Cost-Benefit Analysis of Melanoma Screening

Team: 9078

March 12, 2022

1 Executive Summary

Healthcare plays a critical role in covering medical costs for millions of patients. In 2020, approximately 91.4% of Americans were estimated to have health insurance, accounting for 19.7% of annual GDP. With healthcare providing non-trivial coverage for a vast majority of patients in the US, it is a primary objective for private health insurance companies to maximize profit and minimize costs in the long-term. In our paper, we focused on the effects of melanoma cases on health insurance coverage.

In order to complete this task, we implemented and synthesized two models: a UV exposure model and a melanoma cost model. The UV exposure model uses UV and ozone data from the Sentinel-5 Precursor Satellite. The model forecasts future UV trends by encoding UV and ozone time series data into Gramian Angular Fields (GAFs). The resulting GAFs were fed into a Convolutional Long Short-term Memory (ConvLSTM) network which outputs predicted GAFs for future time series. By applying an inverse GAF function, the predicted GAFs were converted into future UV forecasts.

Our next model, the melanoma cost model is implemented to estimate the expected cost of melanoma for insurance companies. By designing a Markov model, the model was able to adjust for melanoma patients of different demographics. Additionally, prognosis screening for melanoma was implemented to determine whether having screening for a certain population would reduce the expenditure due to melanoma. To determine the effectiveness of screening, the model also computes the minimum rate of new melanoma cases per year. Using a transition matrix, the model tracked average patients through their progression between different health stages over time and the associated cost of their predicted melanoma progression to estimate the final expected treatment cost value.

Finally, the two models were synthesized together in order to analyze the cost-benefit metrics of melanoma screening and the effects it could potentially have on health insurance.

Currently, we estimate that each screening costs $150. By using this value for our models, we found that having a screening is not profitable for health insurance companies. This result could be attributed to multiple factors and could change in the longer term. For example, the diminishing costs of screening would improve the cost-efficiency of early diagnosis of melanoma. Additionally, with the trend of increasing UV exposure per region, cases of melanoma could rise and early screening may provide prevention of later stages of melanoma, thus lowering medical cost.

Our model predicts that the costs of screening for melanoma outweighs the treatment costs saved for a healthcare company, so screening is not directly financially beneficial. We recommend that healthcare insurance companies seek to inform both their stakeholders and the general public about easy ways to prevent and seek diagnosis for melanoma, as well as consider the role of other complex factors — such as patient welfare and peace of mind — in their overall cost-benefit analyses.

Contents

2 Introduction and Background

For private health insurers, inclusion of a disease or treatment in coverage is motivated by the expected profit [1]. A rational actor health insurance company would only insure a procedure if it would increase net revenue, motivating a cost-benefit framework for analysis.

Skin cancer is the most common form of cancer in the United States, and the American Academy of Dermatology estimates that 9,500 people in the US are diagnosed with skin cancer everyday [2] [3]. For the year 2022, the American Cancer Society predicts 99,780 new US cases of cutaneous melanoma — the deadliest form of skin cancer [4]. Over the last decades, melanoma has experienced a rapid increase in cases [4] and saw an overall US death rate of 2.5 per 100,000 between 2012–2016 [5]. Total melanoma treatment costs are approximately $3.3 billion annually, and Guy et al. 2015 estimate that the annual costs for treating melanoma among newly diagnosed cases will rise from $457 million in 2011 to $1.6 billion in 2030 [6]. For context, the 2020 revenue of the United States healthcare industry was estimated to be $4.1 trillion, accounting for 19.7% of the country's GDP [7]. A majority of the cost burden for melanoma treatments falls upon health insurance companies [8], since 91.4% of Americans are estimated to have health insurance [9]. Thus, melanoma is an increasingly pressing health concern in the US that places a significant burden on the nation's healthcare.

With respect to mortality, most detectable forms of cancer such as breast, colorectal, and lung cancers have demonstrated general global declines over the past decades [10]. Conversely, US melanoma mortality rates have consistently risen since the 1980s, while US melanoma incidence has increased exponentially since the 1930s [11]. Although new technologies and their complement clinical examination help to provide early recognition of melanoma, the rising mortality trend suggests that a clinical screening for melanoma may provide an effective advantage in early diagnosis. Additionally, a screening plan where individuals at risk of developing melanoma based on known predisposing risk factors (e.g. familial history, UV exposure, etc.) could present an opportunity to target the right individuals for effective screening. The earlier that melanoma patients are screened, the higher likelihood they have for survival, and the cheaper their treatment will cost [8], saving lives and money.

According to Jhappan et al. 2003, UV exposure and genetic factors play a critical role in the process in which melanocytes are transformed into malignant melanoma [12]. A 2010 study of melanoma cases in the UK also found that 90% of cases in men and 82% in women were attributed to "excess" solar irradiation [13]. Bharath & Turner 2009 and Parker 2021 both identify general climate change as a critical environmental and behavioral factor contributing to rising risk of melanoma [14] [15]. Ozone depletion as part of climate change leads to increased overall UV exposure risk, and warmer weather induces greater time spent outdoors while exposed to UV [14]. As such, decreases in ozone layer thickness correlate with rising incidences in melanoma. With global rises in UV exposure, Parker predicts that climate change will play an increasingly larger role in future cutaneous carcinogenesis [15].

Studies also suggest that sociodemographic factors contribute to melanoma prognosis. Race,

sex, and age have all been shown to have an influential effect on the skin cancer incident rates for a given population [16] [17] [18]. In particular, prognosis was strongly associated with neighborhood racial heterogeneity, according to melanoma cases reported to Surveillance Epidemiology and End Results (SEER) [19].

2.1 Problem Statement

The goal of our project was to analyze the cost-effectiveness of melanoma screening, given relevant prognostic factors such as demographics and UV exposure, and forecast how this cost-effectiveness evolves over time in order to minimize the expected cost for melanoma insurance coverage.

3 Data Methodology

3.1 Data Collection

We collected our melanoma data on a county-by-county basis from the CDC's State Cancer Profiles database spanning 2014–2018 [20]. For every county in the United States, we queried the average annual count of melanoma cases by race (Native American, White non-Hispanic, White Hispanic, Hispanic, African American, and Asian/Pacific Islander), by sex (male and female), and by age (<50, 50+, <65, and 65+ years old).

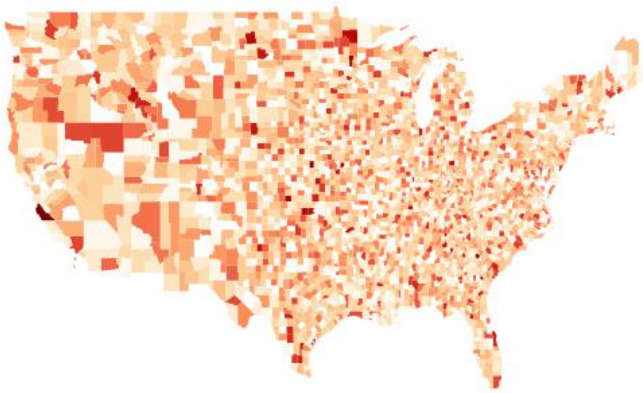

Figure 1: Mainland US color map of CDC average annual melanoma case counts by county (logarithmically-scaled), where darker colors indicate higher melanoma counts (visualization generated using GeoPandas in Python)

Via the Google Earth Engine Data Catalog, our ozone concentration and UV aerosol index data was sourced from measurements collected by the Sentinel-5 Precursor (Sentinel-5P) satellite mission [21]. In total, we collected data for each of the first 50 weeks of 2019. We restricted our queries to the contiguous United States because mainland US counties generally had more complete melanoma data. Furthermore, the climates and geographies of Hawaii, Alaska, and US territories are distinct and discontinuous from those of the mainland US, which can potentially lead to more severe outlier effects if non-contiguous US regions are included in our data collection.

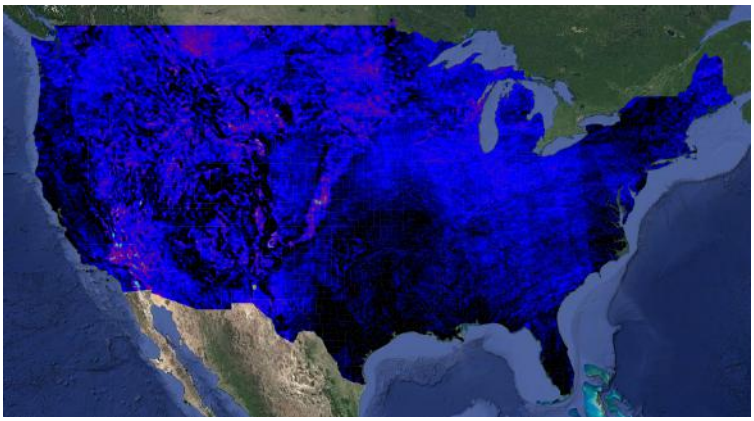

Figure 2: Mainland US heatmap of mean UV aerosol indices over the first week of 2019 (visualization generated using Google Earth Engine)

Furthermore, the COVID-19 pandemic led to extreme quarantine measures throughout the world to slow the spread of the virus, shutting down major industries around the globe, leading to ozone layer recovery [22]. This implies a temporary disruption of pre-pandemic climate trends, which could lead to biased projections of future climate trends from pandemic satellite data. Thus, we chose to restrict the time frame of our historical ozone and UV data to before the pandemic.

Additionally, we collected data to generate the transition matrix, diagnosis matrix, and cost matrix for the Markov chain model described in Section 4.2.1. A study conducted by Wu et al. 2011 described racial and ethnic variations in melanoma incidence and survival from the SEER database. Using [23], demographic-specific matrix was derived based on inputted population. Similarly, Morton et al. 2009 derives transition matrix for melanoma progression by taking account of probability of recurrence after surgery and other treatments and probability of complications from treatments [24]. Finally, to find the expectation value of the cost health insurance have to cover, a metastudy by Guy et al. 2012 provides annual costs of treatment for local, regional, and distant metastatic melanoma [8]. Therefore, the data we assembled quantify the severity and frequency of different stages and outcomes of melanoma disease progression.

3.2 Data Interpolation

In the CDC cancer database, many counties' average annual melanoma counts for specific demographics were unavailable due to data collection and data confidentiality restrictions [20]. Thus, we used a limited nearest-neighbor interpolation for each county, because a majority of the missing data belonged to rural counties, which tend to have similar demographics to those of their neighbors [25] [26]. Furthermore, we assumed immediately adjacent counties share approximately similar geographic and climatic characteristics, which would imply roughly similar ozone and UV levels, which would contribute to similar melanoma incident rates (see Assumption 4 of UV exposure model).

We interpolated each county and demographic's missing melanoma counts by taking the mean of melanoma counts that were collected and/or interpolated for that demographic in the immediately adjacent counties. It was necessary to interpolate these missing numbers so that we could supply our UV exposure model with sufficiently complete training data.

4 Mathematics Methodology

4.1 UV Exposure Model

4.1.1 Assumptions

Assumption 1: Autocorrelations between UV correlations are significant for prediction.

- **Justification:** The purpose of the Gramian Angular Field (GAF) submodel is that these autocorrelations within the GAF images are reflective of some deeper trend, such as seasonal variation, rather than just being noise or a stochastic model.

Assumption 2: UV trends can be predicted from one year of UV data.

- **Justification:** Due to the limits of the data that we have available, we only used UV trends from the year 2019 to train our model.

Assumption 3: UV can be aggregated on the county level.

- **Justification**: Data on melanoma incidence was obtained on a per-county basis. To formulate our GAF-ConvLSTM model, we created a time series of mean UV for each county, without encoding spatial variation in UV within a county. We thus assume that UV exposure is locally homogeneous within the spatial scale of a county.

Assumption 4: UV exposure is a main risk factor for melanoma occurrence.

- **Justification**: According to a research conducted by [13] and [27], UV exposure is the main risk factor for melanoma.

4.1.2 Model Objectives

The depletion of the ozone layer as a consequence of modern industrial activities is of great concern to the healthcare industry as a risk factor. A depletion in ozone, in combination with a multitude of natural and man-made environmental, influences UV exposure. The ozone levels and UV aerosol indices from the Sentinel-5P satellite were only collected from a time range between July 2018 to the present [21]. Due to these temporal limitations, it is necessary to forecast UV indices in order to model future trends in environmentally-induced melanoma cases in the US and project our overall melanoma cost-effectiveness model.

Since these variables vary spatially across different latitudes and climates, as well as temporally, the prediction of UV exposure many years into the future is a complex task. Our model needed to learn spatial trends across the United States, temporal trends across successive measurements from satellite observations, and predict future UV exposure trends in the US over multiple years.

4.1.3 GAF Submodel

To forecast the UV data into future trends, we first encoded the Sentinel-5P satellite time-series data for UV and ozone levels into a Gramian Angular Field (GAF) image for each county. In the original proposal of GAFs by Wang & Oates 2015, it was found that training convolutional neural networks on time-series data encoded as GAFs yielded successful capture multi-level features and competitive classification results [28].

GAFs are images which represent time series in the polar coordinate system. Due to the nature of the polar coordinate system, each point is determined by a distance to a reference point and the angle of the reference direction, a GAF represents a temporal correlation between each element of the time series. According to Wang & Oates, conversion of the time series $\boldsymbol{X} = \{x_1, x_2, \ldots, x_N\}$ into GAF image data first requires $\boldsymbol{X}$ to be normalized as $\tilde{\boldsymbol{X}} = \{\tilde{x}_1, \tilde{x}_2, \ldots, \tilde{x}_N\}$ where:

$$\tilde{x}_i = \frac{(x_i - \max(\boldsymbol{X})) + (x_i - \min(\boldsymbol{X}))}{\max(\boldsymbol{X}) - \min(\boldsymbol{X})} \tag{1}$$

The rescaled time series $\tilde{\boldsymbol{X}}$ is transformed into polar coordinates through the method below:

$$\begin{cases} \phi_i = \arccos(\tilde{x}_i), & -1 \leq \tilde{x}_i \leq 1,\ \tilde{x}_i \in \tilde{\boldsymbol{X}}, \\ r_i = \dfrac{t_i}{R},\ t_i \in \mathbb{N} \end{cases} \tag{2}$$

where R is a regularizing constant factor which normalizes the time series to [0, 1]. After

the polar transformation, the N-by-N Gramian Angular Field $\boldsymbol{G}$ is computed as:

$$\boldsymbol{G} = \begin{pmatrix} \cos(\phi_1 + \phi_1) & \cdots & \cos(\phi_1 + \phi_N) \\ \vdots & \ddots & \vdots \\ \cos(\phi_N + \phi_1) & \cdots & \cos(\phi_N + \phi_N) \end{pmatrix} \tag{3}$$

The benefit of the GAF expression of a time series is that it encodes autocorrelations. Whereas a traditional time series is treated as a Markov series, in that each point is dependent only on its previous points, autocorrelations between points on different part of the trends encode broader variations, such as seasonal or macroscopic fluctuation.

4.1.4 ConvLSTM Submodel

To then generate GAF images predicting future time steps, we fed our GAFs into a Convolutional Long Short-term Memory network (ConvLSTM). The traditional LSTM network is widely used in the literature to forecast one-dimensional time-series data, with applications ranging from tumor growth prediction [29] to drug design [30] to predicting solar radiation from satellite imagery [31]. The LSTM architecture is a neural network based on the following procedure:

Algorithm 1 Training Algorithm for LSTM

$t \leftarrow \alpha$ where α is the memory size
Initialize f neural network
while $N < 50$ **do**
    Train f with $f(\alpha)$
    $t_{N+1} \leftarrow t_N + 1$
    $\alpha \leftarrow t_{N+1}$
    Predict $f(\alpha)$
    Compute $MSE = \frac{1}{M}\sum(\hat{f} - f)^2$
    Compute $\frac{\partial MSE}{\partial \theta}$ where θ are parameters of f
    Update θ

However, when applied to our original time-series data, the LSTM failed to yield an improving loss curve during training, as seen in Figure 3.

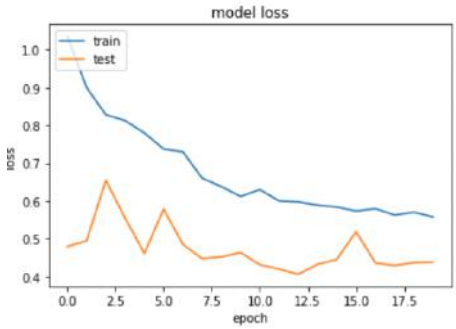

Figure 3: Loss curve when training and testing a traditional LSTM model on one-dimensional UV time-series data

The ConvLSTM architecture enables the LSTM to operate on higher dimensional (spatial) data. It merges the LSTM with the convolutional neural network, which learns features from images through the use of the image convolution operator. GAFs have successfully been used in conjunction with ConvLSTMs for both time-series classification [32] and forecasting tasks [33]. In Hong et al. 2020, it was found that GAFs helped overcome ConvLSTM limitations in one-dimensional forecasting problems [33], and our model seeks to avoid these same limitations. The ConvLSTM learns not only on the time series itself, but also on the autocorrelations encoded in the GAF to provide better predictions of future time series values. This approach yields significant comparative advantages over static time series forecasting approaches, such as Autoregressive Integrated Moving Average (ARIMA) or Vector Autoregression Moving-Average with Exogenous Regressors (VARMAX).

For our model, we employed a series of 2D convolutional LSTM layers with batch normalization steps in between. The architecture is shown in Figure 4.

Layer (type)	Output Shape	Param #
conv_lst_m2d_1 (ConvLSTM2D)	(None, None, 25, 25, 350)	17691800
batch_normalization_1 (Batch	(None, None, 25, 25, 350)	1400
conv_lst_m2d_2 (ConvLSTM2D)	(None, None, 25, 25, 200)	7040800
batch_normalization_2 (Batch	(None, None, 25, 25, 200)	800
conv_lst_m2d_3 (ConvLSTM2D)	(None, None, 25, 25, 200)	2880800
batch_normalization_3 (Batch	(None, None, 25, 25, 200)	800
conv_lst_m2d_4 (ConvLSTM2D)	(None, None, 25, 25, 150)	840600
batch_normalization_4 (Batch	(None, None, 25, 25, 150)	600
conv3d_1 (Conv3D)	(None, None, 25, 25, 1)	18751

Total params: 28,476,351
Trainable params: 28,474,551
Non-trainable params: 1,800

Figure 4: Trainable parameters in ConvLSTM

The optimizer we used was the ADADELTA optimizer, which is a gradient descent method that utilizes adaptive learning rates [34]. The loss functions we used were mean squared error (MSE) and mean absolute error (MAE). These are defined mathematically as follows:

$$\begin{cases} MAE = \dfrac{\sum_{i=1}^{n} |y_i - \hat{y}_i|}{n} \\ MSE = \dfrac{\sum_{i=1}^{n} (y_i - \hat{y}_i)^2}{n} \end{cases} \tag{4}$$

4.1.5 Application of Model

To prepare inputs to the GAF-ConvLSTM model, we first generated raster images of UV and ozone data for the mainland United States. We then found the mean UV and ozone level for the data points within each county at a given time step, denoted $\bar{A}$ and $\bar{B}$ respectively. This generated a two-channel time series of $\bar{A}_t$ and $\bar{B}_t$ at each time step from $t = 1, \ldots, 50$,

denoted as y_t and y_t^*, where time step t represented week t of the 2019 satellite data. We normalized the time series to prepare them for GAF creation. The normalization equation for a time series is described below, where α is the normalization scaling factor:

$$\hat{y_t} = \frac{y_t - \min(y_t)}{\max(y_t) - \min(y_t)} = (y_t - \min(y_t)) * \alpha \tag{5}$$

For each time series for each county, we generated two-dimensional GAF images. We began by specifying a sampling size and stride size, denoted as s_{samp} and s_{str} respectively. The sampling size is the image size of the resulting GAF, and the stride length is the increment between steps. We then generated the evolution sequence of the GAFs for each county, denoted as κ, using $N = 50$ as the number of time steps to subsample over.

$$\kappa = \left\{ \boldsymbol{G}\Big(\{\tilde{x}_{i \cdot s_{str}+j},\ j = 0 \dots (s_{samp} - 1)\} \Big),\ i = 1 \dots \left\lfloor \frac{N - s_{samp} + 1}{s_{str}} \right\rfloor \right\} \tag{6}$$

κ contains a sequence of $\left\lfloor \frac{N - s_{samp} + 1}{s_{str}} \right\rfloor$ GAFs which can be used to train the ConvLSTM network. The ConvLSTM network needs to predict the GAF at time $t_{new} = t + 1$ given the GAF at time t. κ is then the input to the ConvLSTM model for each of the total 3103 counties, after filtering out counties which are outside the US mainland or have zero GAF (all zero values).

After performing predictions with our model, the result was a predicted GAF matrix at some time in the future denoted t^*. To map this back to a numerical value on a time series we apply the inverse GAF transform (invGAF). The GAF transformation is bijective, so a one-to-one mapping exists from any GAF to a time series. The diagonal of the GAF matrix encodes the exact angle of each time step ϕ_i, as shown in equation 3. Specifically the angle can be computed as follows.

$$\phi_i = \frac{\arccos \boldsymbol{G}_{i,i}}{2} \text{ for } i = 1 \dots N \tag{7}$$

For computing the radius, a key insight comes from observing that the normalizing constant is the regularizing constant factor, or $R = \alpha$. Using this, we can compute the radius of each point.

$$r_i = \frac{t}{\alpha} \tag{8}$$

We can therefore calculate the exact time series values for each predicted GAF using the invGAF. We can then undo the normalization to map it back to its original scale by reversing the normalization step to get our final predicted value.

$$\hat{y} = r \cos \phi * \frac{1}{\kappa} \tag{9}$$

4.2 Melanoma Cost Model

4.2.1 Assumptions

Assumption 1: The elements of our transition matrix can have a zero probability.

- **Justification:** We assume in our transition matrix that certain transitions between states are not possible (i.e. have a 0% frequency). We make this assumption largely because the paper from which we derive the transition matrix $\boldsymbol{P}$ makes the same assumption [24].

Assumption 2: Our transition matrices and the transitions are Markovian.

- **Justification:** In our transition matrix, we assume that the transitions are Markovian: that is, the transition matrix is constant regardless of how many time steps (years) have passed. We needed this assumption because detailed data on patient transitions after multiple years is not available to the public. This assumption is reasonable because as time increases and the number of transitions increases, the probability of both death stages become much larger than any other stage in the exponentiated matrix $\boldsymbol{P}^t$.

  For instance, consider $\boldsymbol{P}^{20}$, the matrix of transition frequencies between year t and year $t + 20$.

$$\boldsymbol{P}^{20} = \begin{pmatrix} 54.5\% & 0.9\% & 1.9\% & 3.5\% & 33.0\% & 4.7\% \\ 54.5\% & 0.9\% & 1.9\% & 3.5\% & 33.0\% & 4.7\% \\ 41.5\% & 0.7\% & 1.5\% & 2.7\% & 48.9\% & 3.6\% \\ 0.0\% & 0.0\% & 0.0\% & 0.0\% & 100.0\% & 0.0\% \\ 0.0\% & 0.0\% & 0.0\% & 0.0\% & 100.0\% & 0.0\% \\ 0.0\% & 0.0\% & 0.0\% & 0.0\% & 0.0\% & 100.0\% \end{pmatrix} \tag{10}$$

  As expected, the 2nd, 3rd, and 4th columns which correspond to having melanoma after 20 years have very low frequencies, less than 5% each. In comparison, there is a much higher chance of ending up dead from melanoma or ending up melanoma-free.

- The Markovian approach is common in modeling health processes: for instance, see [24] [35] [36].

- For a real insurance company, this assumption would not be strictly necessary: they would have more comprehensive data that would allow them to alter the transition model as time goes on.

Assumption 3: The same demographic population shares the same diagnosis matrix.

- **Justification:** We assume that any patient from demographic i has the same diagnosis frequency vector $\boldsymbol{D}_i$. Since we only have demographic vectors for different races, we essentially assume that every patient with the same race has the same diagnosis matrix. This is a reasonable assumption because skin color is one of the largest factors in melanoma detection: it is harder to note melanoma's characteristic skin discoloration in darker skin [37].

Assumption 4: Treatment cost vector is independent of the patient's demographics.

- **Justification:** We assume that every patient, no matter their demographic or any other factor, has the same costs of treatment per stage every year. This assumption is probably not too problematic — patients have to undergo similar treatment, which cost similar amounts. Additionally, since we're only estimating the average expected cost of melanoma, we just need our cost vector to be close to the real-world average annual costs.

4.2.2 Model Objectives

The primary objectives of our cost model are to estimate the expected cost of melanoma (and by proxy, the potential savings of melanoma screening) to insurance companies. For this, we design a model to estimate the expected costs incurred through the treatment of a typical melanoma patient, with sufficient customizability to adjust for different demographics and trends in time.

The goals of our cost model are twofold: first, to determine the cost of treatment for a given melanoma patient from time of diagnosis, and second to determine how much money an insurance company might save if they provide precautionary screening for catching melanoma at an earlier stage. The latter goal can be used to determine, in conjunction with demographic data (age, gender, race, geographic location), what demographics are worth providing melanoma screening for.

We accomplish both goals with a transition matrix model that tracks the average patient through their treatment, through different possible stages of melanoma that are associated with a cost of treatment.

The primary advantage of using matrices is to allow us to efficiently consider every possible patient journey (the various stages of disease they go through during their treatment) and their associated costs, and then multiply and sum them all up to arrive at a final expected value.

Since a patient's condition can improve or deteriorate at every time step randomly, the total number possible patient journeys grows exponentially with the number of time steps considered. It would be monstrously inefficient and computationally infeasible to independently consider every possible patient journey. In contrast, our matrix model allows us to efficiently calculate the expected costs of patient journeys with only a few discretizing and common assumptions in Section 4.2.1.

Our model takes as input an initial vector of the frequencies at which different stages of melanoma are diagnosed and outputs the expected cost of treatment over a patient's total treatment cycle. The stage frequency vector is highly customizable and can be altered for different demographics or for the presence of screening procedures.

4.2.3 Transition Matrix Model

Our transition model is based on a Markov chain, which provides the annual transition probabilities between stages of melanoma. The stages we track are disease-free, localized metastases, regional metastases, distant metastases, death from melanoma, and death from other causes. A flowchart of the Markov chain is shown in Figure 5.

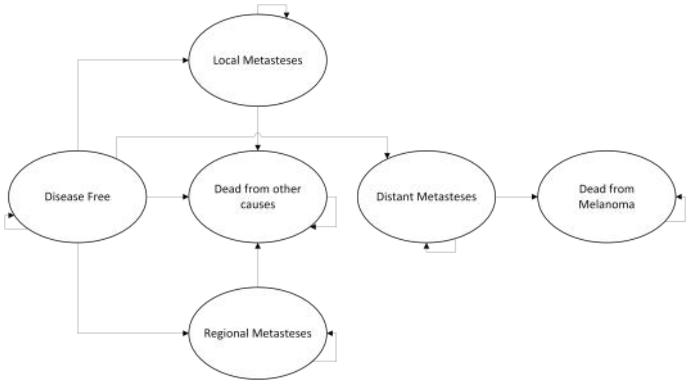

Figure 5: Flowchart representation of the Markov chain for stages of melanoma

The transition matrix of the Markov chain represents the pairwise probability that a patient starting in one stage will end up in another stage after one time step, which is one year in our model. Thus, the matrix is a 6×6 matrix of probabilities, with the value in the mth row and nth column being the probability that a patient transitions from the mth stage to the nth stage. As a consequence, the probabilities in each row will sum to 1. The transitions are shown in Table 1, and they come from [24].

	Disease free	**Local metastases**	**Nodal metastases**
Disease Free	93.1%	1.6%	3.3%
Local metastases	93.2%	1.5%	3.4%
Nodal metastases	72.0%	0.0%	2.8%
Distant metastases	0.0%	0.0%	0.0%
Dead melanoma	0.0%	0.0%	0.0%
Dead other causes	0.0%	0.0%	0.0%

	Distant metastases	**Dead melanoma**	**Dead other causes**
Disease Free	1.6%	0.0%	0.3%
Local metastases	1.6%	0.0%	0.3%
Nodal metastases	24.9%	0.0%	0.3%
Distant metastases	58.2%	41.8%	0.0%
Dead melanoma	0.0%	100.0%	0.0%
Dead other causes	0.0%	0.0%	100.0%

Table 1: Annual Transition Probabilities between Melanoma States

We can transform Table 1 into our transition matrix for the Markov chain model. We denote $\boldsymbol{P}$ as the transition matrix. $\boldsymbol{P}$ is shown below:

$$\boldsymbol{P} = \begin{pmatrix} 93.1\% & 1.6\% & 3.3\% & 1.6\% & 0.0\% & 0.3\% \\ 93.2\% & 1.5\% & 3.4\% & 1.6\% & 0.0\% & 0.3\% \\ 72.0\% & 0.0\% & 2.8\% & 24.9\% & 0.0\% & 0.3\% \\ 0.0\% & 0.0\% & 0.0\% & 58.2\% & 41.8\% & 0.0\% \\ 0.0\% & 0.0\% & 0.0\% & 0.0\% & 100.0\% & 0.0\% \\ 0.0\% & 0.0\% & 0.0\% & 0.0\% & 0.0\% & 100.0\% \end{pmatrix} \tag{11}$$

The benefit of creating a transition matrix is that we can easily find the probability that a patient transitions between two stages after multiple years. The transition probabilities between stages after t years is given by $\boldsymbol{P}^t$ [38].

For example, after 3 years, the transition probabilities are given by $\boldsymbol{P}^3$:

$$\boldsymbol{P}^3 = \begin{pmatrix} 88.0\% & 1.5\% & 3.1\% & 4.2\% & 2.0\% & 0.9\% \\ 88.1\% & 1.5\% & 3.1\% & 4.2\% & 2.0\% & 0.9\% \\ 67.1\% & 1.1\% & 2.4\% & 11.2\% & 17.2\% & 0.7\% \\ 0.0\% & 0.0\% & 0.0\% & 19.7\% & 80.3\% & 0.0\% \\ 0.0\% & 0.0\% & 0.0\% & 0.0\% & 100.0\% & 0.0\% \\ 0.0\% & 0.0\% & 0.0\% & 0.0\% & 0.0\% & 100.0\% \end{pmatrix} \tag{12}$$

The probability that a patient starts with regional metastases and ends up dead from melanoma in 3 years is 17.2% (row 3, column 5).

4.2.4 Stage Diagnosis Frequency Vector

We also consider a vector $\boldsymbol{D}$ of the frequency at which different stages of melanoma are diagnosed. This vector has 6 dimensions, but since the frequencies are after melanoma has been diagnosed, the frequency of no disease is 0% and the frequencies of both death stages are 0%. Diagnosis of melanoma is assumed to be the start of treatment and therefore melanoma costs.

Different demographics have different stage diagnosis frequencies, and the diagnosis frequencies can also change if screening is in place. Screening might catch a case of localized melanoma that otherwise would have gone unnoticed until distant metastases. In addition, different demographics self-diagnose melanoma differently: for example, Black populations tend to self-diagnose melanoma at a later stage than White Hispanic populations [39], and it is reflected in their two diagnosis vectors, $\boldsymbol{D}_b$ and $\boldsymbol{D}_{wh}$ respectively:

$$\begin{cases} \boldsymbol{D}_b = (0, 66.0\%, 19.3\%, 14.7\%, 0, 0) \\ \boldsymbol{D}_{wh} = (0, 78.6\%, 14.0\%, 7.5\%, 0, 0) \end{cases} \tag{13}$$

The diagnosis vectors of other demographics is provided in section 4.2.6

Let $\boldsymbol{D}$ denote some demographic's stage diagnosis frequency vector. Then, we can use matrix

multiplication: $\boldsymbol{D}\ \boldsymbol{P}^t$ represents the expected frequency of the different stages of melanoma after t years of treatment.

4.2.5 Calculating Costs

Finally, to get the expected cost of a melanoma patient, we consider a cost vector $\boldsymbol{C}$ of the costs associated with the different stages of melanoma. We get the values for each element of C from [8].

$$\boldsymbol{C} = (0, \$25,242, \$32,847, \$57,860, 0, 0) \tag{14}$$

The expected patient cost after i years is:

$$c_i = \boldsymbol{D}\ \boldsymbol{P}^i\ \boldsymbol{C}^T \tag{15}$$

where we use $\boldsymbol{C}^T$ as the transpose of $\boldsymbol{C}$ for the matrix dimensions to work for multiplication. The output of the matrix multiplication, c_i, is a scalar value.

Therefore, the total cost is:

$$c_{tot} = \sum_{i=1}^{\infty} c_i \tag{16}$$

The transition matrix over n years tends primarily to deaths as n grows large, which is expected because after enough years the patient will almost certainly have died from either melanoma or natural causes. Therefore, we do not have to evaluate the sum all the way to ∞: instead, we cut off the sum at $t = 20$. The cutoff is reasonable because it is very unlikely that the patient survives and continues to have melanoma 20 years later.

To demonstrate, $\boldsymbol{P}^{20}$ is as follows:

$$\boldsymbol{P}^{20} = \begin{pmatrix} 54.5\% & 0.9\% & 1.9\% & 3.5\% & 33.0\% & 4.7\% \\ 54.5\% & 0.9\% & 1.9\% & 3.5\% & 33.0\% & 4.7\% \\ 41.5\% & 0.7\% & 1.5\% & 2.7\% & 48.9\% & 3.6\% \\ 0.0\% & 0.0\% & 0.0\% & 0.0\% & 100.0\% & 0.0\% \\ 0.0\% & 0.0\% & 0.0\% & 0.0\% & 100.0\% & 0.0\% \\ 0.0\% & 0.0\% & 0.0\% & 0.0\% & 0.0\% & 100.0\% \end{pmatrix} \tag{17}$$

As expected, the 2nd, 3rd, and 4th columns which correspond to having melanoma after 20 years have very low frequencies. In comparison, there is a much higher chance of ending up dead from melanoma or ending up melanoma-free, in column 1 or column 5.

We assume that the transition matrix $\boldsymbol{P}$ and the cost vector $\boldsymbol{C}$ are invariant for different demographics and for the presence of screening procedures. Therefore, we express the total expected cost c_{tot} as a function of the diagnosis vector:

$$c_{tot}(\boldsymbol{D}) \approx \sum_{i=1}^{20} \boldsymbol{D}\ \boldsymbol{P}^i\ \boldsymbol{C}^T \tag{18}$$

4.2.6 Real World Stage Diagnosis Frequencies

We get the stage diagnosis frequency vectors $\boldsymbol{D}_i$ for a given demographic i from real world data in [23]. The source provides data on Black, American Indian/Alaskan Native, Asian/Pacific Islander, non-Hispanic White, and White Hispanic populations. We denote each demographic's frequency vector as $\boldsymbol{D}_b$, $\boldsymbol{D}_{ai}$, $\boldsymbol{D}_{api}$, $\boldsymbol{D}_{wnh}$, and $\boldsymbol{D}_{wh}$ respectively.

[23] provides the number of patients diagnosed with melanoma in the localized, regional, and distant cancer stages. There are also diagnoses with unknown staging due to incomplete data, so for our purposes we ignore the unknown diagnoses (we assume the unknown diagnoses have the same diagnosis distribution as the known diagnoses). The frequency of no disease is 0% because we assume melanoma is present at diagnosis, and the frequency of both death states at diagnosis is 0%.

Here are the frequency vectors:

$$\begin{cases} \boldsymbol{D}_b & = (0.0\%, 66.0\%, 19.3\%, 14.8\%, 0.0\%, 0.0\%) \\ \boldsymbol{D}_{ai} & = (0.0\%, 73.7\%, 17.8\%, 8.5\%, 0.0\%, 0.0\%) \\ \boldsymbol{D}_{api} & = (0.0\%, 72.1\%, 17.0\%, 11.0\%, 0.0\%, 0.0\%) \\ \boldsymbol{D}_{wnh} & = (0.0\%, 87.1\%, 9.0\%, 3.9\%, 0.0\%, 0.0\%) \\ \boldsymbol{D}_{wh} & = (0.0\%, 78.6\%, 14.0\%, 7.5\%, 0.0\%, 0.0\%) \end{cases} \tag{19}$$

4.2.7 Effect of Screening

The stage diagnosis frequency vectors change when screening procedures are in place because screening catches melanoma at an earlier stage. We assume that the frequency vectors provided above ($\boldsymbol{D}_b$, $\boldsymbol{D}_{ai}$, etc.) are the frequency vectors without screening procedures in place.

Then, we introduce $S(\boldsymbol{D}_0)$, a function of the non-screening diagnosis vector $\boldsymbol{D}_0$ that outputs the diagnosis vector if screening was in place. We seek to approximate $S(\boldsymbol{D}_0)$.

There is limited data on screening's effectiveness at catching melanoma early. We use [40], which states: "The percentage of early stage 1 melanoma increased from 52% in the pre-screening period to 64% in the actual screening period."

Early stage 1 melanoma corresponds to local metastases in our model, and an increase from 52% to 64% represents a relative increase of 23%. Thus, in our model we assume that the frequency of local metastases increases by 23%. In the absence of other data, we assume the frequencies of regional and distant metastases decrease by the same proportion so that their ratio is maintained and the sum of the frequencies remains 100%. Therefore, we approximate:

$$S(\boldsymbol{D}_0) \approx \left(0.0\%,\ 1.23 \cdot \boldsymbol{D}_0[1],\ \frac{1 - 1.23 \cdot \boldsymbol{D}_0[1]}{1 - \boldsymbol{D}_0[1]} \cdot \boldsymbol{D}_0[2],\ \frac{1 - 1.23 \cdot \boldsymbol{D}_0[1]}{1 - \boldsymbol{D}_0[1]} \cdot \boldsymbol{D}_0[3],\ 0.0\%,\ 0.0\%\right) \tag{20}$$

where $\boldsymbol{D}_0[1]$ is the frequency of localized melanoma, $\boldsymbol{D}_0[2]$ is the frequency of regional melanoma, and $\boldsymbol{D}_0[3]$ is the frequency of distant melanoma.

With our cost model, we can calculate the savings of screening. $c_{tot}(\boldsymbol{D}_0)$ represents the average cost for a melanoma patient without screening, and $c_{tot}(S(\boldsymbol{D}_0))$ represents the average cost with screening. Therefore, insurance companies save on average

$$c_{tot}(\boldsymbol{D}_0) - c_{tot}(S(\boldsymbol{D}_0)) \tag{21}$$

for every patient they have screened who has melanoma, not accounting for the cost of screening.

To consider if the cost of melanoma outweighs the cost of screening, we consider the overall costs an insurance company would incur if they chose to screen demographic i with diagnosis vector $\boldsymbol{D}_i$, population size n_i, and rate of annual melanoma occurrence r_i (so r_i of the overall population n_i are expected to develop melanoma every year).

Let c_s represent the cost of screening a single person once in a year. Then, the total cost of screening the entire demographic i for a year is $n_i \cdot c_s$. The total number of people expected to develop melanoma that year is $n_i \cdot r_i$, so the total cost saved by screening is $n_i \cdot r_i \cdot [c_{tot}(\boldsymbol{D}_i) - c_{tot}(S(\boldsymbol{D}_i))]$.

Therefore, for screening to be economically feasible to insurance companies, the following inequality must be satisfied:

$$n_i \cdot c_s \leq n_i \cdot r_i \cdot [c_{tot}(\boldsymbol{D}_i) - c_{tot}(S(\boldsymbol{D}_i))] \tag{22}$$

which implies:

$$r_i \geq \frac{c_s}{c_{tot}(\boldsymbol{D}_i) - c_{tot}(S(\boldsymbol{D}_i)} \tag{23}$$

Our model allows insurance companies to calculate a minimum bound on the rate of new cases of melanoma per year for screening to be economically feasible.

4.2.8 Examples with Real Data

We demonstrate each facet of our cost model with our diagnosis frequency vectors $\boldsymbol{D}_b$, $\boldsymbol{D}_{ai}$, $\boldsymbol{D}_{api}$, $\boldsymbol{D}_{wnh}$, and $\boldsymbol{D}_{wh}$, our transition matrix $\boldsymbol{P}$, and our cost vector $\boldsymbol{C}$. Additionally, we use a cost of screening c_s of \$150 as per [41], who estimate that office visits cost \$105 and biopsies cost \$45 on average for a combined \$150.

We can calculate the average cost that a melanoma patient in each demographic would incur in treatment, the cost they would have incurred if they were screened for melanoma before knowing they had melanoma, and the minimum annual threshold for the rate of new melanoma patients per year in that demographic. The values are presented in Table 2.

Demographic	Cost without Screening	Cost with Screening	Amount Saved	Minimum Annual Melanoma Rate
Black	$104,420.34	$99,350.52	$5,069.81	3.0%
American Indian	$101,213.95	$95,952.79	$5,261.16	2.9%
Asian/Pacific Islander	$102,142.16	$96,747.20	$5,394.96	2.8%
White Non-Hispanic	$97,019.65	$90,886.98	$6,132.68	2.4%
White Hispanic	$99,830.33	$94,119.70	$5,710.64	2.6%

Table 2: Table of the costs of melanoma and savings from screening, as well as the minimum rates for new melanoma cases per year for economic feasibility

We find that, as expected, the cost of treating a melanoma patient decreases on average when people are screened for melanoma. However, the savings are not significant, hovering around $5,000–$6,000. Accordingly, the minimum annual melanoma rate threshold for economic feasibility, which is inversely proportional to the savings, is extremely high at 2%–3% of the population becoming melanoma patients every year.

4.3 Synthesis Model

We sought to synthesize the two aforementioned models to generate cost-benefit metrics of screenings for the current melanoma landscape, and forecast how this relationship could evolve in the future based on UV exposure trends.

The Markov model generated critical (minimum annual) melanoma rates, r^*, for each demographic subcategory. If melanoma cases occur more frequently than r^*, then it is beneficial for the healthcare company to offer melanoma screenings for that demographic, and vice versa. Using the five-year average melanoma incidence from the Census Bureau, we can compute melanoma rates for each county for each possible demographic, and compare this to r^*. The necessity of screenings varies greatly between different regions of the United States, based on demographic distribution and melanoma rates.

4.4 Model Evaluation

4.4.1 UV Forecast Model Results

Prediction of UV trends turns out to be an extremely challenging task. Even GAF-ConvLSTM networks with more than 100 million parameters could not achieve robust performance on a 15% subset of the data. We were sucessfully able to encode the time series as GAFs. The example Gramian transform for a time series of average UV data for a sample county is shown below in Figure 6.

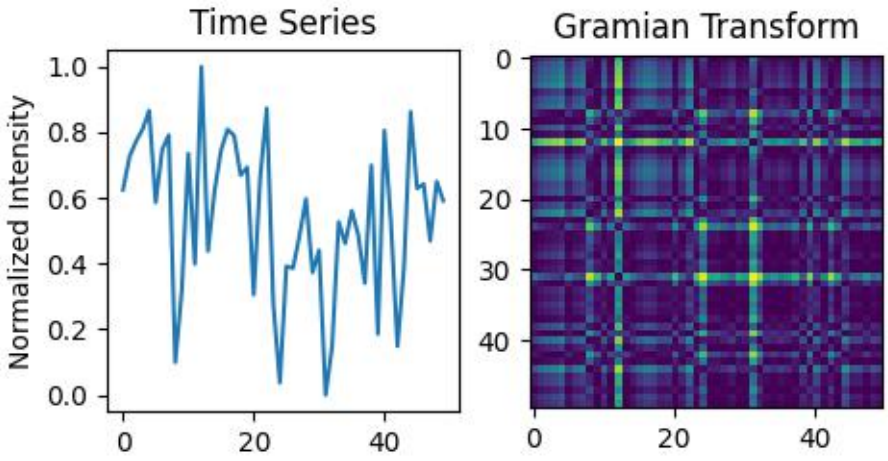

Figure 6: GAF with sampling size 10 and stride 3

From observation of time series alone, it is apparent that the data is not well behaved. There is no obvious trend in UV measurements over time and the GAF as a result is scattered.

After encoding time series into GAF progressions, we performed testing to identify which combination of (s_{samp}, s_{str}) offered the lowest mean squared error (MSE) and mean absolute error (MAE) after 50 epochs of training.

(s_{samp}, s_{str})	MSE	MAE
(10,3)	0.5683	0.6249
(15,3)	0.5538	0.6470
(25,2)	0.5329	0.6388
(25,1)	0.5211	0.5941

Table 3: Table of GAF hyperparameter tuning

Based on Table 3, we chose s_{samp} as 25 and s_{str} as 1. After training our model, we forecasted UV values 5 years into the future for a random sample of 51 counties to get $\boldsymbol{A}$, where $\boldsymbol{\mu}$ is a vector of mean UV exposure indices projected for 2024 and $\boldsymbol{\nu}$ is a vector of mean UV exposure indices collected in 2019:

$$\begin{cases} \boldsymbol{A}_i = \dfrac{\boldsymbol{\mu}_i}{\boldsymbol{\nu}_i}, \quad 1 \leq i \leq 51 \\ \boldsymbol{A} = \{1.4815, 0.8762, 0.8866, \ldots, 1.3733, 1.1947, 0.9540\} \end{cases} \tag{24}$$

Based on Figure 7, the model performed fairly well at predicting the forecasted value, which is the value in the bottom right position of the GAF. For the rest of the matrix, the model predicted the average behavior of the UV exposure GAFs across 3000 samples. The highly stochastic behavior of the evolution of the GAF of UV means that the model could not learn

on the exact features of the GAF given the time and computational power given. Instead, it predicted the average behavior for the samples.

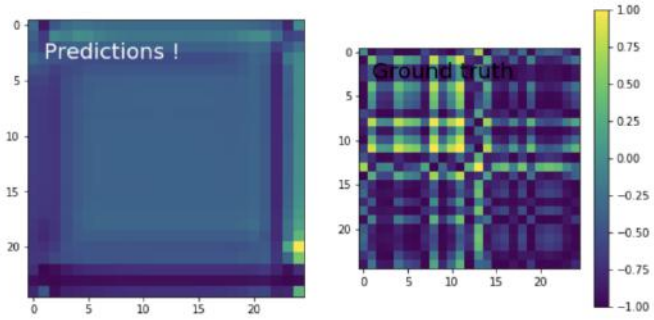

Figure 7: Comparison of predictions to ground truth for a sample data point

In reality, trends in UV are incredibly difficult to learn on. The multitude of factors that influence UV and ozone are effectively random in nature, and even sophisticated time series architecture cannot determine the exact variations of these trends.

4.4.2 Cost Model Results

Sensitivity training by varying the total number of years predicted in the cost model yielded the predicted screening rate thresholds in Figure 8.

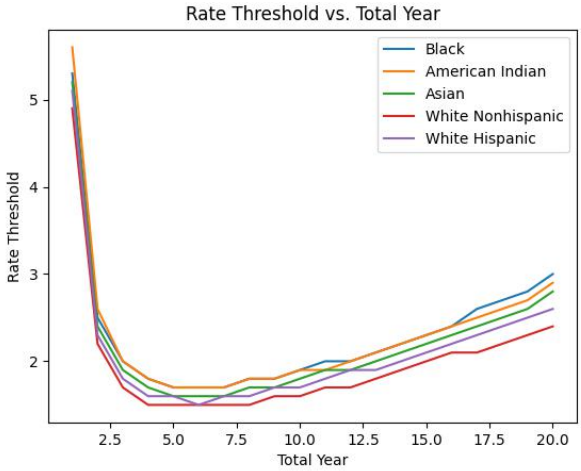

Figure 8: Curve of rate threshold vs. total number of years modeled, broken down by race demographics

We find that the minimum rate threshold is given by a 5-year cost model, and that predicted rate thresholds increase as the model's total years increase. These results indicate that our cost model is unstable and sensitive to the total year parameter.

4.4.3 Sensitivity of Stage Frequency Vectors

We calculate the sensitivity of our model to different D vectors, the stage diagnosis frequency vectors. Instead of perturbing an arbitrary D vector, we compare the total costs $c_{tot}(D)$ for each demographic diagnosis vector (described in Section 4.2.6). The values are provided in Table 4.

Demographic	Vector	Total Cost
Black	(0.0%, 66.0%, 19.3%, 14.8%, 0.0%, 0.0%)	$104,420.34
American Indian	(0.0%, 73.7%, 17.8%, 8.5%, 0.0%, 0.0%)	$101,213.95
Asian/Pacific Islander	(0.0%, 72.1%, 17.0%, 11.0%, 0.0%, 0.0%)	$102,142.16
White Non-Hispanic	(0.0%, 87.1%, 9.0%, 3.9%, 0.0%, 0.0%)	$97,019.65
White Hispanic	(0.0%, 78.6%, 14.0%, 7.5%, 0.0%, 0.0%)	$99,830.33

Table 4: Table for Sensitivity Analysis of Stage Frequency Vectors

We expect that with a higher proportion of melanoma cases detected in lower stages, the overall cost per melanoma patient would decrease. This is reflected in our sensitivity analysis: White Non-Hispanic patients — who have the highest proportion of melanoma cases detected in the local stage — have the least total expected cost, while Black patients — who have the lowest proportion of local melanoma cases — have the highest expected cost. The ordering of the expected cost per demographic is the same as the reverse order of their proportion of local melanoma cases.

The trend can be visualized if we plot the proportion of local melanoma cases vs. the total costs, as shown in Figure 9. The plotted points fall close to the line of best fit (in red), showing a clear near-linear inverse relationship between local stage diagnosis frequency and eventual cost. This is expected because local stage diagnosis corresponds with easier and less costly treatment, so a higher proportion of local stage diagnoses would naturally lead to lower treatment costs on net.

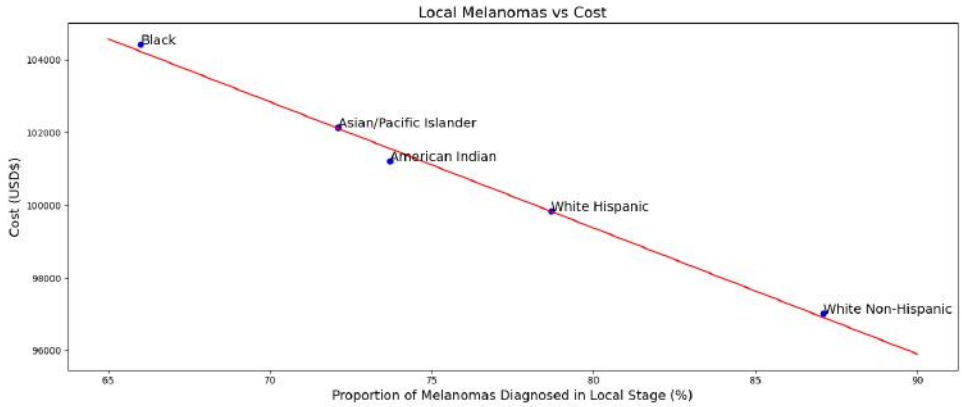

Figure 9: Proportion of local melanoma cases at diagnosis vs. total expected cost of melanoma treatment

By performing a linear regression, we estimate a decrease in expected cost of $1271.22 for every 1% increase in the proportion of melanoma cases detected in the local stage.

4.4.4 Screening Diagnosis Sensitivity Analysis

We calculate the sensitivity of our model to alternative screening diagnosis vectors. As described in Section 4.2.7, we calculate the diagnosis vector with screening $S(\boldsymbol{D}_0)$ given diagnosis vector $\boldsymbol{D}_0$ by scaling the local stage diagnosis frequency up by 23%, and scaling the regional and distant stage frequencies by an equal proportion so the sum of the frequencies is still 1. We consider perturbing the 23% scale, to scalings between 18% and 28%. Figure 10 provides a graph of the savings due to screening plotted against the screen scaling for stage diagnosis frequency vector.

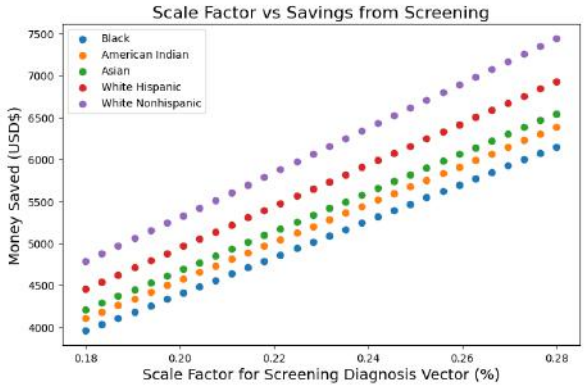

Figure 10: The scale factor of the diagnosis vector vs. the amount saved by screening

As Figure 10 demonstrates, the money saved is linear with the scale factor of screening. This is due to our formula for total cost, which multiplies the D vector by a matrix: perturbing the scale factor linearly, in the manner we do, causes the total cost to be perturbed linearly as well. A linear positive relationship matches with real-life expectations: when screening is less effective (the scale factor is smaller), we expect less money to be saved by screening. By performing a linear regression, we find that we expect an increase in savings of $219.69 for every 1% increase in scale factor.

5 Risk Analysis

5.1 Melanoma Risk

From Assumption 4 of our UV exposure model, we are able to infer the distribution of melanoma risk from the sampling distribution $\boldsymbol{A}$ of UV exposure risk predicted in (24).

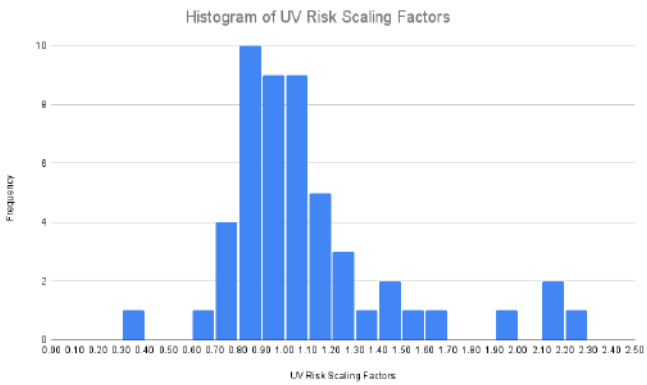

Figure 11: Distribution of 2019–2024 melanoma risk scaling factors for 51 sampled counties

As seen in Figure 11, the distribution of the UV risk scaling factors is skewed right, but in general the values are concentrated between scaling factors of 0.6 to 1.3. The mean of 1.10 suggests that a typical county in the sample will experience a 10% increase in UV exposure risk between 2019 and 2024.

Melanoma incidence risk M as a function of UV exposure risk U is given by the following power law adapted from Moan et al. 2015 [27]:

$$\ln(M) = A_b \cdot \ln(U) \tag{25}$$

where A_b is the biological amplification factor that varies by sex and melanoma site. Given scaling factors α and β for U and M respectively:

$$\ln(\beta M) = A_b \cdot \ln(\alpha U) \quad \Rightarrow \quad \beta M = (\alpha U)^{A_b} \tag{26}$$

we can express melanoma risk scaling factor β in terms of UV risk scaling factor α:

$$\beta = \alpha^{A_b} \tag{27}$$

By definition, the 2019–2024 UV risk scaling factors for each sampled county are given by $\boldsymbol{A}$. From these values, we can compute the distribution of melanoma risk scaling factors for our sampled counties using (27) for different values of A_b found in Moan et al. [27].

	(M, H&N, 1.29)	(F, H&N, 1.52)	(M, Tr, 1.81)	(F, Tr, 1.93)
Mean	1.1575	1.2103	1.2912	1.3301
Median	1.0172	1.0203	1.0242	1.0258
St. Dev.	0.5243	0.6664	0.8813	0.9842

Table 5: Table of 2019–2024 melanoma risk scaling factor statistics predicted from $\boldsymbol{A}$ for different demographics given by (sex, melanoma site, A_b) where: M=male, F=female, H&N=head and neck, Tr=trunk

Table 5 predicts that for a typical county in the sample, all potential demographic outcomes will experience elevated melanoma risk in 2024 when compared to 2019. However, the standard deviations for each demographic suggest a broad spread, likely due to both the inherent geographic diversity of the US and the varying climate trends for different localities predicted in the UV exposure model. For each melanoma site, females will suffer from greater increase in melanoma risk than men.

Finally, the predictions and computations for $\boldsymbol{A}$ can be generalized to sample sizes far greater than the 51 counties examined here, yielding melanoma risk projections that cover a broader population over a longer time frame into the future. Greater separation of potential outcomes can also be achieved by experimentally determining A_b values specific to smaller demographics.

5.2 Financial Risk

Because our proposed models ultimately quantify the impact of melanoma through financial cost, our risk analysis will be most beneficial if performed from the perspective of the healthcare insurance industry — which in most cases, bears the majority of melanoma-induced financial costs [8]. Under the insurance perspective, one main risk of melanoma is having to pay for the treatment of many late-stage melanoma cases, as these cases incur a tremendous cost with poor prognosis.

Another main risk for healthcare insurance companies is net-loss services, such as screening. If covering screening and other services costs more money than it saves in the long term, such a practice would depress the profits of an insurance company. When considering a melanoma screening policy, an insurance company has two options:

1. Implement melanoma screening and incur the costs associated with paying \$150 per person per year for melanoma screening and the costs associated with researching which populations to optimally target with screening. In this case, melanoma would be caught in much earlier stages than it would be without melanoma screening, and therefore the insurance company would incur lower hospitalization costs because it is more expensive to treat later stages of melanoma [8].

2. Do not implement melanoma screening and incur the additional costs associated with higher rates of hospitalization for and death from late-stage melanoma. In this case, the insurance company would not have to pay \$150 per person per year for melanoma screening, but it would also mean that cases of melanoma are more likely to be diagnosed at later stages or even remain undiagnosed, leading to expensive hospitalization costs for the insurance company.

In our model, the expected value of risk is the expected cost for treating a melanoma patient. As seen in Table 2, the expected cost for treating a melanoma patient differs based on the patient's demographics. Of the two options, we infer that an insurance company would reasonably choose to not implement melanoma screening because it would not be economically

feasible for them to take on the risk unless the rate of new melanoma cases rose to 2-3% of the total population.

6 Recommendations

Though we determined in the course of this analysis that melanoma screenings are not yet a cost-effective risk mitigation strategy, if cases were to exceed r^*, screening could become effective at mitigating late-stage melanoma. More generally, cancer screening serves as an effective risk mitigation strategy for early diagnosis and prevention for many types of cancer, such as breast and colon cancer. Further analysis using the method developed in our paper, could be cross-applied to screening for other forms of cancer, or other multi-stage diseases in general.

There is a consensus in literature that melanoma risk is determined by genetic factors and exposure to UV radiation. Public awareness about the dangers of UV radiation, whether solar or through tanning beds [42], could in and of itself mitigate melanoma risks. As such, healthcare companies should make it a priority to inform stakeholders about easy ways to prevent melanoma. Another simple way to reduce late-stage melanoma risk is to inform the public about melanoma self-diagnosis [43]. Self-diagnosis for skin lesions can be difficult, as many tumors can appear benign and vice versa, so many early stage melanoma cases go unreported until later stages, which are evidently more expensive to treat. Informing people about how to determine if skin lesions are malignant, and to go see a doctor if that is the case, would increase the number of early-stage melanoma diagnoses which are much easier to treat and have dramatically improved prognostic outcomes [44].

For cancers like melanoma, where screening is generally not cost-benefit effective, there are two foreseeable advances that could alter the risk mitigation landscape. First, advances in technology could dramatically reduce the cost of screening. Specifically, advances in computer vision could allow patients and doctors to diagnose skin lesions using images alone [45]. Second, the use of genetic screening in order to better quantify predisposition to melanoma could help tailor screening to high-risk individuals [46] [47]. With rapid advances in genome sequencing, the healthcare industry could use genetic analysis to preempt melanoma and other diseases which have high inherent genetic risk.

Finally, it is crucial to note that our model is limited to only the direct financial costs of melanoma screening. Our model does not take into account more complex — and less easily quantifiable — factors, such as the patients' increased welfare and sense of security if the insurance company is willing to incur a loss on paying for screening, or the potential net profit from additional customers the insurance company could attract through a screening program. Thus, we recommend any insurance companies using our model to additionally consider projecting cost-benefit analyses for these factors into the future.

Bibliography

[1] Robert M. Kaplan and Yair M. Babad. "Balancing influence between actors in healthcare decision making". In: *BMC health services research* 11.1 (2011), pp. 1–14.

[2] American Academy of Dermatology. *Skin cancer*. June 2021. URL: https://www.aad.org/media/stats-skin-cancer.

[3] Howard W. Rogers, Martin A. Weinstock, and Steven R. Feldman. "Incidence Estimate of Nonmelanoma Skin Cancer (Keratinocyte Carcinomas) in the US Population". In: *JAMA Dermatology* 151.10 (Oct. 2015), p. 1081.

[4] American Cancer Society. *Melanoma skin cancer statistics*. URL: https://www.cancer.org/cancer/melanoma-skin-cancer/about/key-statistics.html.

[5] Centers for Disease Control and Prevention. *Melanoma incidence and mortality, United States – 2012–2016*. June 2019. URL: https://www.cdc.gov/cancer/uscs/about/data-briefs/no9-melanoma-incidence-mortality-UnitedStates-2012-2016.htm.

[6] Gery P. Guy et al. "Vital signs: melanoma incidence and mortality trends and projections—United States, 1982–2030". In: *MMWR. Morbidity and mortality weekly report* 64.21 (2015), p. 591.

[7] Insider Intelligence. *US healthcare industry in 2022: Analysis of the health sector, healthcare trends, & future of digital health*. URL: https://www.insiderintelligence.com/insights/healthcare-industry/.

[8] Gery P. Guy et al. "Melanoma treatment costs: a systematic review of the literature, 1990–2011". In: *American journal of preventive medicine* 43.5 (2012), pp. 537–545.

[9] KB Keisler-Starkey and Lisa N Bunch. "Health Insurance Coverage in the United States: 2020". In: *United States Census Bureau* (2021).

[10] D. Hashim et al. "The global decrease in cancer mortality: trends and disparities". In: *Annals of Oncology* 27.5 (2016), pp. 926–933.

[11] Darrell S. Rigel. "Epidemiology of Melanoma". In: *Seminars in Cutaneous Medicine and Surgery* 29.4 (Dec. 2010), pp. 204–209.

[12] Chamelli Jhappan, Frances P. Noonan, and Glenn Merlino. "Ultraviolet radiation and cutaneous malignant melanoma". In: *Oncogene* 22.20 (2003), pp. 3099–3112.

[13] D. M. Parkin, D. Mesher, and P. Sasieni. "Cancers attributable to solar (ultraviolet) radiation exposure in the UK in 2010". In: *British Journal of Cancer* 105.2 (2011), S66–S69.

[14] A. K. Bharath and R. J. Turner. "Impact of climate change on skin cancer". In: *Journal of the Royal Society of Medicine* 102.6 (2009), pp. 215–218.

[15] Eva Rawlings Parker. "The influence of climate change on skin cancer incidence – A review of the evidence". In: *International Journal of Women's Dermatology* 7.1 (2021), pp. 17–27.

[16] Esther Erdei and Salina M. Torres. "A new understanding in the epidemiology of melanoma". In: *Expert review of anticancer therapy* 10.11 (2010), pp. 1811–1823.

[17] Raghav Tripathi et al. "Racial differences in time to treatment for melanoma". In: *Journal of the American Academy of Dermatology* 83.3 (Apr. 2020), pp. 854–859.

[18] Melissa Ward-Peterson et al. "Association between race/ethnicity and survival of melanoma patients in the United States over 3 decades". In: *Medicine* 95.17 (Apr. 2016).

[19] Melody J. Eide, Martin A. Weinstock, and Melissa A. Clark. "Demographic and socioeconomic predictors of melanoma prognosis in the United States". In: *Journal of Health Care for the Poor and Underserved* 20.1 (2009), pp. 227–245.

[20] Centers for Disease Control and Prevention. *United States Cancer Statistics: Data Visualizations*. URL: `https://gis.cdc.gov/Cancer/USCS/#/AtAGlance/`.

[21] Google Earth Engine Data Catalog. *Sentinel-5P Datasets in Earth Engine*. URL: `https://developers.google.com/earth-engine/datasets/catalog/sentinel-5p`.

[22] I. Khan, D. Shah, and S. S. Shah. "Covid-19 pandemic and its positive impacts on environment: An updated review". In: *International Journal of Environmental Science and Technology* 18.2 (Nov. 2020), pp. 521–530.

[23] Xiao-Cheng Wu et al. "Racial and ethnic variations in incidence and survival of cutaneous melanoma in the United States, 1999-2006". In: *Journal of the American Academy of Dermatology* 65.5 (2011), S26–e1.

[24] R. L. Morton, K. Howard, and J. F. Thompson. "The cost-effectiveness of sentinel node biopsy in patients with intermediate thickness primary cutaneous melanoma". In: *Annals of surgical oncology* 16.4 (2009), pp. 929–940.

[25] Michael J. R. Martin, Stephen A. Matthews, and Barrett A. Lee. "The spatial diffusion of racial and ethnic diversity across US counties". In: *Spatial demography* 5.3 (2017), pp. 145–169.

[26] Joel Lieske. "Regional subcultures of the United States". In: *The Journal of Politics* 55.4 (1993), pp. 888–913.

[27] Johan Moan et al. "The relationship between UV exposure and incidence of skin cancer". In: *Photodermatology, photoimmunology & photomedicine* 31.1 (2015), pp. 26–35.

[28] Zhiguang Wang and Tim Oates. "Imaging time-series to improve classification and imputation". In: *Twenty-Fourth International Joint Conference on Artificial Intelligence*. 2015.

[29] Ling Zhang et al. "Spatio-temporal convolutional LSTMs for tumor growth prediction by learning 4D longitudinal patient data". In: *IEEE transactions on medical imaging* 39.4 (2019), pp. 1114–1126.

[30] Anvita Gupta et al. "Generative recurrent networks for de Novo Drug Design". In: *Molecular Informatics* 37.1-2 (2017), p. 1700111.

[31] Jong-Min Yeom et al. "Spatial mapping of short-term solar radiation prediction incorporating geostationary satellite images coupled with deep convolutional LSTM networks for South Korea". In: *Environmental Research Letters* 15.9 (2020), p. 094025.

[32] Hoonyong Lee et al. "Detecting excessive load-carrying tasks using a deep learning network with a Gramian Angular Field". In: *Automation in Construction* 120 (2020), p. 103390.

[33] Ying-Yi Hong, John Joel F. Martinez, and Arnel C. Fajardo. "Day-ahead solar irradiation forecasting utilizing Gramian Angular Field and Convolutional Long Short-Term Memory". In: *IEEE Access* 8 (2020), pp. 18741–18753.

[34] Matthew D. Zeiler. "Adadelta: an adaptive learning rate method". In: *arXiv preprint arXiv:1212.5701* (2012).

[35] Richard Shellenberger, Mohammed Nabhan, and Sweta Kakaraparthi. "Melanoma screening: A plan for improving early detection". In: *Annals of medicine* 48.3 (2016), pp. 142–148.

[36] Elena Losina et al. "Visual screening for malignant melanoma: a cost-effectiveness analysis". In: *Archives of Dermatology* 143.1 (2007), pp. 21–28.

[37] Alan N. Houghton and David Polsky. "Focus on melanoma". In: *Cancer cell* 2.4 (2002), pp. 275–278.

[38] Renato Cesar Sato and Désirée Moraes Zouain. "Markov models in Health Care". In: *Einstein (São Paulo)* 8.3 (Sept. 2010), pp. 376–379.

[39] Shasa Hu et al. "Comparison of stage at diagnosis of melanoma among Hispanic, black, and white patients in Miami-Dade County, Florida". In: *Archives of Dermatology* 142.6 (June 2006).

[40] Clara Curiel-Lewandrowski et al. "Survival is not the only valuable end point in melanoma screening". In: *Journal of Investigative Dermatology* 132.5 (2012), pp. 1332–1337.

[41] Martha Matsumoto et al. "Estimating the cost of skin cancer detection by dermatology providers in a large health care system". In: *Journal of the American Academy of Dermatology* 78.4 (2018).

[42] Richard P. Gallagher, John J. Spinelli, and Tim K. Lee. "Tanning beds, sunlamps, and risk of cutaneous malignant melanoma". In: *Cancer Epidemiology and Prevention Biomarkers* 14.3 (2005), pp. 562–566.

[43] Martin A Weinstock et al. "Melanoma early detection with thorough skin self-examination: the "Check It Out" randomized trial". In: *American journal of preventive medicine* 32.6 (2007), pp. 517–524.

[44] Michelle R. Iannacone and Adèle C. Green. "Towards skin cancer prevention and early detection: evolution of skin cancer awareness campaigns in Australia". In: *Melanoma management* 1.1 (2014), pp. 75–84.

[45] Megan Lewis. *An artificial intelligence tool that can help detect melanoma*. en. Apr. 2, 2021. URL: `https://news.mit.edu/2021/artificial-intelligence-tool-can-help-detect-melanoma-0402`.

[46] Lauren G Aoude et al. "Genetics of familial melanoma: 20 years after CDKN 2 A". In: *Pigment cell & melanoma research* 28.2 (2015), pp. 148–160.

[47] Michael F Berger and Levi A Garraway. "Applications of genomics in melanoma oncogene discovery". In: *Hematology/Oncology Clinics* 23.3 (2009), pp. 397–414.