s}] \geq \overline{r}_i(\mathbb{P}_i(\mathbf{s}))\ \forall\, \mathbf{s} \in \mathscr{S}$). This economic model can be adapted to include behavioral factors. For instance, it can account for the psychological stress from genetic tests as a cost while treating the relief from anxiety because of prophylactic surgeries as a reward. These assumptions mean that our model only applies to genetic testing of genes with known strategies to either prevent or treat the underlying diseases and improve a patient's health.

4.3. Single-Gene Cost and Reward Functions

We start with a schematic of the point estimations from the existing literature for the no-testing censored reward. The empirical data present four scenarios: a true noncarrier with no behavior change $r_i^0(0)$, a true noncarrier with censored treatment options $\overline{r}_i^0(1)$, a true carrier with censored treatment options $\overline{r}_i^1(1)$, and a true carrier with no behavior change $r_i^1(0)$. We use net QALY percentage in Figure 3(a) to illustrate these reward functions.

For expository purposes, we will focus on the single-gene setting in this section (i.e., $|\mathcal{J}|=1$). For cost and reward functions in the multigene setting, see Online Appendix C.2. Based on Definition 2, the cost function simplifies to $c_i(\mathbf{s}) = c_i^{\mathrm{s}} + \mathbb{P}_i(\mathbf{s})c_i^{\mathrm{f}}$.

Next, we construct two linear reward functions, $\overline{r}_i^0(\mathbb{P}_i(\mathbf{s}))$ for the true noncarriers and $\overline{r}_i^1(\mathbb{P}_i(\mathbf{s}))$ for the true carriers, as illustrated in Figure 3(b). This linearity is not required by our model. However, data do not exist in the literature at this time to populate a more complicated function.

Assumption 5 (Oracle Linearity). *Assuming linear reward functions, we have*

$$\overline{r}_i^0(\mathbb{P}_i(\mathbf{s})) = (1-\mathbb{P}_i(\mathbf{s}))r_i^0(0) + \mathbb{P}_i(\mathbf{s})\overline{r}_i^0(1)$$

$$\overline{r}_i^1(\mathbb{P}_i(\mathbf{s})) = (1-\mathbb{P}_i(\mathbf{s}))r_i^1(0) + \mathbb{P}_i(\mathbf{s})\overline{r}_i^1(1).$$

Based on Definition 4, the expected no-testing reward $\overline{r}_i(\mathbb{P}_i(\mathbf{s})) \triangleq \mathbb{E}_i[\overline{r}_i^{\Theta_i}(\mathbb{P}_i(\mathbf{s}))|\mathbf{s}]$ is a linear combination of $\overline{r}_i^0(\mathbb{P}_i(\mathbf{s}))$ and $\overline{r}_i^1(\mathbb{P}_i(\mathbf{s}))$, as shown in Figure 3(b): that is,

$$\overline{r}_i(\mathbb{P}_i(\mathbf{s})) = (1-\mathbb{P}_i(\mathbf{s}))\overline{r}_i^0(\mathbb{P}_i(\mathbf{s})) + \mathbb{P}_i(\mathbf{s})\overline{r}_i^1(\mathbb{P}_i(\mathbf{s})). \qquad (11)$$

Note that this is a quadratic function of $\mathbb{P}_i(\mathbf{s})$.

When tested, an individual's reward is no longer censored. As shown in Figure 3(c), we construct the reward function for the expected testing reward $\mathbb{E}_i[r_i(\mathbb{P}_i(\mathbf{S}'))|\mathbf{s}]$ as

$$\mathbb{E}_i[r_i(\mathbb{P}_i(\mathbf{S}'))|\mathbf{s}] = (1-\mathbb{P}_i(\mathbf{s}))r_i^0(0) + \mathbb{P}_i(\mathbf{s})r_i^1(1). \qquad (12)$$

It is a linear combination of the true rewards with no censoring, $r_i^0(0)$ and $r_i^1(1)$, where individuals find out about their true status and behave according to their rational belief (i.e., true status).

Figure 3. (Color online) Illustration of Cost and Reward Functions in the Single Gene Setting

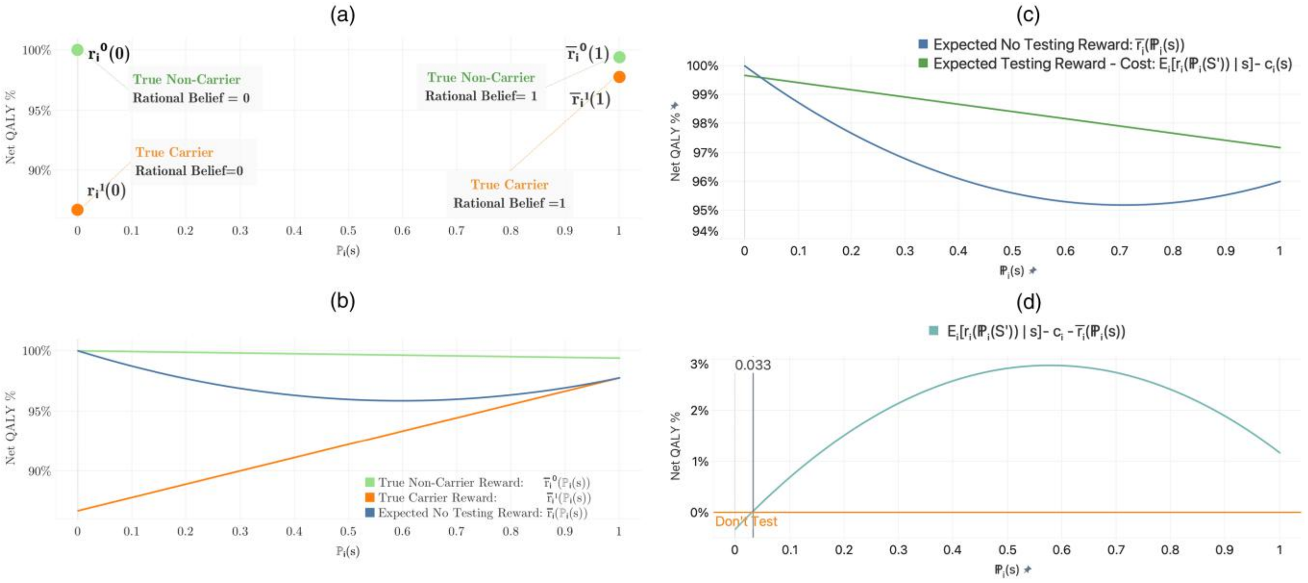

Notes. (a) Empirical Boundary Points for $r_i^0(0), \overline{r}_i^0(1), \overline{r}_i^1(1)$, and $r_i^1(0)$. (b) Reward functions $r_i^0(\mathbb{P}_i(\mathbf{s})), r_i^1(\mathbb{P}_i(\mathbf{s})), \overline{r}_i(\mathbb{P}_i(\mathbf{s}))$. (c) The expected no-testing reward $\overline{r}_i(\mathbb{P}_i(\mathbf{s}))$ vs. the expected testing reward minus cost $\mathbb{E}_i[r_i(\mathbb{P}_i(\mathbf{S}'))|\mathbf{s}] - c_i(\mathbf{s})$. (d) The net benefits of testing $\mathbb{E}_i[r_i(\mathbb{P}_i(\mathbf{S}'))|\mathbf{s}] - c_i(\mathbf{s}) - \overline{r}_i(\mathbb{P}_i(\mathbf{s}))$.

5. Optimal Policies

In this section, we explore the structural properties of optimal testing, continuing with the single-gene setting from the previous section. We decouple the original problem into two decisions: when to stop and whom to test. In the general case, optimal stopping adopts a two-sided threshold policy, segmenting the probability space into three parts. It is optimal to continue when the carrier probability lies between these thresholds. The gap between the two thresholds narrows with increasing costs and widens with increasing rewards. To build intuition about optimal stopping, we also consider an "independent policy" that treats each family member as an unrelated individual. With censoring, optimal individual testing loses the upper threshold and becomes a single-threshold policy. In the general setting with multiple family members, the dual-threshold remains, but the thresholds adjust as individuals are tested, factoring in the future implications of test outcomes.

Regarding whom to test, the choice is highly dependent on the family structure and risk profile. In the special case of sibling-only testing scenarios, we show that the sequence can be determined a priori (i.e., before any testing). Furthermore, the state space reduces dramatically to a count vector of carriers with zero, one, or two copies, enabling optimal solutions in $\mathcal{O}(N^2)$.

5.1. When to Stop

Here, we show that the two-sided threshold policy holds for the optimal stopping problem for family testing in general. Formally, we have Theorem 1.

Theorem 1. *For every state* $\mathbf{s}$, *if there exists an untested individual* $i \in \mathcal{I} \setminus \mathbf{s}$ *such that* $\mathbb{P}_i(\mathbf{s})$ *is between two state-dependent thresholds* $\mathbb{P}_i^{min}(\mathbf{s})$ *and* $\mathbb{P}_i^{max}(\mathbf{s})$, *where* $\mathbb{P}_i^{min}(\mathbf{s}) \le \mathbb{P}_i^{max}(\mathbf{s})$, *then it is optimal to continue testing, and otherwise, it is optimal to stop.*

Consider an "independent policy" that treats each family member as an unrelated individual in (4). When the decision narrows down to just one individual (i.e., $|\mathcal{I}| = 1$), Equations (11) and (12) suggest a simple decision rule for deciding whether to test i: test if the individual expected net benefit is positive. As shown in Figure 3, the expected net benefit (i.e., the difference between testing reward minus cost versus no-testing reward), $\mathbb{E}_i[r_i(\mathbb{P}_i(\mathbf{S}'))|\mathbf{s}] - c_i(\mathbf{s}) - \overline{r}_i(\mathbb{P}_i(\mathbf{s}))$ is a quadratic function with closed form $\mathbb{P}_i^{min}(\mathbf{s})$ and $\mathbb{P}_i^{max}(\mathbf{s})$. Recall that c_i^{s} is the testing cost, c_i^{f} is the follow-up cost, $d_i = r_i^1(1) - \overline{r}_i^1(1)$ is the censoring effect, and $r_i^\delta = r_i^0(0) + \overline{r}_i^1(1) - \overline{r}_i^0(1) - r_i^1(0)$ is the total benefit of censored interventions.

Corollary 1. *For any given individual* $i \in \mathcal{I} \setminus \mathbf{s}$ *treated independently from other family members and given the individual risk* $\mathbb{P}_i(\mathbf{s})$, *an optimal solution follows a risk-based threshold policy:*

$$\pi = \begin{cases} i & \text{if } \mathbb{P}_i(\mathbf{s}) \ge \dfrac{r_i^\delta + d_i - c_i^f - \sqrt{(r_i^\delta + d_i - c_i^f)^2 - 4c_i^s r_i^\delta}}{2r_i^\delta} \\ & \text{and } (r_i^\delta + d_i - c_i^f)^2 \ge 4c_i^s r_i^\delta \\ \varnothing & o.w., \end{cases} \tag{13}$$

where $r_i^\delta = r_i^0(0) + \overline{r}_i^1(1) - \overline{r}_i^0(1) - r_i^1(0)$ *and* $d_i = r_i^1(1) - \overline{r}_i^1(1)$.

Because of censored reward, only the lower threshold $\mathbb{P}_i^{min}(\mathbf{s})$ exists in $[0, 1]$. The upper threshold $\mathbb{P}_i^{max}(\mathbf{s}) = \left(r_i^\delta + d_i - c_i^{\mathrm{f}} + \sqrt{(r_i^\delta + d_i - c_i^{\mathrm{f}})^2 - 4c_i^{\mathrm{s}} r_i^\delta}\right)/2r_i^\delta$ is greater than one and therefore, is not meaningful (see the proof of Corollary 1 in Online Appendix B). This policy is a function of both the testing and follow-up costs and the individualized rewards at the boundary conditions. Critically, this policy is significantly different from any policy or guideline with a static threshold that is the same for everyone.

If we were to access full treatment options without censoring (i.e., $d_i = 0$), then the upper bound $\mathbb{P}_i^{max}(\mathbf{s})$ exists within $[0, 1]$ (Corollary EC.1 in the Online Appendix). If there exist $\mathbb{P}_i^{min}(\mathbf{s})$ and $\mathbb{P}_i^{max}(\mathbf{s})$, the gap between the two thresholds is $\left(\sqrt{(r_i^\delta - c_i^{\mathrm{f}})^2 - 4c_i^{\mathrm{s}} r_i^\delta}\right)/r_i^\delta$. Its derivative on r_i^δ is always positive. That is, as r_i^δ increases, the gap between $\mathbb{P}_i^{min}(\mathbf{s})$ and $\mathbb{P}_i^{max}(\mathbf{s})$ widens. The derivatives with respect to c_i^{s} and c_i^{f} are always negative. Therefore, as costs increase, the gap between $\mathbb{P}_i^{min}(\mathbf{s})$ and $\mathbb{P}_i^{max}(\mathbf{s})$ narrows. For the general case, the gap between the thresholds in Theorem 1 narrows as costs increase, but the impact of the information is not trivial; see Online Appendix B for more information.

5.2. Whom to Test Next

We next show the sequence of whom to test. In this section, we show that in general, this sequence cannot be fixed a priori. We start by discussing the example family from Figure 1 in detail.

5.2.1. Family Example. We now illustrate how an optimal policy outperforms other heuristics via information sharing using the family shown in Figure 2 as an example. For this family, two of the children have tested positive, and the grandmother 1 has tested negative. Under the current guidelines, depending on their ages, the mom 6 and the maternal grandmother 3 will likely be recommended for BRCA1/2 testing.

However, testing the paternal grandfather 2 can simultaneously resolve uncertainties on both sides of the family. If he is not a carrier, through the laws of inheritance, we can definitively conclude that 6 carries the mutation and that 5 does not as shown in Figure 2(b). If he is a carrier, the probabilities will update, and 5 will be more likely to be a carrier; however, we would not know his true status conclusively. Testing any other individual in Figure 2(a) would also be nonconclusive for other members. Therefore, only by testing 2 can we conclusively determine other family members' statuses. Despite the cost and limited QALY improvement for 2, the optimal policy is likely to improve more QALYs for the family compared with an uncoordinated approach.

Thus, the decision maker may prioritize information value over individual QALY gains. For instance, although testing those with higher individual benefits, like 3 and 6, may seem optimal, selecting a family member with significant informational value can benefit the entire family.

5.2.2. Family Testing Sequence. From the example, one might think that testing older generations first for their informational value is optimal. However, that is not always the case. In Figure 4(a), the optimal policy begins by testing those with higher QALY gains, typically younger family members. If a young member tests positive, older generations are then tested; if not, testing ceases for the family. This approach is clinically termed "cascade upward."

An optimal genetic testing policy often alternates between cascading upward and downward, navigating through generations. Using another family with a different risk profile but identical structure, we show that a simply structured policy is unlikely to be optimal for an arbitrary family (Figure 4(b)). Here, the optimal policy starts with older generations (e.g., great grandparents) who offer limited personal QALY gains but hold considerable informational value. Depending on their results (zero, one, or two mutations), the subsequent tests can target any remaining generation, with each subsequent test contingent on the system's state. This family structure, a basic four-generation model with a nuclear family spanning two generations, is not exclusive in exhibiting such behaviors.

5.3. Sibling Testing

Sibling testing refers to testing only the siblings in a family, typically when parents are unavailable for tests and the next generation is too young. In sibling testing as opposed to the general case, we show that the value function is nondecreasing in r_i^δ of the tested siblings. This enables us to rank the siblings by their r_i^δ's and reduce the dynamic family testing problem to a fixed sequence testing problem. The state space for sibling testing also simplifies to a space of possible counts of test results. In addition, there exist equivalent classes of sibling results. These three properties allow us to break down the dynamic decision making into individual decisions on each sibling, making optimal sibling testing solvable in polynomial time, as detailed in Theorem 2. We further show in Corollary 2 that the optimal sibling testing follows a two-sided count-dependent threshold policy.

5.3.1. Fixed Sequence Testing.

Lemma 1. *Assuming a censoring effect proportional to* r_i^δ, *substituting testing a sibling with another sibling who has a no less* r_i^δ *will not decrease the value function. That is,*

$$r_k^\delta \le r_{k+1}^\delta \Rightarrow V^{\pi_1}(\mathbf{s}) \le V^{\pi_2}(\mathbf{s}), \tag{14}$$

where sibling k *is tested and* $k+1$ *is untested following policy* π_1 *in one of the sample paths after* $\mathbf{s}$ *and sibling* $k+1$ *is tested and* k *is untested following policy* π_2 *in the same sample path. The two policies are identical otherwise.*

When considering testing a given sibling, if there exists a sibling with a higher r_i^δ, then without loss of

Figure 4. (Color online) (a) An Optimal Policy for a Family of Seven with Family Risk ~0.2% and Cost $2,500 and (b) an Optimal Policy for the Same Family of Structure with Family Risk ~4% and Cost $1,000

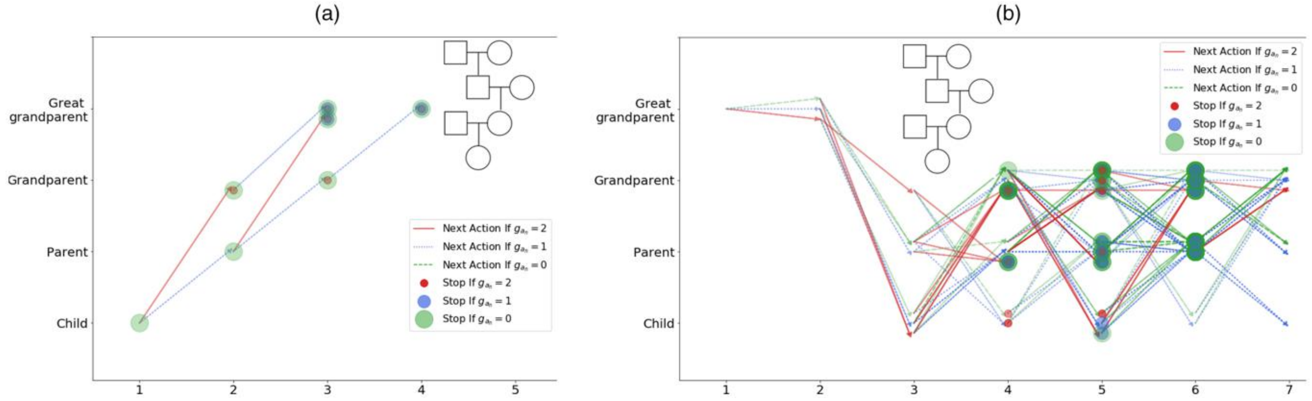

Notes. Solid arrows indicate the next person to test when the current individual has two diseased copies; dotted lines indicate one copy, and dash lines indicate negative results. Colored circles signify stopping the test with the current results.

optimality, the decision maker can switch their testing decisions.

Lemma 1 follows because collecting rewards at each epoch is equivalent to collecting all the rewards after testing ends (Lemma EC.2 in the Online Appendix). In a given end state, swapping siblings between the untested set and the tested set changes the value function. Additionally, the difference is monotone in r_i^δ. This monotonicity reduces dynamic family testing to fixed sequence testing, where we can determine the order of optimal sibling testing a priori.

5.3.2. Counts of the Test Results. To reduce the state space for sibling testing, we represent posterior probabilities and the value functions of sibling testing policies as functions of the counts of the test results. The chain rule of conditional probabilities permits that the posterior probabilities remain the same regardless of the order of events (Koller and Friedman 2009, definition 2.3). Given that and the fact that the siblings carry identical risks, we know that the posterior probability is only dependent on the counts of those tested with zero, one, and two mutations, (n_0, n_1, n_2).

Proposition 1. *In sibling testing, the posterior probability depends only on the counts of test results, that is, for* $i \in \mathbf{s}^n$,

$$\mathbb{P}(G_i = g_i | \mathbf{s}^n) = \mathbb{P}(G_i = g_i | n_0, n_1, n_2) \;\; \forall\, g_i \in \{0, 1, 2\},$$

where (n_0, n_1, n_2) *are the counts of test results* 0, 1, *and* 2, *respectively, and* $n_0 + n_1 + n_2 = n$.

Because the posterior probabilities only depend on the counts (n_0, n_1, n_2) and because the order of testing is also fixed in optimal sibling testing, we can consider states with the same counts of zero, one, and two as *equivalent* classes of states $\mathbf{s} \equiv \mathbf{s}'$ if and only if $(n_0(\mathbf{s}), n_1(\mathbf{s}), n_2(\mathbf{s})) = (n_0(\mathbf{s}'), n_1(\mathbf{s}'), n_2(\mathbf{s}'))$.

Proposition 2. *In sibling testing,* $V^*(\mathbf{s}^n)$ *is a function of* (n_0, n_1, n_2), *that is,*

$$V^*(\mathbf{s}^n) = V^*(n_0, n_1, n_2). \tag{15}$$

From Propositions 1 and 2, we can represent posterior probabilities and the value functions as functions of the counts of the test results. We can further show the transitions between different equivalent classes of states in sibling testing as a graph (Figure 5). Each node represents a type of equivalent class of test results (i.e., states). Each curve means a new test result that can change the posterior probabilities. For example, once both zero and two are present in the tested siblings, regardless of new test results, the posterior probabilities of the untested siblings no longer change. One can think of node {0, 1, 2} as an absorbing state in terms of *posterior probabilities*. However, the *value functions* will continue to change for node {0, 1, 2} as more siblings get tested.

Figure 5. Equivalent Classes-Type Transition Diagram for Sibling Testing

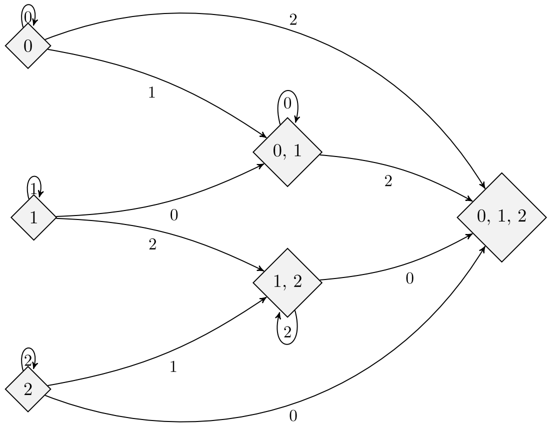

Notes. An arrow (→) signifies a new result that alters the posterior probabilities of untested siblings. For the node {0,1}, only a result of zero impacts the posterior probabilities, whereas results of one do not. When a new result of two emerges, the state transitions to the equivalent class type {0,1,2}.

5.3.3. Optimal Sibling Testing. Using the graph shown in Figure 5, the total number of equivalent classes equals the sum of size 1, size 2, and size 3 sets (i.e., $3N + 2\sum_{k=1}^{N-2} k + N - 3 = N^2 - N + 3$, where k is the number of ones in node (0, 1) or (1, 2)). As such, a backward induction algorithm solves testing with N siblings in $\mathcal{O}(N^2)$.

Theorem 2. *Given an ordered set of* $\{r_i^\delta\}_{i\in\mathcal{I}}$ *for a set of siblings* $\mathcal{I}$ *for testing, it is optimal for the decision maker to only consider testing the remaining sibling who has the largest* r_i^δ. *Furthermore, the optimal value function for* N *siblings can be found in* $\mathcal{O}(N^2)$.

Theorem 2 is a direct consequence of Lemma 1, Lemma EC.2 in the Online Appendix, Proposition 1, and Proposition 2. The ordered policy guarantees that whenever stopping happens, no sibling has a higher r_i^δ in the untested set. Therefore, in any end state, the siblings in the untested set are those with the lowest r_i^δ, and the value function cannot be further improved.

Furthermore, optimal sibling testing follows a two-sided count-dependent threshold policy (i.e., Corollary 2).

Corollary 2. *Given an ordered set of* $\{r_i^\delta\}_{i\in\mathcal{I}}$ *for a set of siblings* $\mathcal{I}$ *for testing, for every* (n_0, n_1, n_2), *there exists a* $\mathbb{P}^{min}(n_0, n_1, n_2)$ *and a* $\mathbb{P}^{max}(n_0, n_1, n_2)$ *where* $\mathbb{P}^{min}(n_0, n_1, n_2) \le \mathbb{P}^{max}(n_0, n_1, n_2)$ *such that if the untested individual with the largest* r_i^δ *has posterior probability* $\mathbb{P}_i(\mathbf{s})$ *between them, then it is optimal to continue testing, and otherwise, it is optimal to stop.*

This corollary follows the general case in Theorem 1 and the fixed order testing in Theorem 2.

6. Numerical Experiments

To better understand the impact of optimal policy on family genetic testing, we conduct several numerical experiments in a realistic setting of BRCA1/2 testing. Our study involves families with different sizes, structures, and risk profiles. We also conduct sensitivity analyses around the family members' ages and testing costs. We first examine current guidelines for individual testing and distinguish two versions: namely the independent policy and the myopic updating policy.

- Independent policy. The independent policy, or batch testing, is a policy given by the RHS of (4) considering each person individually: $\forall\, i \in \mathcal{I}$, test if $\mathbb{P}_i(\mathbf{s}) \geq r_i^\delta + d_i - c_i^{\mathrm{f}} - \sqrt{(r_i^\delta + d_i - c_i^{\mathrm{f}})^2 - 4c_i^{\mathrm{s}} r_i^\delta}/2r_i^\delta$ (Corollary 1). It tests each individual whose net benefits are positive when considered independently. This policy represents the best current practice in terms of cost-effectiveness as it rationalizes the best guidelines for individual testing. However, there is no sharing of information across the family.
- Myopic updating policy. The myopic updating policy derives from the RHS of (4) assuming $V(\mathbf{S}') = 0$. It maximizes the individual net benefits at the current state. The myopic policy updates the test results for the rest of the family, but the independent policy does not. The myopic updating policy is considered the best available alternative when only the needs of individuals and not the potential family-wide consequences are taken into account, but the results are shared.

6.1. Parameters for BRCA1/2 Testing

Our BRCA1/2 testing relies on two primary information sources. The first is the clinical tool BRCAPRO, which defines the state space, initializes the joint distribution μ^0, updates the posterior risk $\mathbb{P}_i(\mathbf{s})$, and computes state transition probabilities $p_{a^{n-1}}(\mathbf{s}^{n-1}, \mathbf{s}^n)$. The second source is the work presented in Nohdurft et al. (2017), which is instrumental in determining individual reward functions for BRCA1/2 testing. To integrate these sources in parameterizing the MDP, we modify the methods described in Nohdurft et al. (2017).

6.1.1. Establishing State Space and Calculating $\boldsymbol{\mu}^0$, $\mathbb{P}_i(\mathbf{s})$, and $p_{a^{n-1}}(\mathbf{s}^{n-1}, \mathbf{s}^n)$ Using BRCAPRO. The BRCAPRO model, which is widely used and has been validated (Berry et al. 2002) and updated (Chen et al. 2020), assumes a uniform risk for all pathogenic variants. Additionally, it posits that an individual cannot have two mutations: either in the BRCA1 or in the BRCA2 gene. Based on these assumptions, we simplify our MDP's state space by limiting the values for $\boldsymbol{\theta}_i$ (carrier status) to four: negative (0, 0), BRCA1 positive (1, 0), BRCA2 positive (0, 1), or both BRCA1 and BRCA2 positive (1, 1).

The initial prior distribution μ^0 is the posterior distribution output from BRCAPRO utilizing population priors, individual cancer status, and family history of breast and ovarian cancer as inputs. As new test results emerge and the state evolves, BRCAPRO updates the posterior risk $\mathbb{P}_i(\mathbf{s})$. Moreover, the probability $p_{a^{n-1}}(\mathbf{s}^{n-1}, \mathbf{s}^n)$ governing state transitions adheres to Mendelian inheritance laws, and it is determined by $\mathbb{P}_i(\mathbf{s})$.

6.1.2. Parameterizing Reward Functions Based on Procedures from Nohdurft et al. (2017). Our model adopts the QALYs calculation procedures outlined in Nohdurft et al. (2017), incorporating factors such as BRCA1/2 status, age, preventative surgeries, cancer type, stage, and cause of death (whether because of cancer or because of other causes). For a comprehensive understanding, refer to the original study. As we extend upon this work, our first step is to replicate the optimal values from their model strictly following the QALY-maximizing surgery sequence specific to BRCA1 and BRCA2 carriers as depicted in Figure EC.1 in the Online Appendix. The results, presented in Table EC.3 in the Online Appendix, are consistent with previously documented life expectancies and cumulative cancer risks, considering both the presence and absence of preventative surgeries.

To ensure a seamless integration of the reward functions with our MDP, we make the following modifications to the procedures from Nohdurft et al. (2017).

1. Recalibrate the reward functions using BRCAPRO's recently updated cancer risks and the latest clinical and statistical data.
2. Extend the model's scope to include both male and female patients, accounting for unique risks and preventive measure impacts.
3. Transform QALYs into net QALYs, adapting to a cost-effective framework. We adopt a \$50,000/QALY willingness-to-pay threshold, aligning with the American College of Cardiology and the American Heart Association standards. These net QALYs set reward parameters for each genetic status.
4. Combine single-gene functions to create multigene reward functions and introduce crossapplication for BRCA1 and BRCA2 preventive measures. These facilitate expected reward function calculations as per Definition 3.
5. Formulate censored reward functions with a 50% heuristic censoring effect and use Equations (EC.13)–(EC.16) in the Online Appendix to generate censored rewards.

Table 2. Risk Profiles of a Nuclear Family with Two Daughters

Individual	Age, years	Personal history	Tested	$\mathbb{P}_i(\mathbf{s})$, % BRCA1	BRCA2	BRCA1 &2
Father	62	No pre-existing conditions	No	2.12	1.07	0.0020
Mother	64	No pre-existing conditions	No	0.54	0.49	0.0002
Daughter 1	44	**Ovarian cancer**	No	2.61	1.50	0.0012
Daughter 2	40	No pre-existing conditions	No	0.91	0.62	0.0005

Note. Daughter 1's cancer diagnosis is shown in bold.

For detailed explanations of the implementation and modifications, please refer to Online Appendix C.

6.2. Policy Comparisons

In this section, we will explore how the optimal policy improves cost-effectiveness over competing policies. To provide insight into the various settings where our framework can be applied, we have selected three families with contrasting characteristics—a small, high-risk nuclear family; a small, low-risk nuclear family; and a larger family with complex relationships and diverse risk profiles.

6.2.1. Example in Which the Optimal Policy Tests Someone Who the Myopic/Independent Policies Would Not. To illustrate this, consider a nuclear family with two daughters as an example, as shown in Table 2. As the cost of testing decreases, the family is weighing the option of getting tested now. Daughter 1, who had ovarian cancer in the past but did not undergo genetic testing, underwent bilateral salpingo-oophorectomy surgery as part of her treatment. At the age of 44, if she chooses not to have any further prophylactic surgery, her remaining QALYs will be 31.73 for BRCA1 ($r_i^{1,0}(0,0)$) or 32.78 for BRCA2 ($r_i^{0,1}(0,0)$) on average across all ovarian cancer subtypes and stages. However, if she decides to have another bilateral mastectomy surgery based on the results of a genetic test, her remaining QALYs will increase slightly to 32.98 for BRCA1 ($r_i^{1,0}(1,0)$) or 33.10 for BRCA2 ($r_i^{0,1}(0,1)$). For comparison, if she is not a carrier, her remaining QALYs will be 37.57 ($r_i^{0,0}(0,0)$). The other reward parameters for this family are shown in Table EC.7 in the Online Appendix.

However, daughter 1 has not undergone genetic testing to determine her carrier status for BRCA1 and BRCA2 yet. Her current risk of being positive for BRCA1 is 2.61%, for BRCA2 is 1.50%, and for both is 0.0012%. Under a myopic or independent policy, where the rest of her family is not considered, her remaining QALYs would be 37.37 ($\mathbb{E}_i[r_i(\mathbb{P}_i(\mathbf{S}'))|\mathbf{s}]$) if she chooses to get tested and follow the QALY-maximizing surgery sequence in case of a positive result. If she decides to not get tested and only takes limited preventative options, her expected remaining QALYs are 37.35 ($\bar{r}_i(\mathbb{P}_i(\mathbf{s}))$). The benefit of testing for her is a bit less than 0.02 QALYs or at \$50,000/QALY, \$1,000. Therefore, at a cost of \$1,000, the myopic and independent policies do not recommend testing her. Her sister (daughter 2), on the other hand, is currently cancer free at the age of 40 with a lower risk profile, but she has tested remaining QALYs of 41.88 and no-test remaining QALYs of 41.80. The benefit is 0.08 QALYs, about four times as much as that of daughter 1.

However, the optimal policy recognizes that daughter 1's higher risk could provide valuable information for the family (i.e., a larger $\mathbb{E}[V(\mathbf{S}')]$). This yields a family value of 0.0914 unconditional on any test results (Table 3). With the BRCA1 (2.6%) and BRCA2 (1.5%) incidence rates, this could represent 3.5 QALYs for BRCA1 positives, 6.1 QALYs for BRCA2 positives, or any values in between. If the added QALYs are proportional between BRCA1 and BRCA2 upon daughter 1 testing positive, we observe a potential increase of 2.8 QALYs for BRCA1 and 1.2 QALYs for BRCA2 ($2.6\% \times 2.8 + 1.5\% \times 1.2 \approx 0.0914$). Although the benefits in QALYs seem subtle in expectation, the influence on each diagnosed case is pronounced. Identifying the right family can mean notable quality-of-life enhancements (Abbott 2023).

6.2.2. Example in Which the Optimal Policy Tests at a Higher Price Point than the Myopic/Independent Policies and Outperforms Them Arbitrarily. Considering the testing policy as a function of cost, for a low-risk

Table 3. Reward Functions and Bellman Equation Breakdown

Individual	$r_i^{1,0}(0,0)$	$r_i^{0,1}(0,0)$	$r_i^{1,0}(1,0)$	$r_i^{0,1}(0,1)$	$r_i^{0,0}(0,0)$	$\mathbb{E}_i[r_i(\mathbb{P}_i(\mathbf{S}'))\|\mathbf{s}]$	$\bar{r}_i(\mathbb{P}_i(\mathbf{s}))$	$c_i(\mathbf{s})$	Individual net QALY	$\mathbb{E}[V(\mathbf{S}')]$
Father	15.84	14.42	15.86	14.71	20.00	19.85	19.84	0.022	−0.0162	0.0897
Mother	14.09	15.21	16.62	16.75	20.97	20.93	20.91	0.021	−0.0029	0.0738
Daughter 1	**31.73**	**32.78**	**32.98**	**33.10**	**37.57**	**37.37**	**37.35**	**0.022**	**−0.0003**	**0.0917**
Daughter 2	28.68	33.27	35.25	35.79	41.99	**41.88**	**41.80**	0.021	0.0600	**0.0132**

Note. Numbers referenced in the text are in bold.

Table 4. Risk Profiles of a Nuclear Family with Two Daughters Without a History of Cancer

Individual	Age, years	Personal history	Tested	BRCA1, %	BRCA2, %	BRCA1 &2, %
Mother	62	No pre-existing conditions	No	0.06	0.09	0.0001
Father	64	No pre-existing conditions	No	0.11	0.13	0.0001
Daughter 1	44	No pre-existing conditions	No	0.09	0.11	0.0001
Daughter 2	40	No pre-existing conditions	No	0.09	0.11	0.0001

family with two daughters and no history of cancer regarding BRCA1/2 testing (Table 4), the independent policy, which lacks a specific testing order, tests daughters at \$465, the mother at \$340, and the father at \$45 (Figure 6(a)). In contrast, the myopic updating policy, which tests sequentially and shares test results within the family, tests daughter 2 at \$465. If she tests positive or as the cost goes down, daughter 1 and the mother are also tested. The father is only deemed cost effective to test when the price falls below \$45. Such a policy, contingent on familial test results, effectively prevents unnecessary tests, for example, when a member has tested negative.

The optimal policy, which is even more cost effective, permits testing at costs up to \$860. Initially, it mirrors the myopic policy's sequence at higher costs. However, when the price drops to \$400, it pivots, prioritizing parents' tests and using their results to decide on testing the daughters. This shift might stem from the higher costs making it inefficient to test the father initially.

At a \$550 testing cost, both existing policies, the independent policy and the myopic updating policy, do not test any family members, whereas the optimal policy tests for the family's benefit (Figure 6). The total value difference in net QALYs between the optimal and myopic updating policies grows linearly with the number of daughters (Figure 7(b)), suggesting that an optimal policy outperforms these existing policies arbitrarily.

The overall structure of the optimal policy for low-risk families, like the one described in this scenario, is not sensitive to the age of the family members. As demonstrated in Figure 7(a), the value function difference between the optimal and myopic updating policies increases at a cost of around \$900 and then decreases as the cost drops to around \$465, where the myopic

Figure 6. (Color online) Independent, Myopic Updating, and Optimal Testing at \$0–\$1,000 for the Family in Table 4

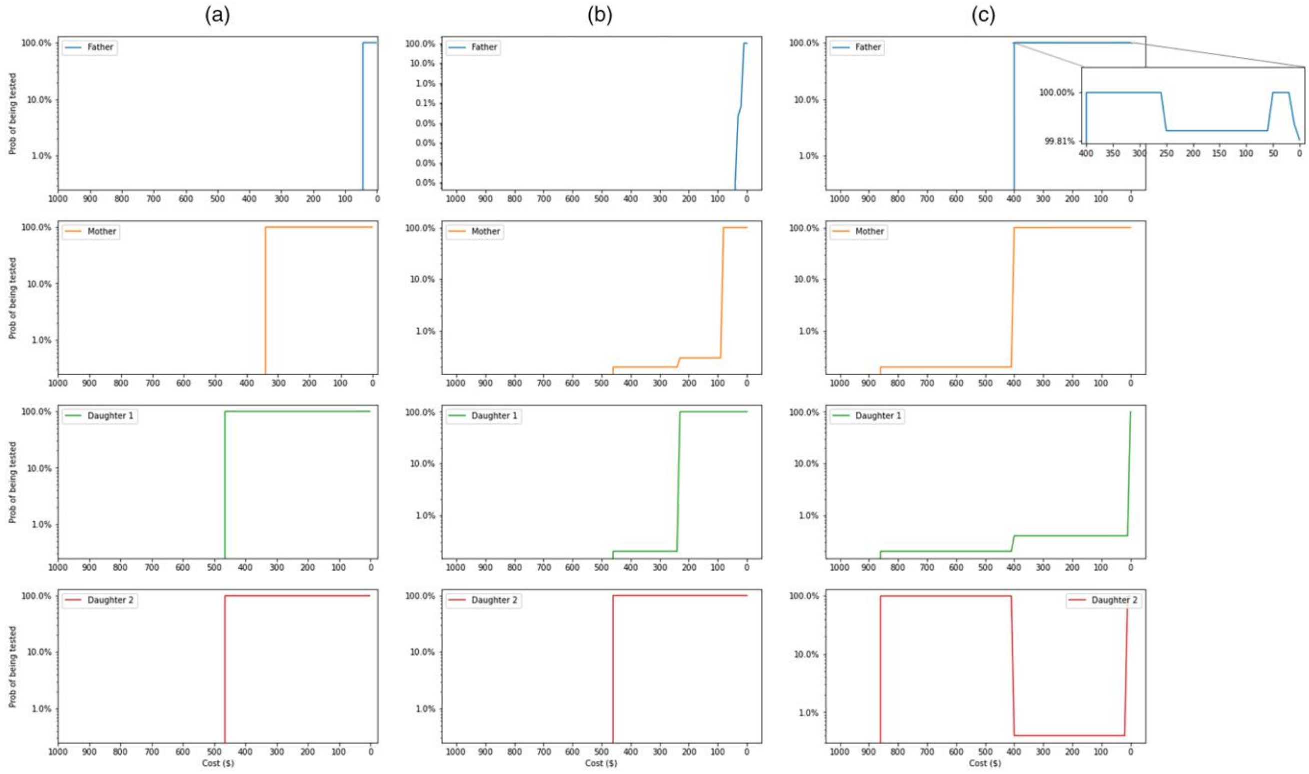

Notes. (a) Independent. (b) Myopic updating. (c) Optimal testing.

Figure 7. (Color online) QALYs Differences Between Optimal and Myopic Updating Policy by Varying Characteristics of Family Shown in Table 4.

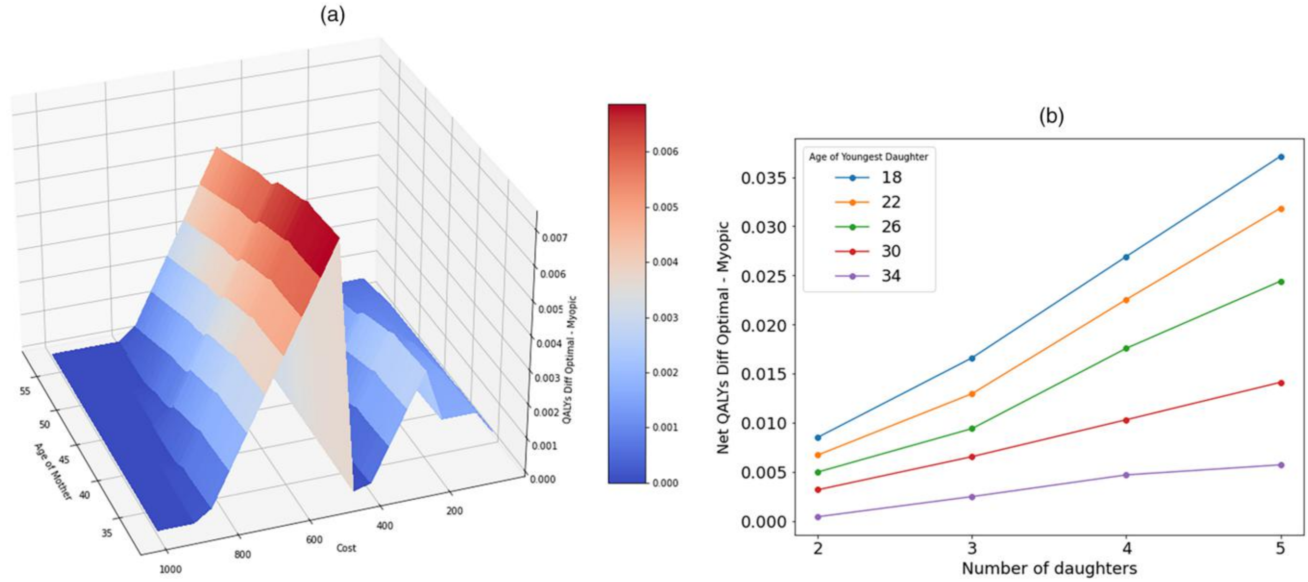

Notes. (a) QALYs Differences (Optimal Minus Myopic) by Various Mother's Ages and Price Points of the Family Shown in Table 4. (b) QALYs differences with an increasing number of daughters in the family at $550 per test by varying the youngest daughter's age, assuming that each daughter is born two years apart.

strategy begins testing the family. This pattern reappears at a lesser degree before disappearing again as the cost decreases to $0, at which point everyone in the family is tested regardless of the policy being followed. This piecewise linear structure is a result of both testing at higher costs provided by the optimal policy, which leads to a gain in net QALYs, and the cost savings resulting from decreasing testing costs. The robustness of the structure to changes in the mother's age is established in Figure 7(a). In practical terms, the optimal policy allows low-risk families like this one to maximize their benefits when the cost is intermediate, and sharing test results among family members can lead to additional cost savings by limiting unnecessary testing.

The value gains from the optimal policy align with reported QALYs gains from BRCA1/2 testing. At age 30, high-risk populations, such as Ashkenazi Jewish women, average a 0.09-QALY increase from BRCA1/2 testing (Manchanda et al. 2015), whereas low-risk groups gain between 0.003 and 0.007 QALYs testing at age 30 (Manchanda et al. 2018, Guzauskas et al. 2020). This comparison is striking; our optimal policy, if widely adopted, could lead to health benefits akin to direct testing but without the associated costs—a significant value proposition for both families and the broader healthcare system.

6.2.3. Example in Which an Optimal Policy Coordinates a Large Complex Family for Testing.

To demonstrate the limitations of existing guidelines in determining the optimal testing strategy for a complex family, we use an example of a large family with various relationships

Table 5. Risk Profiles of a Family with 11 Individuals

Individual	Age, years	Personal history	Tested	BRCA1, %	BRCA2, %	BRCA1 &2, %
Maternal grandmother	60	No pre-existing conditions	No	0.64	0.55	0.0004
Maternal grandfather	62	No pre-existing conditions	No	1.97	1.03	0.0022
Paternal grandmother	66	No pre-existing conditions	Negative	0.00	0.00	0.00
Paternal grandfather	68	No pre-existing conditions	No	0.11	0.11	0.0001
Mother	40	Ovarian	No	2.55	1.51	0.0020
Father	42	No pre-existing conditions	No	0.06	0.05	0.0000
Daughter	20	No pre-existing conditions	No	0.55	0.37	0.0004
Mother's sister 1	42	No pre-existing conditions	No	1.05	0.69	0.0008
Mother's brother-in-law 1	42	No cancer	BRCA1	100.00	0.00	0.0013
Cousin 1	18	No pre-existing conditions	No	49.98	0.38	0.3772
Mother's sister 2	36	No pre-existing conditions	No	1.14	0.73	0.0009

Figure 8. (Color online) Panel (a) Shows the Stopping Time Distribution of Independent, Myopic Updating, and Optimal Policies, and Panel (b) Shows the Value Function Differences Among the Three Policies

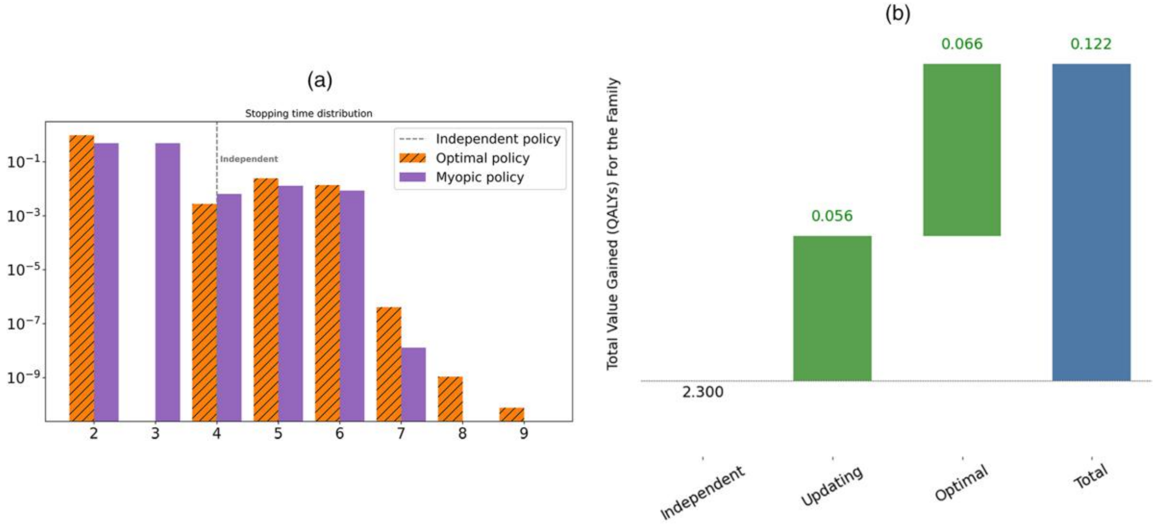

and health conditions (as shown in Table 5). Because of the complex interactions between family members and the dynamic nature of familial risks in such a family, the optimal testing strategy for the rest of the family may change drastically as new test results are revealed. Existing guidelines that focus on individuals are not equipped to adapt to these changes. We have chosen a high testing cost of $2,500 for this example not to make conclusions about this specific family but to better understand the contingent nature of the optimal policy in comparison with other policies.

Based on the independent policy, only four individuals (the mother, the mother's sister, the daughter, and the cousin) are cost effective to test individually. Therefore, the independent policy stops testing after these four individuals (as shown in Figure 8(a)). The myopic updating policy, on the other hand, will test as many as seven individuals, but it is most likely to stop after testing two or three. The optimal policy will sometimes test all nine remaining family members. The optimal policy has a higher probability of stopping after testing two individuals than the myopic updating policy, but once it decides to continue testing, at least two more individuals will need to be tested before stopping.

The myopic updating and optimal policies have a wider range of tested individuals compared with the independent policy. This is reflected in the value function as seen in Figure 8(b). The myopic updating policy mainly increases the value through cost savings as the sharing of information leads to more informative testing.

Compared with the independent policy, the optimal policy is able to increase the value of the genetic testing for this family by an additional 5.2% (0.122 of 2.3 QALYs). This demonstrates the ability of the optimal policy to not only identify individuals who would benefit from testing but also, maximize the overall value of the testing by considering the informational value of the test results and the potential impacts on future states. In other words, the optimal policy takes a more comprehensive approach to testing, considering not only the immediate benefits of testing but also the potential family benefits and costs of the testing. As a result, the optimal policy increases the value of the testing beyond what the myopic updating policy or the independent policy could achieve.

7. Conclusion

This study presents a comprehensive modeling framework for genetic testing aimed at maximizing the aggregated social welfare of a family. Our Markov decision process framework incorporates family relationships with a state space capturing the laws of inheritance and uses a Bayesian network of genetic statuses to compute and update family members' risks after each test result. Additionally, we provided a detailed characterization of the reward functions for making rational decisions on whether to test or remain untested. Using this framework, we show that optimal policies can outperform current individualistic testing policies in terms of cost-effectiveness across a diverse set of families.

In the general family setting, we show that optimal stopping follows a structure with two-sided thresholds, previously known only for individual testing. In the special case of sibling testing, we also determine a priori the optimal sequence to test. To numerically test our framework in a realistic setting, we integrate into our model a widely used clinical tool that calculates a person's probability of carrying a deleterious mutation of

BRCA1 and BRCA2 based on the person's cancer status and personal or family history of breast and ovarian cancer (Chen et al. 2004, 2020).

To make optimal genetic testing of families a reality, two key steps must be taken. First, we need to evaluate the overall impact of our model on society at large, taking into account factors such as population awareness and stigma surrounding genetic testing, which can vary widely by region. Additionally, we need to gather more data on family structures and ages to extrapolate our findings to the population level. Currently, data on family structures and ages are limited, with the U.S. Census Bureau only covering nuclear families (U.S. Census Bureau 2020) and having limited data on the ages of mothers (U.S. Census Bureau 2012).

Second, in order for optimal genetic testing to be implemented, there must be a number of legal and policy changes, provided that the study demonstrates substantial benefits to society. These include convincing payers to test for the benefit of other family members and addressing potential barriers under the Health Insurance Portability and Accountability Act (HIPAA). Additionally, clinical tools must be developed for easy use by clinicians. There are existing tools, such as BRCAPRO, that can be used as a starting point given that our model is no more computationally complex in terms of time than the exact inference problem solved using BRCAPRO.

In order to ensure the relevance and clarity of our study's results to clinical practice, we have made a number of assumptions regarding the reward function, even though the framework does not strictly require them. For example, we have focused on modeling rational behaviors in this study; however, future researchers may wish to study irrational behaviors. This can be done by relaxing the assumption of rationality (Assumption 1) without changing the overall structure of the model. Similarly, relaxing the do-no-harm assumption (Assumption 2) would allow for the inclusion of patients with negative remaining net QALYs. Relaxing the assumption of the superiority of truthful results (Assumption 3) would allow for potentially harmful tests. In addition, we assumed that a positive genetic result is a prerequisite for accessing the full range of treatment options (Assumption 4). Loosening this assumption would lower the upper limit of testing, potentially leading to cost savings. Furthermore, we assume that there is a linear relationship between carrier probability and behavior change (Assumption 5). This assumption serves the purpose of highlighting the nonlinear structure of health benefits as a function of risk (as seen in Figure 3 and Figure EC.4(a) in the Online Appendix), but it is not a necessary assumption. For example, researchers could consider using piecewise linear functions $r_i^{\theta_i}(p)$, in which rewards change at different carrier probability breakpoints corresponding with different treatment protocols.

Although the proposed model has the flexibility to incorporate any exogenous cost into the individual cost function $c_i(\mathbf{s})$, it does not explicitly take into account the emotional cost of genetic testing. Recent research has provided mixed results on the emotional impact of genetic testing, with some studies finding no significant adverse effects and others reporting negative effects. For example, a survey of genetic tests for cardiovascular diseases found no significant adverse effects (Oliveri et al. 2018). However, when it comes to genetic tests for cancer, the existing literature reported conflicting results, with some studies reporting negative effects, no effects, and positive effects. Considering the emotional cost of testing, if the psychological characteristics of patients, such as their anxiety level, were accurately captured, our model may not recommend testing for them even if there would be an information benefit for the rest of the family.

Despite these limitations, this work makes progress in rationalizing decision making around genetic testing problems and provides a foundation for future research in this area.

Acknowledgments

The authors thank the department editor, the associate editor, and the referees for their thoughtful and detailed comments, which greatly improved the presentation of the paper. The authors also thank Dr. David H. Ledbetter and Joe Leader at Geisinger, whose MyCode program of genetic testing (Carey et al. 2016) provided inspiration for this work.

Endnotes

[1] This is the net health benefit in QALYs, which consists of two parts: (1) a patient's health gain (in QALYs) and (2) the health opportunity cost of funding the treatment for this patient. The cost gets converted to QALYs after dividing by a willingness-to-pay threshold. The net QALYs gained is the health gain minus the health cost (Paulden 2020) that a person gains based on the belief in the probability of having a particular genetic variant using empirical outcomes when people know their genetic status with certainty.

[2] See Remark EC.2 in the Online Appendix for more details.

[3] In genetics, the transmission of one gene and the transmission of another gene are generally considered to be probabilistically independent unless they are in close proximity on the same chromosome.

[4] Section 3 (Bayesian network) and Section 4 (economics) can be read in any preferred order.

References

Abbott B (2023) Cancer runs in families. Too few are getting tested. *Wall Street Journal* (August 29), A3.

Alagoz O, Maillart LM, Schaefer AJ, Roberts MS (2004) The optimal timing of living-donor liver transplantation. *Management Sci.* 50(10):1420–1430.

Bergin J (2019) DNA sequencing market: Size, trends, share & research report. Technical report, BCC Research, Wellesley, MA.

Berry DA, Iversen ES, Gudbjartsson DF, Hiller EH, Garber JE, Peshkin BN, Lerman C, et al. (2002) BRCAPRO validation, sensitivity of genetic testing of BRCA1/BRCA2, and prevalence of other breast cancer susceptibility genes. *J. Clinical Oncology* 20(11):2701–2712.

Carey DJ, Fetterolf SN, Davis FD, Faucett WA, Kirchner HL, Mirshahi U, Murray MF, Smelser DT, Gerhard GS, Ledbetter DH (2016) The Geisinger MyCode Community Health Initiative: An electronic health record-linked biobank for Precision Medicine research. *Genetics Medicine* 18(9):906–913.

Center for the Evaluation of Value and Risk in Health (2020) The cost-effectiveness analysis registry. Institute for Clinical Research and Health Policy Studies, Tufts Medical Center. Accessed April 17, 2020, www.cearegistry.org.

Chen J, Bae E, Zhang L, Hughes K, Parmigiani G, Braun D, Rebbeck TR (2020) Penetrance of breast and ovarian cancer in women who carry a BRCA1/2 mutation and do not use risk-reducing salpingo-oophorectomy: An updated meta-analysis. *JNCI Cancer Spectrum* 4(4):pkaa029.

Chen S, Wang W, Broman KW, Katki HA, Parmigiani G (2004) BayesMendel: An R environment for Mendelian risk prediction. *Statist. Appl. Genetics Molecular Biology* 3(1).

Chhatwal J, Alagoz O, Burnside ES (2010) Optimal breast biopsy decision-making based on mammographic features and demographic factors. *Oper. Res.* 58(6):1577–1591.

Cooper GF (1990) The computational complexity of probabilistic inference using Bayesian belief networks. *Artificial Intelligence* 42(2):393–405.

Daly MB, Pal T, Berry MP, Buys SS, Dickson P, Domchek SM, Elkhanany A, et al. (2021) Genetic/familial high-risk assessment: Breast, ovarian, and pancreatic, version 2.2021, NCCN clinical practice guidelines in oncology. *J. Natl. Comprehensive Cancer Network* 19(1):77–102.

Drakopoulos K, Randhawa RS (2021) Why perfect tests may not be worth waiting for: Information as a commodity. *Management Sci.* 67(11):6678–6693.

El Hajj H, Bish DR, Bish EK, Kay DM (2022) Novel pooling strategies for genetic testing, with application to newborn screening. *Management Sci.* 68(11):7994–8014.

Guzauskas GF, Garbett S, Zhou Z, Spencer SJ, Smith HS, Hao J, Hassen D, et al. (2020) Cost-effectiveness of population-wide genomic screening for hereditary breast and ovarian cancer in the United States. *JAMA Network Open* 3(10):e2022874.

Koller D, Friedman N (2009) *Probabilistic Graphical Models: Principles and Techniques* (MIT Press, Cambridge, MA).

Kong N, Zhang S, eds. (2018) *Decision Analytics and Optimization in Disease Prevention and Treatment* (Wiley, Hoboken, NJ).

Kschischang FR, Frey BJ, Loeliger H (2001) Factor graphs and the sum-product algorithm. *IEEE Trans. Inform. Theory* 47(2):498–519.

Lee HK, Li J, Bain PA, Baker C (2022) Optimal intervention policies for total joint replacement postoperative care process. *Manufacturing Service Oper. Management* 24(1):617–631.

Manahan ER, Kuerer HM, Sebastian M, Hughes KS, Boughey JC, Euhus DM, Boolbol SK, Taylor WA (2019) Consensus guidelines on genetic testing for hereditary breast cancer from the American Society of Breast Surgeons. *Ann. Surgical Oncology* 26(10):3025–3031.

Manchanda R, Legood R, Burnell M, McGuire A, Raikou M, Loggenberg K, Wardle J, et al. (2015) Cost-effectiveness of population screening for BRCA mutations in Ashkenazi Jewish women compared with family history-based testing. *J. Natl. Cancer Inst.* 107(1):380.

Manchanda R, Patel S, Gordeev VS, Antoniou AC, Smith S, Lee A, Hopper JL, et al. (2018) Cost-effectiveness of population-based BRCA1, BRCA2, RAD51C, RAD51D, BRIP1, PALB2 mutation testing in unselected general population women. *J. Natl. Cancer Inst.* 110(7):714–725.

Mayo Clinic (2021) *Center for Individualized Medicine—Genomic Sequencing: Tapestry Study* (Mayo Clinic, Rochester, MN).

Mendel G (1866) *Experiments in Plant Hybridisation* (Royal Horticultural Society, London).

Miller DT, Lee K, Chung WK, Gordon AS, Herman GE, Klein TE, Stewart DR, et al. (2021) ACMG SF v3.0 list for reporting of secondary findings in clinical exome and genome sequencing: A policy statement of the American College of Medical Genetics and Genomics (ACMG). *Genetics Medicine* 23(8):1381–1390.

Mills A, Ziya S (2021) Testing with limited capacity and pooling. Preprint, submitted March 18, https://papers.ssrn.com/abstract=3689028.

Nohdurft E, Long E, Spinler S (2017) Was Angelina Jolie right? Optimizing cancer prevention strategies among BRCA mutation carriers. *Decision Anal.* 14(3):139–169.

Oliveri S, Ferrari F, Manfrinati A, Pravettoni G (2018) A systematic review of the psychological implications of genetic testing. *Frontiers Genetics* 9:624.

Örmeci EL, Güneş ED, Kunduzcu D (2016) A modeling framework for control of preventive services. *Manufacturing Service Oper. Management* 18(2):227–244.

Paulden M (2020) Calculating and interpreting ICERs and net benefit. *PharmacoEconomics* 38(8):785–807.

Pearl J (1982) Reverend Bayes on inference engines: A distributed hierarchical approach. *Proc. Natl. Conf. Artificial Intelligence* (AAAI Press, Palo Alto, CA), 133–136.

Phillips KA, Deverka PA, Hooker GW, Douglas MP (2018) Genetic test availability and spending: Where are we now? Where are we going? *Health Affairs* 37(5):710–716.

Puterman ML (1994) *Markov Decision Processes: Discrete Stochastic Dynamic Programming* (John Wiley & Sons, Inc., New York).

Rand LZ, Kesselheim AS (2021) Controversy over using quality-adjusted life-years in cost-effectiveness analyses: A systematic literature review. *Health Affairs* 40(9):1402–1410.

Rehm HL, Bale SJ, Bayrak-Toydemir P, Berg JS, Brown KK, Deignan JL, Friez MJ, Funke BH, Hegde MR, Lyon E (2013) ACMG clinical laboratory standards for next-generation sequencing. *Genetics Medicine* 15(9):733–747.

Richards S, Aziz N, Bale S, Bick D, Das S, Gastier-Foster J, Grody WW, et al. (2015) Standards and guidelines for the interpretation of sequence variants: A joint consensus recommendation of the American College of Medical Genetics and Genomics and the Association for Molecular Pathology. *Genetics Medicine* 17(5):405–423.

Robbins R (2019) Top U.S. medical centers roll out DNA sequencing clinics for healthy clients. *Scientific American* (August 18), https://www.scientificamerican.com/article/top-u-s-medical-centers-roll-out-dna-sequencing-clinics-for-healthy-clients/.

Sturm AC, Knowles JW, Gidding SS, Ahmad ZS, Ahmed CD, Ballantyne CM, Baum SJ, et al. (2018) Clinical genetic testing for familial hypercholesterolemia: JACC scientific expert panel. *J. Amer. College Cardiology* 72(6):662–680.

Terry MB, Liao Y, Whittemore AS, Leoce N, Buchsbaum R, Zeinomar N, Dite GS, et al. (2019) 10-year performance of four models of breast cancer risk: A validation study. *Lancet Oncology* 20(4):504–517.

The 100000 Genomes Project Pilot Investigators (2021) 100,000 Genomes pilot on rare-disease diagnosis in healthcare—Preliminary report. *New England J. Medicine* 385(20):1868–1880.

Toriello A, Haskell WB, Poremba M (2014) A dynamic traveling salesman problem with stochastic arc costs. *Oper. Res.* 62(5):1107–1125.

U.S. Census Bureau (2012) Census 2000 data for the United States. Report, U.S. Census Bureau, Washington, DC.

U.S. Census Bureau (2020) Historical households tables. Report, U.S. Census Bureau, Washington, DC.

US Preventive Services Task Force, Owens DK, Davidson KW, Krist AH, Barry MJ, Cabana M, Caughey AB, et al. (2019) Risk assessment, genetic counseling, and genetic testing for BRCA-related cancer: US Preventive Services Task Force recommendation statement. *JAMA* 322(7):652.

Weiler N (2020) Free genome sequencing now available to all UCSF patients. Accessed October 11, 2020, https://www.ucsf.edu/news/2020/08/418286/free-genome-sequencing-now-available-all-ucsf-patients.