

---

# Counterfactual Analysis in Dynamic Latent-State Models

---

Martin Haugh<sup>1</sup> Raghav Singal<sup>2</sup>

## Abstract

We provide an optimization-based framework to perform counterfactual analysis in a dynamic model with hidden states. Our framework is grounded in the “abduction, action, and prediction” approach to answer counterfactual queries and handles two key challenges where (1) the states are *hidden* and (2) the model is *dynamic*. Recognizing the lack of knowledge on the underlying causal mechanism and the possibility of infinitely many such mechanisms, we optimize over this space and compute upper and lower bounds on the counterfactual quantity of interest. Our work brings together ideas from causality, state-space models, simulation, and optimization, and we apply it on a breast cancer case study. To the best of our knowledge, we are the first to compute lower and upper bounds on a counterfactual query in a dynamic latent-state model.

## 1. Introduction

*Counterfactual analysis*, falling on the third rung of Pearl’s ladder of causation (Pearl & Mackenzie, 2018), is a fundamental problem in causality. It requires us to *imagine* a world where a certain policy was enacted with a corresponding outcome *given* that a different policy and outcome were actually observed. It is performed via the 3-step framework of *abduction* (conditioning on the observed data), *action* (changing the policy), and *prediction* (computing the counterfactual quantity of interest (CQI)), and has wide-ranging applications (Pearl, 2009a;b).

As a concrete application in healthcare and legal reasoning, consider someone who recently died from breast cancer. The exact progression of her disease is unknown. What is known, however, is that over a period of time prior to her diagnosis, her insurance company adopted a strategy of denying her regular scans (e.g., mammograms) even though

these scans should have been covered by her policy. Had these scans gone ahead, the cancer may have been found earlier and the patient’s life saved. Now a court wants to know the probability that her life would have been saved had the routine scans been permitted.

On top of the challenges posed by standard counterfactual analysis, there are two that are particular to such a setting. First, it’s possible the underlying state of the patient (e.g., stage of cancer) is *hidden / latent* and we only observe a noisy signal depending on the accuracy of the scan (e.g., sensitivity and specificity of a mammogram). Second, the underlying model is *dynamic* as the patient’s state evolves over time. As such, our goal in this work is to perform counterfactual analysis in dynamic latent-state models.<sup>1</sup>

Two streams of work are closely related to ours. The first relates to works on constructing bounds on CQIs (Balke & Pearl, 1994; Tian & Pearl, 2000; Kaufman et al., 2005; Cai et al., 2008; Pearl, 2009b; Mueller et al., 2021; Zhang et al., 2021). These papers focus on *static* models. We note that despite some similarities of our work with Zhang et al. (2021), the two approaches are quite different. In particular, while both papers recognize the relevance of polynomial optimization for bounding CQIs, Zhang et al. (2021) do not solve polynomial optimization problems but instead propose Monte-Carlo algorithms as a work-around. In contrast, we actually solve polynomial optimization problems via sample average approximations (SAAs), which we generate via Monte-Carlo. As such, Monte-Carlo serves as an “input” to our polynomial programs whereas Zhang et al. (2021) use it as a “substitute” for polynomial programs. As mentioned above, another difference is our focus on dynamic models whereas Zhang et al. (2021) focus on static models. The second stream is more recent and concerns counterfactual analysis in dynamic models (Buesing et al., 2019; Oberst & Sontag, 2019; Lorberbom et al., 2021; Tsirtsis et al., 2021). Except Buesing et al. (2019), none of these works allows for *latent* states. In addition, these works perform counterfactual analysis by embedding assumptions that are strong enough to restrict the underlying set of causal mechanisms to a singleton. In particular, Buesing et al. (2019) explicitly fix

---

<sup>1</sup>Imperial College <sup>2</sup>Dartmouth College. Correspondence to: MH <m.haugh@imperial.ac.uk>, RS <rsingal@dartmouth.edu>.

---

<sup>1</sup>The key feature distinguishing a static model from a dynamic model with  $T$  periods say, is that the single-period structure is repeated  $T$  times. As we shall see, our framework takes advantage of this repeated structure in several ways.

a single causal mechanism whereas Oberst & Sontag (2019) and Tsirtsis et al. (2021) invoke *counterfactual stability* and implicitly fix the causal mechanism (via the Gumbel-max distribution). Lorberbom et al. (2021) extend the Gumbel-max approach but their choice of causal mechanism is the one that minimizes variance when estimating the CQI. In summary, none of these approaches explicitly account for all possible causal mechanisms, and therefore, they do not consider the construction of lower and upper bounds on the CQI, which is our focus.

Our key contribution is to provide a principled framework for counterfactual analysis in dynamic latent-state models. We define our problem in §2 and discuss counterfactual stability in §3. In §4, we present our solution approach and we describe our numerics in §5. We conclude in §6.

## 2. Problem Definition

We first define the underlying dynamic latent-state model (§2.1) and then describe the counterfactual analysis problem (§2.2). We will use a breast cancer application as a vehicle for explaining ideas throughout but it should be clear our framework is quite general.

### 2.1. A Dynamic Latent-State Model

The model, visualized in Figure 1, has  $T$  discrete periods. In each period  $t$ , the system is in a hidden state  $H_t \in \mathbb{H}$  (finite). As a stochastic function of  $H_t$  and the policy  $X_t \in \mathbb{X}$  (finite), we observe an emission  $O_t \in \mathbb{O}$  (finite). The *emission probability* is denoted by  $e_{hxi} := \mathbb{P}(O_t = i \mid H_t = h, X_t = x)$  for all  $(h, x, i)$ . This is followed by the state  $H_t$  transitioning to  $H_{t+1}$  with *transition probability*  $q_{hih'} := \mathbb{P}(H_{t+1} = h' \mid H_t = h, O_t = i)$ . The model  $\mathbf{M}$  comprises three primitives:  $\mathbf{M} \equiv (\mathbf{p}, \mathbf{E}, \mathbf{Q})$ , where  $\mathbf{p} := [p_h]_h$  denotes the *initial state distribution* with  $p_h := \mathbb{P}(H_1 = h)$  for all  $h$ ,  $\mathbf{E} := [e_{hxi}]_{h,x,i}$ , and  $\mathbf{Q} := [q_{hih'}]_{h,i,h'}$ .




Figure 1. A dynamic latent-state model. States  $H_{1:T}$  are hidden (red). Emissions  $O_{1:T}$  are observed (blue).  $X_{1:T}$  represents the policy (observed). (We use the notation  $X_{1:T} := (X_t)_{t=1}^T$ .)

In the breast cancer application, the time periods map to the frequency of mammograms (e.g., 6 months) and the hidden state  $H_t \in \{1, \dots, 7\}$  denotes the patient’s condition. State 1 equates to the patient being healthy whereas states 2 and 3 correspond to *undiagnosed* in-situ and invasive breast cancer, respectively. States 4 and 5 correspond to *diagnosed*

in-situ and invasive breast cancer respectively, with the understanding that the cancer treatment has begun (since it has been diagnosed). States 6 and 7 are absorbing and denote recovery from cancer (due to treatment) and death from cancer, respectively. The observation  $O_t \in \{1, \dots, 7\}$  captures the mammogram result. A value of 1 means no screening took place, whereas 2 denotes a negative screening result (possibly a false negative). A value of 3 corresponds to a positive mammogram result, but followed by a negative biopsy (i.e., the patient is healthy and the mammogram produced a false positive). Observations 4 and 5 map to correctly diagnosed in-situ and invasive cancer respectively, i.e., a positive mammogram followed by a positive biopsy. Observations 6 and 7 are used to denote patient recovery and death from breast cancer, respectively. The variable  $X_t \in \{0, 1\}$  models the insurance company’s coverage policy for the mammograms, with 0 denoting the company covers it and 1 denoting the company (incorrectly) denies the coverage. If the coverage is denied, then the mammogram is not performed and hence, the observation cannot be 2, 3, 4, or 5. (In this application, the  $X_t$ ’s are deterministic but in general, they could be the result of a randomized policy.)

Our model is therefore a generalization of a hidden Markov model (HMM) since  $H_{t+1}$  depends not only on  $H_t$  but also on  $O_t$ . The dependence on  $O_t$  is needed to capture the fact that if cancer was detected during period  $t$  and treatment began at that point of time, i.e.,  $O_t \in \{4, 5\}$ , then  $H_{t+1}$  depends on the fact that the treatment began in period  $t$ . For example, if  $H_t = 2$  (in-situ cancer) and  $O_t = 4$  (in-situ diagnosed and hence, treatment began), then  $H_{t+1}$  would be different compared to when  $H_t = 2$  and  $O_t = 2$  (false negative and hence, treatment did not begin).

**Remark 1.** Ayer et al. (2012) employed a similar model for determining an optimal screening strategy for breast cancer but as their goal was to optimize over screening strategies, their model was a partially observable Markov decision process (POMDP). In contrast, our goal is not to find an optimal strategy but to evaluate CQIs. As such, our model is not a POMDP although it is easily related to a POMDP setting. For example, we can view the insurance company’s observed coverage strategy and the counterfactual strategy where coverage is always provided, as being feasible strategies from a POMDP. Finally, we also note that a practical justification for our model comes from the simulation model used by the National Cancer Institute (UWBCS, 2013).

### 2.2. The Counterfactual Analysis Problem

We now use our dynamic model to state the counterfactual analysis problem.

**Observed data.** Suppose we observe emissions  $o_{1:T}$  with the underlying policy being  $x_{1:T}$ . The true hidden states  $h_{1:T}$  are not observed. In the context of breast cancer, the

observations for a particular patient might be as follows:

$$\underbrace{o_1, \dots, o_{\tau_s-1}}_{\in \{2,3\}}, \underbrace{o_{\tau_s}, \dots, o_{\tau_e}}_{=1}, \underbrace{o_{\tau_e+1}, \dots, o_{\tau_d-1}}_{\in \{4,5\}}, \underbrace{o_{\tau_d:T}}_{=7}. \quad (1)$$

That is, the patient was screened ( $x_{1:\tau_s-1} = 0$ ) and appeared healthy ( $o_{1:\tau_s-1} \in \{2, 3\}$ ) up to and including time  $\tau_s - 1$ . Coverage was denied during periods  $\tau_s$  to  $\tau_e$ , i.e.  $x_{\tau_s:\tau_e} = 1$ ; (see red font in (1)). Hence, screening was not performed during those periods ( $o_{\tau_s:\tau_e} = 1$ ). As soon as the coverage for screening was re-approved (period  $\tau_e + 1$  and hence,  $x_{\tau_e+1} = 0$ ), the patient was found to have cancer (either in-situ or invasive) and the corresponding treatment began; thus,  $o_{\tau_e+1} \in \{4, 5\}$ . Unfortunately, the patient died at  $\tau_d$ .

**CQI.** We focus on the well-known *probability of necessity* (PN) (Pearl, 2009a) as our CQI. It is the probability the patient would have not died (counterfactual state  $\tilde{H}_T \neq 7$ ) had the screening been covered in every period (interventional policy  $\tilde{x}_{1:T} = 0$ ) given the observed data ( $o_{1:T}, x_{1:T}$ ). (“Tilde” notation denotes quantities in the counterfactual world.) The interpretation of  $\tilde{x}_{1:T}$  is straightforward as it is fixed exogenously. The counterfactual state  $\tilde{H}_T$  is obtained via the 3 steps of abduction, action, and prediction (Pearl, 2009b). Step 1 (abduction) involves conditioning on the observed data ( $o_{1:T}, x_{1:T}$ ) to form a posterior belief over the hidden states. Step 2 (action) changes the policy from  $x_{1:T}$  to  $\tilde{x}_{1:T}$  and brings us to the counterfactual world  $\tilde{M}$ . Step 3 (prediction) computes PN in the counterfactual model:

$$\text{PN} = \mathbb{P}(\tilde{H}_T \neq 7), \quad (2)$$

with the understanding that the event  $\{\tilde{H}_T \neq 7\}$  is conditional on  $(o_{1:T}, x_{1:T})$ . Though we focus on PN, it is easy to extend our framework to a broad class of CQIs as the abduction and action steps do not depend on the CQI.

Given our focus on counterfactual analysis, we will assume the primitives  $(\mathbf{p}, \mathbf{E}, \mathbf{Q})$  are known. We discuss their calibration to real-world data in §5 and emphasize that even with known  $(\mathbf{p}, \mathbf{E}, \mathbf{Q})$ , counterfactual analysis is challenging. This is because we are interested in counterfactuals at an *individual* level (i.e., conditioning on the patient-level data  $(o_{1:T}, x_{1:T})$  via abduction), as opposed to the *population* level. A population-level counterfactual analysis would ignore the first step of abduction but simply change the policy to  $\tilde{x}_{1:T}$  to predict the CQI (by simulating the resulting model and obtaining a Monte-Carlo estimate of PN or doing it in closed-form if analytically tractable). However, this is very different from the task at hand, which falls on the highest rung of Pearl’s ladder of causation (Pearl & Mackenzie, 2018). For instance, consider a patient who dies immediately after the coverage was denied versus a patient who dies a couple of years after the coverage was denied. Clearly, the first patient had a more “aggressive”

cancer and hence we expect that her PN would be lower. By conditioning on individual-level data  $(o_{1:T}, x_{1:T})$ , we are able to account for such differences. However, it makes the problem considerably more challenging.

In our dynamic latent-state model, each of the three steps of abduction, action, and prediction presents its own set of challenges<sup>2</sup>, which we discuss when presenting our methodology in §4. Before doing so, we discuss the notion of counterfactual stability (CS), which has become a popular approach in some settings (Oberst & Sontag, 2019).

### 3. Limitations of Counterfactual Stability

Instead of discussing CS in our dynamic latent-state model, we do so using the following simple model:  $X \rightarrow Y$ . Suppose we observe an outcome  $Y = y$  under policy  $X = x$ . With  $Y_x := Y \mid (X = x)$ , CS requires that the counterfactual outcome under an interventional policy  $\tilde{x}$  (denoted by  $\tilde{Y} := Y_{\tilde{x}} \mid Y_x = y$ ) cannot be  $y'$  (for  $y' \neq y$ ) if  $\mathbb{P}(Y_{\tilde{x}} = y) / \mathbb{P}(Y_x = y) \geq \mathbb{P}(Y_{\tilde{x}} = y') / \mathbb{P}(Y_x = y')$ . In words, CS states that if  $y$  was observed and this outcome becomes relatively more likely than  $y'$  under the intervention, then the counterfactual outcome  $\tilde{Y}$  can not be  $y'$ .

Though somewhat appealing, the appropriateness of CS depends on the application and should be justified by domain specific knowledge. Moreover, we show in Example 1 that CS can permit counterfactuals that it was seemingly designed to exclude.

**Example 1.** Consider the  $X \rightarrow Y$  model and suppose  $X \in \{0, 1\}$  denotes a medical treatment and  $Y \in \{\text{bad}, \text{better}, \text{best}\}$  the patient outcome. For illustration, suppose the outcome  $Y_x$  obeys the following distribution:  $Y_0 \sim \{\text{bad}, \text{better}, \text{best}\}$  w.p.  $\{0.2, 0.3, 0.5\}$  and  $Y_1 \sim \{\text{bad}, \text{better}, \text{best}\}$  w.p.  $\{0.2, 0.2, 0.6\}$ . That is, under treatment ( $x = 1$ ), the “best” outcome becomes more likely but the likelihood of the “bad” outcome does not change. Consider a patient whose outcome  $Y$  was “better” under no treatment ( $x = 0$ ). Suppose also that domain specific knowledge tell us that even at the individual level, the counterfactual outcome  $\tilde{Y}$  should not be worse under treatment ( $\tilde{x} = 1$ ) than under no treatment ( $x = 0$ ). However, since

$$\frac{\mathbb{P}(Y_1 = \text{better})}{\mathbb{P}(Y_0 = \text{better})} = \frac{0.2}{0.3} < \frac{0.2}{0.2} = \frac{\mathbb{P}(Y_1 = \text{bad})}{\mathbb{P}(Y_0 = \text{bad})},$$

“bad” is a feasible counterfactual outcome under CS.

Even if CS is appropriate, its current operationalization has a key limitation. In particular, instead of considering all possible structural causal models (SCMs) that obey CS, both Oberst & Sontag (2019) and Tsiirtsis et al. (2021) pick

<sup>2</sup>Instead of using the “twin networks” approach (Pearl, 2009b), we perform the counterfactual analysis directly by leveraging the structure in our model.

one SCM via the Gumbel-max distribution. Ideally, one should characterize the space of all SCMs obeying CS, and map that space into appropriate bounds on the CQI.

We present our optimization-based framework to perform counterfactual analysis next. Our framework does not rely on CS. However, *if* CS is deemed appropriate for one or more components of the SCM (see §4), our approach allows us to encode CS via linear constraints in the optimization and characterize the *entire* space of solutions that obey CS. We do this in §5 to negatively answer the open question of Oberst & Sontag (2019) regarding whether Gumbel-max obeys CS uniquely. Further, if enforcing the so-called *path-wise monotonicity* (PM) is desirable, i.e., ensuring the counterfactual outcome does not worsen under a better intervention (as we assumed in Example 1), then we can embed it in our optimization via linear constraints as well.

## 4. Counterfactual Analysis via Optimization

We now present our solution methodology for the counterfactual analysis problem introduced in §2. We first discuss the underlying SCM (§4.1), which is a precursor to defining the counterfactual model  $\tilde{M}$  (§4.2), which feeds into our optimization framework for counterfactual analysis (§4.3).

### 4.1. The Structural Causal Model (SCM)




Figure 2. The SCM underlying the dynamic latent-state model. The only difference between the SCM here and Figure 1 is the addition of (grey) exogenous noise nodes  $[\mathbf{U}_t, \mathbf{V}_t]_t$ .

To understand the SCM (Figure 2), consider  $O_t$  for any  $t$ , which is a stochastic function of  $(H_t, X_t)$ . The stochasticity is driven by the exogenous noise vector  $\mathbf{V}_t := [V_{thx}]_{h,x}$ , which comprises of  $|\mathbb{H}||\mathbb{X}|$  noise variables. We model the exogenous node as a vector (as opposed to a scalar) to capture the fact that each  $O_{thx} := O_t | (H_t = h, X_t = x)$  defines a *distinct* random variable for all  $(h, x)$ . Moreover, these random variables might be independent or not. One way to handle this is to associate each  $O_{thx}$  with a noise variable  $V_{thx}$ . The *dependence structure* among these noise variables  $[V_{thx}]_{h,x}$  is then what determines the dependence structure among  $[O_{thx}]_{h,x}$ . The structural equation obeys

$$O_t = f(H_t, X_t, \mathbf{V}_t) = \sum_{h,x} f_{hx}(V_{thx}) \mathbb{I}_{\{H_t=h, X_t=x\}} \quad (3a)$$

where  $f_{hx}(\cdot)$  is defined using the emission distribution  $[e_{hxi}]_i$  and  $V_{thx} \sim \text{Unif}[0, 1]$  wlog. Similarly, for  $t > 1$ , recognizing that each  $H_{thi} := H_t | (H_{t-1} = h, O_{t-1} = i)$  is a distinct random variable for all  $(h, i)$ , we associate each  $H_{thi}$  with its own noise variable  $U_{thi}$ :

$$H_t = g(H_{t-1}, O_{t-1}, \mathbf{U}_t) = \sum_{h,i} g_{hi}(U_{thi}) \mathbb{I}_{\{H_{t-1}=h, O_{t-1}=i\}} \quad (3b)$$

where  $g_{hi}(\cdot)$  is defined using the transition distribution  $[q_{h'ih'}]_{h'}$  and  $U_{thi} \sim \text{Unif}[0, 1]$  wlog.

The representation in (3a) allows us to model  $[O_{thx}]_{h,x}$  and capture any dependence structure among these random variables by specifying the joint multivariate distribution of  $\mathbf{V}_t$ . Since the univariate marginals of  $\mathbf{V}_t$  are known ( $\text{Unif}[0, 1]$ ), specifying the multivariate distribution amounts to specifying the dependence structure or *copula*. (Of course, the same comment applies to (3b) and  $\mathbf{U}_t$  as well.) For example, if the  $V_{thx}$ 's are mutually independent (the independence copula) and we have  $(H_t = h', X_t = x')$ , then inferring the conditional distribution of  $V_{th'x'}$  will tell us nothing about the  $V_{thx}$ 's for  $(h, x) \neq (h', x')$ . Alternatively, if  $V_{thx} = V_{th'x'}$  for all pairs  $(h, x)$  and  $(h', x')$ , then this models perfect positive dependency (the comonotonic copula) and inferring the conditional distribution of  $V_{th'x'}$  amounts to simultaneously inferring the conditional distribution of all the  $V_{thx}$ 's. We emphasize that we must work with the exogenous vectors  $(\mathbf{U}_t, \mathbf{V}_t)$  when doing a *counterfactual* analysis since different joint distributions of  $(\mathbf{U}_t, \mathbf{V}_t)$  will lead to (possibly very) different values of PN. If we are not doing a counterfactual analysis and only care about the joint distribution of a (subset of)  $(O_{1:T}, H_{1:T})$  then our analysis will only depend on the joint distribution of the  $(\mathbf{U}_t, \mathbf{V}_t)$ 's via their known univariate marginals. We note the  $\mathbf{U}_t$ 's and  $\mathbf{V}_t$ 's must be mutually independent in order for the SCM to be consistent with the dependence / independence relationships implied by the model of Figure 1.

In our model, the emissions and the state transitions are time-independent. Thus, it is natural to assume the copulas underlying  $\mathbf{V}_t$  and  $\mathbf{U}_t$  are time-independent (*time invariance*). As such, we define the notation  $O_{hx} := O_t | (H_t = h, X_t = x)$  and  $H_{hi} := H_{t+1} | (H_t = h, O_t = i)$ .<sup>3</sup> Then,  $e_{hxi} = \mathbb{P}(O_{hx} = i)$  and  $q_{h'ih'} = \mathbb{P}(H_{hi} = h')$ .

While the copula view is useful from a conceptual point of view (since specifying copulas for  $\mathbf{U}_t$  and  $\mathbf{V}_t$  amounts to specifying an SCM), it is more convenient to work with an alternative construction of the SCM. This is because in discrete-state space models, there will be infinitely many joint distributions of  $\mathbf{V}$  (and  $\mathbf{U}$ ) that all lead to the same joint distribution of  $[O_{hx}]_{h,x}$  (and  $[H_{hi}]_{h,i}$ ). In other words, the joint distribution of  $[V_{hx}]_{h,x}$  does not uniquely identify

<sup>3</sup> $O_t | (H_t = h, X_t = x)$  is time-independent and hence, we use the notation  $O_{hx}$  instead of  $O_{thx}$ . Same logic holds for  $H_{hi}$ .

the joint distribution of  $[O_{hx}]_{h,x}$ . This is a consequence of Sklar’s Theorem from the theory of copulas and is discussed<sup>4</sup> further in §C. We will therefore take a more direct approach by modeling the unknown joint distribution of  $[O_{hx}]_{h,x}$  (and  $[H_{hi}]_{h,i}$ ). As such, we define

$$\theta_{\tilde{h}\tilde{x},hx}(\tilde{i}, i) := \mathbb{P}(O_{\tilde{h}\tilde{x}} = \tilde{i}, O_{hx} = i) \quad (4a)$$

$$\pi_{\tilde{h}\tilde{i},hi}(\tilde{h}', h') := \mathbb{P}(H_{\tilde{h}\tilde{i}} = \tilde{h}', H_{hi} = h') \quad (4b)$$

and observe that

$$\theta_{hx,hx}(i, i) = e_{hxi} \quad \forall (h, x, i) \quad (5a)$$

$$\pi_{hi,hi}(h', h') = q_{hih'} \quad \forall (h, i, h'). \quad (5b)$$

This holds because  $\pi_{hi,hi}(h', h') = \mathbb{P}(H_{hi} = h', H_{hi} = h') = \mathbb{P}(H_{hi} = h') = q_{hih'}$ . We also have *symmetry*, i.e.,

$$\theta_{\tilde{h}\tilde{x},hx}(\tilde{i}, i) = \theta_{hx,\tilde{h}\tilde{x}}(i, \tilde{i}) \quad \forall (\tilde{h}, \tilde{x}, \tilde{i}) \quad \forall (h, x, i) \quad (6a)$$

$$\pi_{\tilde{h}\tilde{i},hi}(\tilde{h}', h') = \pi_{hi,\tilde{h}\tilde{i}}(h', \tilde{h}') \quad \forall (\tilde{h}, \tilde{i}, \tilde{h}') \quad \forall (h, i, h'). \quad (6b)$$

This is because  $\pi_{\tilde{h}\tilde{i},hi}(\tilde{h}', h') = \mathbb{P}(H_{\tilde{h}\tilde{i}} = \tilde{h}', H_{hi} = h') = \mathbb{P}(H_{hi} = h', H_{\tilde{h}\tilde{i}} = \tilde{h}') = \pi_{hi,\tilde{h}\tilde{i}}(h', \tilde{h}')$ . We only defined the “pairwise marginals” in (4) but define the joint PMFs in (9). We are now ready to discuss the counterfactual model.

## 4.2. The Counterfactual Model $\tilde{\mathbf{M}}$

Recall from §2 that  $\tilde{\mathbf{M}}$  is obtained after the two steps of abduction (conditioning on the observed data  $(o_{1:T}, x_{1:T})$ ) and action (changing the policy from  $x_{1:T}$  to  $\tilde{x}_{1:T}$ ). Understanding the dynamics underlying  $\tilde{\mathbf{M}}$  are non-trivial, primarily due to the abduction step where the goal is to obtain the posterior distribution of the hidden path  $H_{1:T}$ . It is not possible to provide a closed-form expression for this distribution but we can use filtering / smoothing methods to describe the posterior dynamics of  $H_{1:T}$ . (See §B for details.)

We can therefore use these dynamics to generate  $B$  Monte-Carlo samples  $[h_{1:T}(b)]_{b=1}^B$  from the posterior, i.e., from the distribution of  $H_{1:T} \mid (o_{1:T}, x_{1:T})$ . Then, by conditioning on each sample  $b$ , it is possible to characterize  $\tilde{\mathbf{M}}$ . In particular, denote by  $\tilde{\mathbf{M}}(b) \equiv (\tilde{\mathbf{p}}(b), [\tilde{\mathbf{E}}^{(t)}(b)]_t, [\tilde{\mathbf{Q}}^{(t)}(b)]_t)$  the counterfactual model corresponding to posterior sample  $h_{1:T}(b)$ . Similar to the primitives  $(\mathbf{p}, \mathbf{E}, \mathbf{Q})$  in §2, the counterfactual primitives  $(\tilde{\mathbf{p}}(b), [\tilde{\mathbf{E}}^{(t)}(b)]_t, [\tilde{\mathbf{Q}}^{(t)}(b)]_t)$  correspond to initial state, emission, and transition distributions. As  $H_1$  in Figure 1 has no parents,  $\tilde{\mathbf{p}}(b)$  is such that the counterfactual hidden state in period 1 equals the posterior sample  $h_1(b)$ , i.e.,  $\tilde{h}_1(b) = h_1(b)$ . In contrast with  $\mathbf{E}$  and  $\mathbf{Q}$ , both  $\tilde{\mathbf{E}}^{(t)}(b)$  and  $\tilde{\mathbf{Q}}^{(t)}(b)$  are time-dependent (note

<sup>4</sup>In §C, we also discuss specific copulas (e.g., independence and comonotonic copulas) that can be used to provide benchmark values of PN.

the super-script “ $(t)$ ”). This is because the period  $t$  counterfactual emission  $\tilde{\mathbf{E}}^{(t)}(b) := [\tilde{e}_{\tilde{h}\tilde{x}\tilde{i}}^{(t)}(b)]_{\tilde{h},\tilde{x},\tilde{i}}$  and transition  $\tilde{\mathbf{Q}}^{(t)}(b) := [\tilde{q}_{\tilde{h}\tilde{i}h'}^{(t)}(b)]_{\tilde{h},\tilde{i},h'}$  probabilities are as follows:

$$\tilde{e}_{\tilde{h}\tilde{x}\tilde{i}}^{(t)}(b) = \mathbb{P}(O_{\tilde{h}\tilde{x}} = \tilde{i} \mid O_{h_t(b)x_t} = o_t) \quad (7a)$$

$$\tilde{q}_{\tilde{h}\tilde{i}h'}^{(t)}(b) = \mathbb{P}(H_{\tilde{h}\tilde{i}} = \tilde{h}' \mid H_{h_t(b)o_t} = h_{t+1}(b)). \quad (7b)$$

(The  $O_{hx}$  and  $H_{hi}$  notation is defined above (4).) The dependence on  $t$  is through the observed data  $(o_t, x_t)$  and the posterior samples  $(h_t(b), h_{t+1}(b))$ . As such, for each posterior path  $b$ ,  $\tilde{\mathbf{M}}(b)$  is a time-dependent dynamic latent-state model. If we knew  $\tilde{\mathbf{E}}^{(t)}(b)$  and  $\tilde{\mathbf{Q}}^{(t)}(b)$ , then we could simulate  $\tilde{\mathbf{M}}(b)$  to obtain a Monte-Carlo estimate of our CQI by averaging the CQI over the  $B$  posterior sample paths. However,  $\tilde{\mathbf{E}}^{(t)}(b)$  and  $\tilde{\mathbf{Q}}^{(t)}(b)$  are unknown as they depend on the joint distributions of  $\mathbf{U}_t$  and  $\mathbf{V}_t$ .

Towards this end, we can combine (7) with (4) to obtain

$$\tilde{e}_{\tilde{h}\tilde{x}\tilde{i}}^{(t)}(b) = \frac{\theta_{\tilde{h}\tilde{x},h_t(b)x_t}(\tilde{i}, o_t)}{e_{h_t(b)x_t o_t}}$$

$$\tilde{q}_{\tilde{h}\tilde{i}h'}^{(t)}(b) = \frac{\pi_{\tilde{h}\tilde{i},h_t(b)o_t}(\tilde{h}', h_{t+1}(b))}{q_{h_t(b)o_t h_{t+1}(b)}},$$

which express the unknown and time-dependent emission and transition distributions in terms of the unknown  $\theta$  and  $\pi$  that are time-independent.

## 4.3. Polynomial Optimization

We now propose an optimization model where we treat the unknowns  $(\theta, \pi)$  as decisions and maximize (minimize) the CQI to obtain an upper bound (lower bound). We present our optimization model in terms of the objective and constraints, followed by a discussion on how we can enforce CS and PM (if indeed they were deemed appropriate).

**Objective.** As in (2), we wish to understand the PN, which equals  $\mathbb{P}(\tilde{H}_T \neq 7)$ , where  $\tilde{H}_T$  is the hidden state at time  $T$  under  $\tilde{\mathbf{M}}$ . The randomness in  $\tilde{H}_T$  depends on the randomness in (i) the true hidden path  $H_{1:T}$  (captured by  $[h_{1:T}(b)]_b$ ) and (ii) the counterfactual model  $\tilde{\mathbf{M}} \mid H_{1:T}$  after conditioning on  $H_{1:T}$  (captured by  $\tilde{\mathbf{M}}(b)$ ). Lemma 1 decomposes PN using these two uncertainties. (All proofs are in §A.)

**Lemma 1.** *We have*

$$PN = 1 - \lim_{B \rightarrow \infty} \frac{1}{B} \sum_{b=1}^B \mathbb{P}_{\tilde{\mathbf{M}}(b)}(\tilde{H}_T = 7).$$

We next express  $\mathbb{P}_{\tilde{\mathbf{M}}(b)}(\tilde{H}_T)$  in terms of  $(\theta, \pi)$  from (4).

**Lemma 2.** For  $t \in \{T, T-1, \dots, 2\}$ ,  $\mathbb{P}_{\tilde{\mathbf{M}}(b)}(\tilde{h}_t) := \mathbb{P}_{\tilde{\mathbf{M}}(b)}(\tilde{H}_t = \tilde{h}_t)$  obeys the following recursion (over  $t$ ):

$$\mathbb{P}_{\tilde{\mathbf{M}}(b)}(\tilde{h}_t) = \sum_{\tilde{h}_{t-1}, \tilde{o}_{t-1}} \frac{\pi_{\tilde{h}_{t-1}, \tilde{o}_{t-1}, h_{t-1}(b)}^{o_{t-1}}(\tilde{h}_t, h_t(b))}{q_{h_{t-1}(b)}^{o_{t-1}, h_t(b)}} \times \frac{\theta_{\tilde{h}_{t-1}, \tilde{x}_{t-1}, h_{t-1}(b)}^{x_{t-1}}(\tilde{o}_{t-1}, o_{t-1})}{e_{h_{t-1}(b)}^{x_{t-1}, o_{t-1}}} \times \mathbb{P}_{\tilde{\mathbf{M}}(b)}(\tilde{h}_{t-1}).$$

The recursion breaks at  $t = 1$ :

$$\mathbb{P}_{\tilde{\mathbf{M}}(b)}(\tilde{h}_1) = \begin{cases} 1 & \text{if } \tilde{h}_1 = h_1(b) \\ 0 & \text{otherwise.} \end{cases}$$

Putting together Lemmas 1 and 2 allows us to express PN in terms of the various primitives, all of which except  $(\boldsymbol{\theta}, \boldsymbol{\pi})$  are known (or can be sampled). Thus, we use the notation  $\text{PN}(\boldsymbol{\theta}, \boldsymbol{\pi} \mid [h_{1:T}(b)]_b)$ . As soon as we fix  $(\boldsymbol{\theta}, \boldsymbol{\pi})$ , we can estimate PN. However, it is unclear apriori what we should fix  $(\boldsymbol{\theta}, \boldsymbol{\pi})$  at. We might have some information on the structure of  $(\boldsymbol{\theta}, \boldsymbol{\pi})$  that can help us shrink their feasibility space but in general, there can be many  $(\boldsymbol{\theta}, \boldsymbol{\pi})$ s that are “valid”. To overcome this lack of knowledge, we take an agnostic view and compute bounds over PN. The upper (lower) bound is computed by maximizing (minimizing)  $\text{PN}(\boldsymbol{\theta}, \boldsymbol{\pi} \mid [h_{1:T}(b)]_b)$  over the set of  $(\boldsymbol{\theta}, \boldsymbol{\pi})$  that are “valid”. Denoting by  $\mathcal{F}$  the set of “valid”  $(\boldsymbol{\theta}, \boldsymbol{\pi})$  (discussed below), we define

$$\text{PN}^{\text{ub}}(B) := \max_{(\boldsymbol{\theta}, \boldsymbol{\pi}) \in \mathcal{F}} \text{PN}(\boldsymbol{\theta}, \boldsymbol{\pi} \mid [h_{1:T}(b)]_b) \quad (8a)$$

$$\text{PN}^{\text{lb}}(B) := \min_{(\boldsymbol{\theta}, \boldsymbol{\pi}) \in \mathcal{F}} \text{PN}(\boldsymbol{\theta}, \boldsymbol{\pi} \mid [h_{1:T}(b)]_b). \quad (8b)$$

Both optimizations in (8) are sample average approximations (SAA) due to the use of the Monte-Carlo samples  $[h_{1:T}(b)]_b$ . Thus,  $\text{PN}^{\text{ub}}(B)$  and  $\text{PN}^{\text{lb}}(B)$  are estimates of the “true”  $\text{PN}^{\text{ub}}$  and  $\text{PN}^{\text{lb}}$ . However, given  $[h_{1:T}(b)]_b$  are iid samples, the following consistency result is immediate (cf. Proposition 5.2 in Shapiro et al. (2021)).

**Proposition 1.**  $(\text{PN}^{\text{lb}}(B), \text{PN}^{\text{ub}}(B))$  converges to  $(\text{PN}^{\text{lb}}, \text{PN}^{\text{ub}})$  w.p. 1 as  $B \rightarrow \infty$ .

In addition, we can characterize the asymptotics of  $(\text{PN}^{\text{lb}}(B), \text{PN}^{\text{ub}}(B))$  via results in the SAA theory and we refer the reader to §5.1.2 of Shapiro et al. (2021).

**Constraints (feasibility set  $\mathcal{F}$ ).** We now discuss the feasibility set  $\mathcal{F}$ . Recall from (4) that we used  $\boldsymbol{\theta}$  and  $\boldsymbol{\pi}$  to denote the pairwise marginal distributions over  $[O_{hx}]_{h,x}$  and  $[H_{hi}]_{h,i}$  respectively. We will now also use them to represent the full joint distributions of  $[O_{hx}]_{h,x}$  and  $[H_{hi}]_{h,i}$  respectively. To simplify notation, let  $k \equiv (h, x)$  and

$m \equiv (h, i)$ . Hence,

$$O_k \equiv O_{hx}, \quad e_{ki} \equiv e_{hxi} \\ H_m \equiv H_{hi}, \quad q_{mh'} \equiv q_{hih'}.$$

We have  $k \in [K]$  and  $m \in [M]$ , where  $K := |\mathbb{H}||\mathbb{X}|$  and  $M := |\mathbb{H}||\mathbb{O}|$ . The  $K$  and  $M$  dimensional joint PMFs for all  $i_1, \dots, i_K \in \mathbb{O}$  and  $h_1, \dots, h_M \in \mathbb{H}$  are defined as

$$\theta_{1, \dots, K}(i_1, \dots, i_K) := \mathbb{P}(O_1 = i_1, \dots, O_K = i_K) \quad (9a)$$

$$\pi_{1, \dots, M}(h_1, \dots, h_M) := \mathbb{P}(H_1 = h_1, \dots, H_M = h_M). \quad (9b)$$

Note that we only have *one* joint  $\theta_{1, \dots, K}$  among  $K$  random variables in contrast to *multiple* pairwise marginals  $[\theta_{k\ell}]_{(k, \ell)}$ . Each of these joint PMFs are decision variables in the optimization (in addition to the pairwise decision variables) and must obey the following set of constraints. First, the 1-dimensional marginals of  $\boldsymbol{\theta}$  and  $\boldsymbol{\pi}$  must equal the given 1-dimensional marginals  $[e_{ki}]_{(k, i)}$  and  $[q_{mh}]_{(m, h)}$ :

$$\sum_{\{i_1, \dots, i_K\} \setminus \{i_k\}} \theta_{1, \dots, K}(i_1, \dots, i_K) = e_{1i_k} \quad \forall i_k, k \quad (10a)$$

$$\sum_{\{h_1, \dots, h_M\} \setminus \{h_m\}} \pi_{1, \dots, M}(h_1, \dots, h_M) = q_{1h_m} \quad \forall h_m, m. \quad (10b)$$

Recall that  $(\mathbf{Q}, \mathbf{E})$ , i.e., the right-hand-sides of (10), are known. Moreover, since  $\mathbf{Q}$  and  $\mathbf{E}$  themselves define 1-dimensional probability distributions and therefore sum to 1, (10) ensures the same will be true of both the joint PMFs, i.e., they will also sum to 1. Second, we must link the pairwise marginals to the joints:

$$\theta_{k\ell}(i_k, i_\ell) = \sum_{\{i_1, \dots, i_K\} \setminus \{i_k, i_\ell\}} \theta_{1, \dots, K}(i_1, \dots, i_K) \quad (11a)$$

$$\pi_{mn}(h_m, h_n) = \sum_{\{h_1, \dots, h_M\} \setminus \{h_m, h_n\}} \pi_{1, \dots, M}(h_1, \dots, h_M). \quad (11b)$$

(11a) holds for all  $i_k, i_\ell \in \mathbb{O}$  and  $k, \ell \in [K]$  s.t.  $k < \ell$  whereas (11b) holds for all  $h_m, h_n \in \mathbb{H}$  and  $m, n \in [M]$  s.t.  $m < n$ . The “ $k < \ell$ ” and “ $m < n$ ” conditions avoid unnecessary duplication (recall (5) and (6)).<sup>5</sup> Finally, we need to ensure non-negativity:

$$\boldsymbol{\theta}, \boldsymbol{\pi} \geq 0 \quad (12)$$

<sup>5</sup>In fact, given (5) and (6), we do not need to define all pairwise marginals as decision variables but only for “ $k < \ell$ ” and “ $m < n$ ”. This is because if an optimization has two decision variables  $x$  and  $y$  and the constraint  $x = y$ , we can eliminate  $y$  and the constraint  $x = y$  by replacing  $y$  with  $x$  everywhere in the optimization.

where we now use  $(\boldsymbol{\theta}, \boldsymbol{\pi})$  to denote all of the corresponding, i.e., joint and pairwise, decision variables.

Let  $\mathcal{F}$  be the feasible region over  $(\boldsymbol{\theta}, \boldsymbol{\pi})$  defined by the constraints (10), (11), and (12). Observe that  $\text{PN}(\boldsymbol{\theta}, \boldsymbol{\pi})$  is a polynomial in  $(\boldsymbol{\theta}, \boldsymbol{\pi})$  (cf. Lemmas 1 and 2) and the constraints in  $\mathcal{F}$  are linear. Thus, each of the problems in (8) fall within the class of polynomial optimization (Anjos & Lasserre, 2011). Denoting by  $\text{PN}^*$  the PN under the true (unknown)  $(\boldsymbol{\theta}, \boldsymbol{\pi})$ , we obtain the following inequalities.

**Proposition 2.**  $\text{PN}^{\text{lb}} \leq \text{PN}^* \leq \text{PN}^{\text{ub}}$ .

**Enforcing CS and PM via linear constraints.** Suppose that at some time, the patient was in state  $h$ , the emission was  $i$ , followed by a transition to state  $h'$ . This maps to the realization  $H_{hi} = h'$ . For  $\tilde{h}' \neq h'$ , CS requires that if  $\mathbb{P}(H_{\tilde{h}i} = h') / \mathbb{P}(H_{hi} = h') \geq \mathbb{P}(H_{\tilde{h}i} = \tilde{h}') / \mathbb{P}(H_{hi} = \tilde{h}')$ , then  $\mathbb{P}(H_{\tilde{h}i} = \tilde{h}' | H_{hi} = h') = 0$ . Observe that the “if” condition is equivalent to  $q_{\tilde{h}ih'} / q_{hih'} \geq q_{\tilde{h}i\tilde{h}'} / q_{hi\tilde{h}'}$  and the LHS of “then” equals  $\pi_{\tilde{h}i,hi}(\tilde{h}', h') / q_{hih'}$ . Hence, for the state transitions, CS is equivalent to adding the following linear constraints for all  $(h, i, h', \tilde{h}, \tilde{i}, \tilde{h}')$ :

$$\pi_{\tilde{h}i,hi}(\tilde{h}', h') = 0 \text{ if } \frac{q_{\tilde{h}ih'}}{q_{hih'}} \geq \frac{q_{\tilde{h}i\tilde{h}'}}{q_{hi\tilde{h}'}}. \quad (13a)$$

Similarly, for emissions, CS can be modeled by adding the following linear constraints for all  $(h, x, i, \tilde{h}, \tilde{x}, \tilde{i})$ :

$$\theta_{\tilde{h}\tilde{x},hx}(\tilde{i}, i) = 0 \text{ if } \frac{e_{\tilde{h}\tilde{x}i}}{e_{hxi}} \geq \frac{e_{\tilde{h}\tilde{x}\tilde{i}}}{e_{h\tilde{x}\tilde{i}}}. \quad (13b)$$

Hence, we can characterize the space of all SCMs that obey CS, which is in contrast to picking just one such SCM (Oberst & Sontag, 2019). Enforcing CS naturally leads to tighter bounds, but the bounds may not be “legitimate” if the true  $(\boldsymbol{\theta}, \boldsymbol{\pi})$  does not satisfy CS. Denoting by  $\text{PN}_{\text{cs}}^{\text{ub}}$  and  $\text{PN}_{\text{cs}}^{\text{lb}}$  the bounds obtained by adding CS constraints (13) to the optimizations in (8), we have the following result.

**Proposition 3.**  $\text{PN}^{\text{lb}} \leq \text{PN}_{\text{cs}}^{\text{lb}} \leq \text{PN}_{\text{cs}}^{\text{ub}} \leq \text{PN}^{\text{ub}}$ .

PM can also be enforced via linear constraints. To see this, suppose the patient has in-situ cancer in period  $t$  which is not detected but the patient’s state remains at in-situ in period  $t + 1$ . Then, in the counterfactual world, if the cancer is detected in period  $t$ , then PM would require that the cancer can not be worse than in-situ in period  $t + 1$ , i.e.,

$$\mathbb{P}(H_{\tilde{h}i} = \tilde{h}' | H_{hi} = h') = 0$$

for  $h = 2$ ,  $i \in \{1, 2\}$ ,  $h' = 2$ ,  $\tilde{h} \in \{2, 4\}$ ,  $\tilde{i} = 4$ ,  $\tilde{h}' \in \{5, 7\}$ . There can be multiple such cases to consider and we can enforce all the PM constraints by setting the corresponding  $\pi_{\tilde{h}i,hi}(\tilde{h}', h')$  variables equal to 0 as  $\pi_{\tilde{h}i,hi}(\tilde{h}', h') = \mathbb{P}(H_{hi} = h')\mathbb{P}(H_{\tilde{h}i} = \tilde{h}' | H_{hi} = h')$ . As

with CS (Proposition 3), PM will result in bounds  $\text{PN}_{\text{pm}}^{\text{ub}}$  and  $\text{PN}_{\text{pm}}^{\text{lb}}$  tighter than  $\text{PN}^{\text{ub}}$  and  $\text{PN}^{\text{lb}}$ .

---

### Algorithm 1 Counterfactual analysis via optimization

---

**Require:**  $(\mathbf{E}, \mathbf{Q})$ ,  $(o_{1:T}, x_{1:T})$ ,  $B$ ,  $\tilde{x}_{1:T}$   
 1:  $h_{1:T}(b) \sim H_{1:T} | (o_{1:T}, x_{1:T}) \forall b = 1, \dots, B$   
 2:  $\text{PN}^{\text{ub}}(B) = \max_{(\boldsymbol{\theta}, \boldsymbol{\pi}) \in \mathcal{F}} \text{PN}(\boldsymbol{\theta}, \boldsymbol{\pi} | [h_{1:T}(b)]_b)$   
 3:  $\text{PN}^{\text{lb}}(B) = \min_{(\boldsymbol{\theta}, \boldsymbol{\pi}) \in \mathcal{F}} \text{PN}(\boldsymbol{\theta}, \boldsymbol{\pi} | [h_{1:T}(b)]_b)$   
 4: **return**  $(\text{PN}^{\text{lb}}(B), \text{PN}^{\text{ub}}(B))$

---

We summarize our developments in Algorithm 1, which outputs the bounds  $(\text{PN}^{\text{lb}}(B), \text{PN}^{\text{ub}}(B))$ <sup>6</sup>. Line 1 (sampling) can be executed efficiently (cf. §B), and we discuss three computational considerations behind solving the polynomial optimizations (lines 2 and 3).

First, though the constraints are linear, the objective is polynomial, making it a non-trivial non-convex optimization problem. To solve it, we leverage state-of-the-art developments in optimization. In particular, we use the BARON solver (Sahinidis, 2023), which relies on a polyhedral branch-and-cut approach, allowing it to achieve global optima (Tawarmalani & Sahinidis, 2005). We found it to work well in our numeric experiments (§5).

Second, in terms of the problem size, it follows from (4) and (9) that we have *at most*  $|\mathbb{H}|^4|\mathbb{O}|^2 + |\mathbb{H}|^2|\mathbb{O}|^2|\mathbb{X}|^2$  pairwise variables and  $|\mathbb{O}|^{|\mathbb{H}||\mathbb{X}|} + |\mathbb{H}|^{|\mathbb{H}||\mathbb{O}|}$  joint variables. Similarly, it follows from (10) and (11) that the feasible region  $\mathcal{F}$  is defined by *at most*  $|\mathbb{O}||\mathbb{H}||\mathbb{X}| + |\mathbb{H}|^2|\mathbb{O}| + |\mathbb{O}|^2|\mathbb{H}|^2|\mathbb{X}|^2 + |\mathbb{H}|^4|\mathbb{O}|^2$  constraints. However, these are merely upper bounds and we can exploit the sparsity inherent in the underlying application (along with the variable and constraint elimination discussed in Footnote 5) to drastically reduce the problem size. For instance, in our breast cancer application, we have  $(|\mathbb{H}|, |\mathbb{O}|, |\mathbb{X}|) = (7, 7, 2)$ , with the above formulae giving over  $10^{41}$  variables and  $10^5$  constraints. After we exploit sparsity (discussed in §5), they are reduced to 16,124 and 610, respectively. Further, as CS and PM can be modeled by setting appropriate variables to 0, they allow for further sparsity as we can delete those variables.

Third, observe that a naive expansion of the recursion in Lemma 2 results in a number of terms that is exponential in  $T$ , which would result in memory issues for moderate to large values of  $T$ . Nonetheless, as we elaborate in §D.1, it is possible to remove this exponential dependence on  $T$  by a reformulation of the optimization problem. This comes at the cost of introducing polynomial constraints. Nonetheless, this reformulation allowed us to obtain high-quality solutions in the breast cancer setting with as many as

---

<sup>6</sup>We can output  $(\text{PN}_{\text{cs}}^{\text{lb}}(B), \text{PN}_{\text{cs}}^{\text{ub}}(B))$  and  $(\text{PN}_{\text{pm}}^{\text{lb}}(B), \text{PN}_{\text{pm}}^{\text{ub}}(B))$  as well by solving the same optimization problems but with additional linear constraints.

$T = 100$  periods (§D.3). In contrast, we run into memory issues for  $T$  as small as 11 with the original formulation. In fact, we discuss an alternative approach at the end of §D.1. This approach allows us to compute the objective function efficiently without having to add any additional constraints. Unfortunately, the BARON solver does not allow us to use this approach and so we leave this issue for future research.

## 5. Numerical Experiments

We now apply our approach to the breast cancer application we described in §1.

**Setup.** We described the elements of the underlying dynamic latent-state model  $\mathbf{M} \equiv (\mathbf{p}, \mathbf{E}, \mathbf{Q})$  in §2. It has a total of 7 states, 7 emissions, and 2 actions. Given patient-level data  $(o_{1:T}, x_{1:T})$ , we wish to estimate the PN as defined in (2). The primitives  $(\mathbf{p}, \mathbf{E}, \mathbf{Q})$  are calibrated to real-data using a mix of sources, which we discuss in §E.1.

We consider two paths with the first path defined as:

$$\underbrace{o_1}_{=2}, \underbrace{o_2, \dots, o_{T-2}, o_{T-1}}_{=1}, \underbrace{o_T}_{=7} \\ \underbrace{x_1}_{=0}, \underbrace{x_2, \dots, x_{T-2}, x_{T-1}}_{=1}, \underbrace{x_T}_{=0}.$$

That is, we observe a negative test result in period 1 after which screening was not performed for  $T - 2$  periods (red font). The patient died from breast cancer in period  $T$ . Note that under this path, given the calibrated primitives in §E.1, it has to be the case that  $h_{T-1} = 3$  (undiagnosed invasive) since a transition from state 2 (undiagnosed in-situ) to 7 is impossible. Further, the transition from 3 to 7 is not unlikely ( $q_{317} \approx 0.15$ ). The second path is similar but with one difference: screening was performed in period  $T - 1$  and invasive cancer was detected:

$$\underbrace{o_1}_{=2}, \underbrace{o_2, \dots, o_{T-2}}_{=1}, \underbrace{o_{T-1}}_{=5}, \underbrace{o_T}_{=7} \\ \underbrace{x_1}_{=0}, \underbrace{x_2, \dots, x_{T-2}}_{=1}, \underbrace{x_{T-1}, x_T}_{=0}.$$

Hence, in contrast with path 1, the final transition from invasive to death was under treatment with probability  $q_{357} \approx 0.01$  (§E.1), which is much smaller than  $q_{317}$  from above. Given that this low probability transition did occur, this suggests the patient had an “aggressive” cancer in path 2. As such, regardless of what the optimal  $\theta$  and  $\pi$  are, the chances of survival on the counterfactual path would be low because of this “aggressive” nature of the cancer. This doesn’t hold on path 1 as the cancer was less “aggressive”.

We vary  $T \in \{4, \dots, 10\}$ , with a larger value of  $T$  suggesting the cancer may have progressed more slowly. We

compute PN bounds using our framework (Algorithm 1), which we implemented in MATLAB (MATLAB, 2021). The feasibility set  $\mathcal{F}$  over  $(\theta, \pi)$  corresponds to (10), (11), and (12). To solve the polynomial optimizations, we use the MATLAB-BARON interface (Sahinidis, 2023) with CPLEX (IBM, 2017) as the “LP / MIP solver”. It solved each of our problem instances to global optimality within minutes / hours (depending on  $T$ ), with an “absolute termination tolerance” of 0.01 (on an Intel Xeon E5 processor with 16 GB RAM). Optimizations for  $T = 10$  took the longest time on average ( $\sim 2$  hours). We generated  $B = 100$  samples using our sampling method in §B. It took less than a second and we found  $B = 100$  was large enough to produce stable results for our SAA. We ensured this stability by computing our results for 20 seeds (for each (path,  $T$ ) pair) and verifying the standard deviations to be small. Though stability over the seeds is important, our PN estimates may still be biased for a finite  $B$  (recall Proposition 1 only holds asymptotically). As a check, we also generated results for  $B = 500$  and observed them to be very similar to the ones for  $B = 100$ . As noted below Algorithm 1, the sparse structure of  $\mathbf{E}$  and  $\mathbf{Q}$  drastically reduces the size of the problem. For example, when considering the  $\pi_{\tilde{h}i, hi}(\tilde{h}', h')$  variables, we rule out the ones that map to impossible  $(h, i, h')$  or  $(\tilde{h}, \tilde{i}, \tilde{h}')$  combinations (refer to §E.1.2). The same observation also applies to all the joint variables (details in §E.2).

**Results.** The results for path 1 are displayed in Figure 3 (and for path 2 in Figure 10 (§E.5)), where we show the PN bounds as we vary  $T$ . In addition to the bounds ( $\text{PN}^{\text{rb}}, \text{PN}^{\text{ub}}$ ) computed via our baseline optimization (UB and LB), we show the bounds obtained when we encode CS (UB (CS) and LB (CS)) and PM (UB (PM) and LB (PM))<sup>7</sup>. We also show the PN estimate when we perform counterfactual simulations using the two copulas discussed in §C (independence and comonotonic<sup>8</sup>). Finally, the naive estimate completely ignores the information in the observations, i.e., it does not execute the abduction step and is therefore an invalid estimate of PN.

To simplify matters, we adopt an all-or-nothing approach whereby either CS is imposed for both hidden-state transitions *and* observations or not at all. We do the same for PM. Of course, it is possible to consider various combinations, e.g., imposing PM for hidden-state transitions only or imposing CS only for the observations, etc. This is also true of our copulas when we estimate PN for a particular SCM. In Figure 3, for example, the independence (comonotonic) curve corresponds to assuming the independence (comonotonic) copula for both hidden-state transitions *and* observations. But we could of course have assumed one copula for

<sup>7</sup>Details on the PM constraints for breast cancer are in §E.3.

<sup>8</sup>Further details on the comonotonic copula specific to the breast cancer model are in §E.4.






Figure 3. PN results for path 1 as we vary  $T \in \{4, \dots, 10\}$ . Observe that the LB, LB (CS), and LB (PM) curves coincide (the lowest curve in each figure). We show the average over 20 seeds and note that the standard deviation (s.d.) for every data point is smaller than 0.01. Given this small magnitude, we omit the  $\pm 1$  s.d. bars to reduce the clutter in the sub-figures. Note that to simulate the naive estimate and the two copulas (independence and comonotonic), we use  $10^4$  Monte-Carlo samples. These simulations were fast (a couple of minutes). Using  $10^4$  samples is in contrast to the 100 samples we use for SAA and we did so to ensure a low s.d.

the hidden-state transitions and an entirely different one for the observations. Each such combination of copulas would yield a different SCM and therefore a *feasible* value of PN.

The naive estimate is independent of the observed path and can fall outside the bounds. This makes sense as it does not perform abduction but simply simulates the original model  $M$  under the intervention policy  $\tilde{x}_{1:T} = 0$ . The naive estimates are very close to 1 as dying of breast cancer in any 5-year period<sup>9</sup> is highly unlikely.

For path 1, we obtain relatively tight bounds, with PN always above 0.85. This means that in the counterfactual world, the patient would have not died with high probability, consistent with our discussion around  $q_{317}$  above. Even in the absence of any additional structure such as CS or PM, the gap between the lower and upper bounds is within  $\sim 10$  percentage points. The gap gets tighter with CS (within  $\sim 5$  percentage points) and PM (within  $\sim 1$  percentage point!). The fact that the LB and UB under CS do not coincide resolves the open question of Oberst & Sontag (2019) regarding the uniqueness of the Gumbel-max mechanism w.r.t. CS – it is not unique. It is not surprising that the comonotonic estimate falls close to the PM bounds. Interestingly, the estimated PN for the two copulas roughly cover the range of possibilities in terms of the bounds (Figure 3(a)).

For path 2, the lower bounds are close to 0. This aligns with the fact that despite being diagnosed in period  $T - 1$  (and hence, provided treatment), the patient eventually died (which suggests that the patient had an “aggressive” cancer). The bounds without CS and PM are relatively loose, simply reflecting the lack of knowledge to reason in a counterfactual world. As soon as we inject knowledge via CS or PM, the bounds become much tighter.

The experiments discussed so far are for up to  $T = 10$  and we run into memory issues for  $T > 10$  (recall the discussion

<sup>9</sup>Each period maps to 6 months so  $T = 10$  maps to 5 years.

at the end of §4.3). Nonetheless, as we show in §D, we can enhance the scalability of the polynomial optimizations in (8) via a reformulation and an approximation. In fact, as we demonstrate via numerics, these ideas allow us to obtain high-quality solutions for  $T$  as large as 100 in just a few hours of compute time.

## 6. Concluding Remarks

We have provided a framework for performing counterfactual analysis in dynamic latent-state models and in particular, computing lower and upper bounds on CQIs. There are several interesting directions for future research. First, we would like to handle the objective function in the optimization more efficiently as discussed at the end of §4.3. Specifically, BARON’s solver appears to explicitly expand the objective function which results in a number of terms that is exponential in  $T$ . We were able to finesse this issue in §D via a reformulation but we suspect the approach outlined at the end of §4.3 might provide a better solution. All told, it may therefore be worthwhile developing an optimization algorithm specifically tailored to the problem (a polynomial objective with linear constraints) rather than using an off-the-shelf solver. Another possible direction is exploring the use of variance reduction methods and other Monte-Carlo techniques to improve our basic Monte Carlo approach for generating posterior sample paths. Finally, on the practical front, it would be of interest to apply our framework to real-world medical applications and use domain-specific knowledge to obtain (via the imposition of additional constraints) tighter bounds on the CQIs.

## Acknowledgements

We thank the ICML review team, Madhumitha Shridharan, and Jim Smith for taking the time to read the paper and providing very useful feedback. We also thank Nick Sahinidis for his support with BARON-related issues.

## References

- Anjos, M. F. and Lasserre, J. B. *Handbook on semidefinite, conic and polynomial optimization*, volume 166. Springer Science & Business Media, 2011.
- Ayer, T., Alagoz, O., and Stout, N. K. A POMDP approach to personalize mammography screening decisions. *Operations Research*, 60(5):1019–1034, 2012.
- Balke, A. and Pearl, J. Counterfactual probabilities: Computational methods, bounds and applications. In *Uncertainty Proceedings*, pp. 46–54. San Francisco (CA), 1994.
- Barber, D. *Bayesian Reasoning and Machine Learning*. Cambridge University Press, 2012.
- Buesing, L., Weber, T., Zwols, Y., Heess, N., Racaniere, S., Guez, A., and Lespiau, J.-B. Woulda, coulda, shoulda: Counterfactually-guided policy search. In *International Conference on Learning Representations*, 2019.
- Cai, Z., Kuroki, M., Pearl, J., and Tian, J. Bounds on direct effects in the presence of confounded intermediate variables. *Biometrics*, 64(3):695–701, 2008.
- Haugh, M. B. and Lacedelli, O. R. Information relaxation bounds for partially observed Markov decision processes. *IEEE Transactions on Automatic Control*, 65(8):3256–3271, 2019.
- IBM. ILOG CPLEX Optimizer Version 12.8. 2017.
- Johnstone, P. A., Norton, M. S., and Riffenburgh, R. H. Survival of patients with untreated breast cancer. *Journal of surgical oncology*, 73(4):273–277, 2000.
- Kaufman, S., Kaufman, J., MacLenose, R., Greenland, S., and Poole, C. Improved estimation of controlled direct effects in the presence of unmeasured confounding of intermediate variables. *Statistics in Medicine*, 25:1683–1702, 2005.
- Lorberbom, G., Johnson, D. D., Maddison, C. J., Tarlow, D., and Hazan, T. Learning generalized gumbel-max causal mechanisms. In *Advances in Neural Information Processing Systems*, volume 34, pp. 26792–26803, 2021.
- MATLAB. *Version 9.10.0 (R2021b)*. The MathWorks Inc., Natick, Massachusetts, 2021.
- McNeil, A. J., Frey, R., and Embrechts, P. *Quantitative Risk Management: Concepts, Techniques and Tools*. Princeton University Press, 2 edition, 2015.
- Mueller, S., Li, A., and Pearl, J. Causes of effects: Learning individual responses from population data. *arXiv*, 2021.
- Nelsen, R. *An Introduction to Copulas*. Springer, 2 edition, 2006.
- NIH. SEER Cancer Statistics Review (CSR) 1975-2017. 2020. URL [https://seer.cancer.gov/archive/csr/1975\\_2017/results\\_merged/sect\\_04\\_breast.pdf](https://seer.cancer.gov/archive/csr/1975_2017/results_merged/sect_04_breast.pdf).
- Oberst, M. and Sontag, D. Counterfactual off-policy evaluation with Gumbel-max structural causal models. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 4881–4890. PMLR, 09–15 Jun 2019.
- Pearl, J. Causal inference in statistics: An overview. *Statistics surveys*, 3:96–146, 2009a.
- Pearl, J. *Causality*. Cambridge University Press, 2 edition, 2009b.
- Pearl, J. and Mackenzie, D. *The Book of Why*. Penguin Books, 2018.
- Sahinidis, N. V. *BARON 2023.1.5: Global Optimization of Mixed-Integer Nonlinear Programs*, User’s Manual, 2023.
- Shapiro, A., Dentcheva, D., and Ruszczyński, A. *Lectures on stochastic programming: modeling and theory*. SIAM, 2021.
- Sklar, A. Fonctions de répartition à n dimensions et leurs marges. *Publ. Inst. Statist. Univ. Paris*, 8:229–231, 1959.
- Sprague, B. L. and Trentham-Dietz, A. Prevalence of breast carcinoma in situ in the United States. *JAMA: the journal of the American Medical Association*, 302(8):846, 2009.
- Tawarmalani, M. and Sahinidis, N. V. A polyhedral branch-and-cut approach to global optimization. *Mathematical programming*, 103(2):225–249, 2005.
- Tian, J. and Pearl, J. Probabilities of causation: Bounds and identification. *Annals of Mathematics and Artificial Intelligence*, 8:287–313, 2000.
- Tsirtsis, S., De, A., and Rodriguez, M. Counterfactual explanations in sequential decision making under uncertainty. In Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 30127–30139. Curran Associates, Inc., 2021.
- UWBCS. University of Wisconsin Breast Cancer Simulation Model. 2013. URL <https://resources.cisnet.cancer.gov/registry/packages/uwbc-wisconsin/>.
- Zhang, J., Tian, J., and Bareinboim, E. Partial counterfactual identification from observational and experimental data, 2021.

## A. Proofs

**Lemma 1.** *We have*

$$PN = 1 - \lim_{B \rightarrow \infty} \frac{1}{B} \sum_{b=1}^B \mathbb{P}_{\tilde{\mathbf{M}}(b)}(\tilde{H}_T = 7).$$

**Proof.** Observe that

$$\begin{aligned} PN &= \mathbb{P}_{\tilde{H}_T}(\tilde{H}_T \neq 7) && \text{[by definition]} \\ &= 1 - \mathbb{P}_{\tilde{H}_T}(\tilde{H}_T = 7) && [\mathbb{P}(Y \neq y) = 1 - \mathbb{P}(Y = y)] \\ &= 1 - \mathbb{E}_{\tilde{H}_T}[\mathbb{I}\{\tilde{H}_T = 7\}] && [\mathbb{P}(Y = y) = \mathbb{E}[\mathbb{I}\{Y = y\}]] \\ &= 1 - \mathbb{E}_{H_{1:T}}[\mathbb{E}_{\tilde{\mathbf{M}}|H_{1:T}}[\mathbb{I}\{\tilde{H}_T = 7\}]] && \text{[law of total expectation]} \\ &= 1 - \frac{1}{B} \sum_{b=1}^B \mathbb{P}_{\tilde{\mathbf{M}}(b)}(\tilde{H}_T = 7) \text{ as } B \rightarrow \infty. && \text{[law of large numbers]} \end{aligned}$$

The proof is now complete.  $\square$

**Lemma 2.** *For  $t \in \{T, T-1, \dots, 2\}$ ,  $\mathbb{P}_{\tilde{\mathbf{M}}(b)}(\tilde{h}_t) := \mathbb{P}_{\tilde{\mathbf{M}}(b)}(\tilde{H}_t = \tilde{h}_t)$  obeys the following recursion (over  $t$ ):*

$$\begin{aligned} \mathbb{P}_{\tilde{\mathbf{M}}(b)}(\tilde{h}_t) &= \sum_{\tilde{h}_{t-1}, \tilde{o}_{t-1}} \frac{\pi_{\tilde{h}_{t-1}\tilde{o}_{t-1}, h_{t-1}(b) o_{t-1}}(\tilde{h}_t, h_t(b))}{q_{h_{t-1}(b) o_{t-1} h_t(b)}} \times \\ &\quad \frac{\theta_{\tilde{h}_{t-1}\tilde{x}_{t-1}, h_{t-1}(b) x_{t-1}}(\tilde{o}_{t-1}, o_{t-1})}{e_{h_{t-1}(b) x_{t-1} o_{t-1}}} \times \\ &\quad \mathbb{P}_{\tilde{\mathbf{M}}(b)}(\tilde{h}_{t-1}). \end{aligned}$$

The recursion breaks at  $t = 1$ :

$$\mathbb{P}_{\tilde{\mathbf{M}}(b)}(\tilde{h}_1) = \begin{cases} 1 & \text{if } \tilde{h}_1 = h_1(b) \\ 0 & \text{otherwise.} \end{cases}$$

**Proof.** For  $t \in \{T, T-1, \dots, 2\}$ , observe that

$$\begin{aligned} \mathbb{P}_{\tilde{\mathbf{M}}(b)}(\tilde{h}_t) &= \sum_{\tilde{h}_{t-1} \in \mathbb{H}} \sum_{\tilde{o}_{t-1} \in \mathbb{O}} \mathbb{P}_{\tilde{\mathbf{M}}(b)}(\tilde{h}_t, \tilde{h}_{t-1}, \tilde{o}_{t-1}) \\ &= \sum_{\tilde{h}_{t-1} \in \mathbb{H}} \sum_{\tilde{o}_{t-1} \in \mathbb{O}} \mathbb{P}_{\tilde{\mathbf{M}}(b)}(\tilde{h}_t | \tilde{h}_{t-1}, \tilde{o}_{t-1}) \mathbb{P}_{\tilde{\mathbf{M}}(b)}(\tilde{o}_{t-1} | \tilde{h}_{t-1}) \mathbb{P}_{\tilde{\mathbf{M}}(b)}(\tilde{h}_{t-1}) \\ &= \sum_{\tilde{h}_{t-1} \in \mathbb{H}} \sum_{\tilde{o}_{t-1} \in \mathbb{O}} \tilde{q}_{\tilde{h}_{t-1}\tilde{o}_{t-1}\tilde{h}_t}^{(t-1)}(b) \times \tilde{e}_{\tilde{h}_{t-1}\tilde{x}_{t-1}\tilde{o}_{t-1}}^{(t-1)}(b) \times \mathbb{P}_{\tilde{\mathbf{M}}(b)}(\tilde{h}_{t-1}) \\ &= \sum_{\tilde{h}_{t-1} \in \mathbb{H}} \sum_{\tilde{o}_{t-1} \in \mathbb{O}} \frac{\pi_{\tilde{h}_{t-1}\tilde{o}_{t-1}, h_{t-1}(b) o_{t-1}}(\tilde{h}_t, h_t(b))}{q_{h_{t-1}(b) o_{t-1} h_t(b)}} \times \frac{\theta_{\tilde{h}_{t-1}\tilde{x}_{t-1}, h_{t-1}(b) x_{t-1}}(\tilde{o}_{t-1}, o_{t-1})}{e_{h_{t-1}(b) x_{t-1} o_{t-1}}} \times \mathbb{P}_{\tilde{\mathbf{M}}(b)}(\tilde{h}_{t-1}). \end{aligned}$$

The base case ( $t = 1$ ) holds since the counterfactual hidden state in period 1 equals the posterior sample  $h_1(b)$  (recall from §4.2). The proof is now complete.  $\square$

## B. Sampling Hidden Paths from the Posterior Distribution

In this section, we show how one can efficiently perform filtering, smoothing, and sampling for the dynamic latent-state model in Figure 1. As our model is a generalization of an HMM, these algorithms are simple generalizations of the standard variants corresponding to an HMM (Barber, 2012).

**Filtering.** We first compute  $\alpha(h_t) := \mathbb{P}(h_t, o_{1:t}, x_{1:t})$  which will yield the un-normalized filtered posterior distribution. We can then easily normalize it to compute  $\mathbb{P}(h_t | o_{1:t}, x_{1:t}) \propto \alpha(h_t)$ . We begin with  $\alpha(h_1) := \mathbb{P}(o_1 | h_1, x_1)\mathbb{P}(h_1 | x_1)\mathbb{P}(x_1) = \mathbb{P}(o_1 | h_1, x_1)\mathbb{P}(h_1)\mathbb{P}(x_1)$ . For  $t > 1$ , note that

$$\begin{aligned} \alpha(h_t) &= \sum_{h_{t-1}} \mathbb{P}(h_t, h_{t-1}, o_{1:t-1}, o_t, x_{1:t}) \\ &= \sum_{h_{t-1}} \mathbb{P}(o_t | h_t, h_{t-1}, o_{1:t-1}, x_{1:t}) \mathbb{P}(h_t | h_{t-1}, o_{1:t-1}, x_{1:t}) \mathbb{P}(x_t | h_{t-1}, o_{1:t-1}, x_{1:t-1}) \mathbb{P}(h_{t-1}, o_{1:t-1}, x_{1:t-1}) \\ &= \sum_{h_{t-1}} \mathbb{P}(o_t | h_t, x_t) \mathbb{P}(h_t | h_{t-1}, o_{t-1}) \mathbb{P}(x_t) \mathbb{P}(h_{t-1}, o_{1:t-1}, x_{1:t-1}) \\ &= \mathbb{P}(x_t) \mathbb{P}(o_t | h_t, x_t) \sum_{h_{t-1}} \mathbb{P}(h_t | h_{t-1}, o_{t-1}) \alpha(h_{t-1}). \end{aligned}$$

**Smoothing.** We now compute  $\beta(h_t) := \mathbb{P}(o_{t+1:T}, x_{t+1:T} | h_t, o_t)$  with the understanding that  $\beta(h_T) = 1$ . For  $t < T$ , we have

$$\begin{aligned} \beta(h_t) &= \sum_{h_{t+1}} \mathbb{P}(o_{t+1}, x_{t+1}, o_{t+2:T}, x_{t+2:T}, h_{t+1} | h_t, o_t) \\ &= \sum_{h_{t+1}} \mathbb{P}(o_{t+2:T}, x_{t+2:T} | h_t, o_t, h_{t+1}, o_{t+1}, x_{t+1}) \mathbb{P}(h_{t+1}, o_{t+1}, x_{t+1} | h_t, o_t) \\ &= \sum_{h_{t+1}} \mathbb{P}(o_{t+2:T}, x_{t+2:T} | h_{t+1}, o_{t+1}) \mathbb{P}(o_{t+1} | h_{t+1}, x_{t+1}, h_t, o_t) \mathbb{P}(h_{t+1}, x_{t+1} | h_t, o_t) \\ &= \sum_{h_{t+1}} \beta(h_{t+1}) \mathbb{P}(o_{t+1} | h_{t+1}, x_{t+1}) \mathbb{P}(x_{t+1} | h_{t+1}, h_t, o_t) \mathbb{P}(h_{t+1} | h_t, o_t) \\ &= \mathbb{P}(x_{t+1}) \sum_{h_{t+1}} \beta(h_{t+1}) \mathbb{P}(o_{t+1} | h_{t+1}, x_{t+1}) \mathbb{P}(h_{t+1} | h_t, o_t). \end{aligned}$$

Now, note that

$$\begin{aligned} \mathbb{P}(h_t, o_{1:T}, x_{1:T}) &= \mathbb{P}(h_t, o_{1:t}, x_{1:t}) \mathbb{P}(o_{t+1:T}, x_{t+1:T} | h_t, o_{1:t}, x_{1:t}) \\ &= \mathbb{P}(h_t, o_{1:t}, x_{1:t}) \mathbb{P}(o_{t+1:T}, x_{t+1:T} | h_t, o_t) \\ &= \alpha(h_t) \beta(h_t). \end{aligned}$$

We therefore obtain the hidden state marginal

$$\mathbb{P}(h_t | o_{1:T}, x_{1:T}) = \frac{\alpha(h_t) \beta(h_t)}{\sum_{h_t} \alpha(h_t) \beta(h_t)},$$

which solves the smoothing problem.

**Pairwise marginal.** We can compute  $\mathbb{P}(h_t, h_{t+1} | o_{1:T}, x_{1:T})$  by noting that  $\mathbb{P}(h_t, h_{t+1} | o_{1:T}, x_{1:T})$  is proportional to

$$\begin{aligned} &\propto \mathbb{P}(o_{1:t}, o_{t+1}, o_{t+2:T}, x_{1:t}, x_{t+1}, x_{t+2:T}, h_{t+1}, h_t) \\ &= \mathbb{P}(o_{t+2:T}, x_{t+2:T} | o_{1:t}, o_{t+1}, x_{1:t}, x_{t+1}, h_{t+1}, h_t) \mathbb{P}(o_{1:t}, o_{t+1}, x_{1:t}, x_{t+1}, h_{t+1}, h_t) \\ &= \mathbb{P}(o_{t+2:T}, x_{t+2:T} | h_{t+1}, o_{t+1}) \mathbb{P}(o_{t+1} | o_{1:t}, h_{t+1}, h_t, x_{1:t}, x_{t+1}) \mathbb{P}(o_{1:t}, h_{t+1}, h_t, x_{1:t}, x_{t+1}) \\ &= \mathbb{P}(o_{t+2:T}, x_{t+2:T} | h_{t+1}, o_{t+1}) \mathbb{P}(o_{t+1} | h_{t+1}, x_{t+1}) \mathbb{P}(h_{t+1}, x_{t+1} | o_{1:t}, x_{1:t}, h_t) \mathbb{P}(o_{1:t}, x_{1:t}, h_t) \\ &= \mathbb{P}(o_{t+2:T}, x_{t+2:T} | h_{t+1}, o_{t+1}) \mathbb{P}(o_{t+1} | h_{t+1}, x_{t+1}) \mathbb{P}(h_{t+1}, x_{t+1} | h_t, o_t) \mathbb{P}(o_{1:t}, x_{1:t}, h_t). \end{aligned} \quad (16)$$

We can rearrange (16) to obtain

$$\mathbb{P}(h_t, h_{t+1} | o_{1:T}, x_{1:T}) \propto \alpha(h_t) \mathbb{P}(o_{t+1} | h_{t+1}, x_{t+1}) \mathbb{P}(h_{t+1}, x_{t+1} | h_t, o_t) \beta(h_{t+1}). \quad (17)$$

Therefore,  $\mathbb{P}(h_t, h_{t+1} | o_{1:T}, x_{1:T})$  is easy to compute once the forward-backward, i.e. the filtering and smoothing, recursions have been completed.

**Sampling.** We would like to sample from the posterior  $\mathbb{P}(h_{1:T} \mid o_{1:T}, x_{1:T})$ . We can do this by first noting that

$$\begin{aligned} \mathbb{P}(h_{1:T} \mid o_{1:T}, x_{1:T}) &= \mathbb{P}(h_1 \mid h_{2:T}, o_{1:T}, x_{1:T}) \dots \mathbb{P}(h_{T-1} \mid h_T, o_{1:T}, x_{1:T}) \mathbb{P}(h_T \mid o_{1:T}, x_{1:T}) \\ &= \mathbb{P}(h_1 \mid h_2, o_{1:T}, x_{1:T}) \dots \mathbb{P}(h_{T-1} \mid h_T, o_{1:T}, x_{1:T}) \mathbb{P}(h_T \mid o_{1:T}, x_{1:T}). \end{aligned}$$

We can therefore sample sequentially via the following two steps:

- First, draw  $h_T$  from  $\mathbb{P}(h_T \mid o_{1:T}, x_{1:T})$ , which we know from the smoothed distribution of  $h_T$ .
- Second, observe that for any  $t < T$ , we have

$$\begin{aligned} \mathbb{P}(h_t \mid h_{t+1}, o_{1:T}, x_{1:T}) &\propto \mathbb{P}(h_t, h_{t+1} \mid o_{1:T}, x_{1:T}) \\ &\propto \alpha(h_t) \mathbb{P}(h_{t+1}, x_{t+1} \mid h_t, o_t) && \text{[by (17)]} \\ &= \alpha(h_t) \mathbb{P}(h_{t+1} \mid h_t, o_t) \mathbb{P}(x_{t+1}), \end{aligned}$$

from which it is easy to sample.

Hence, we can efficiently generate samples  $[h_{1:T}(b)]_{b=1}^B$  from the posterior  $\mathbb{P}(h_{1:T} \mid o_{1:T}, x_{1:T})$ .

## C. A Brief Introduction to Copulas and Counterfactual Simulations

Copulas are functions that enable us to separate the marginal distributions from the dependency structure of a given multivariate distribution. They are particularly useful in applications where the marginal distributions are known (either from domain specific knowledge or because there is sufficient marginal data) but a joint distribution with these known marginals is required. In our application in this paper, we know the marginal distribution of each random variable in  $[O_{hx}]_{h,x}$  and  $[H_{hi}]_{h,i}$ , which is dictated by the model primitives  $(\mathbf{E}, \mathbf{Q})$  as follows:  $e_{hxi} = \mathbb{P}(O_{hx} = i)$  and  $q_{hih'} = \mathbb{P}(H_{hi} = h')$ . Indeed, these marginal distributions can be estimated from data, but the *joint distribution* must be specified in order to compute counterfactuals.

In each of these cases, one needs to work with a joint distribution with fixed or pre-specified marginal distributions. Copulas and Sklar's Theorem (see below) can be very helpful in these situations. We only briefly review some of the main results from the theory of copulas here but [Nelsen \(2006\)](#) can be consulted for an introduction to the topic. [McNeil et al. \(2015\)](#) also contains a nice introduction but in the context of financial risk management.

**Definition 1.** A  $d$ -dimensional copula,  $C : [0, 1]^d \rightarrow [0, 1]$  is a cumulative distribution function with uniform marginals.

We write  $C(\mathbf{u}) = C(u_1, \dots, u_d)$  for a generic copula. It follows immediately from Definition 1 that  $C(u_1, \dots, u_d)$  is non-decreasing in each argument and that  $C(1, \dots, 1, u_i, 1, \dots, 1) = u_i$ . It is also easy to confirm that  $C(1, u_1, \dots, u_{d-1})$  is a  $(d-1)$ -dimensional copula and, more generally, that all  $k$ -dimensional marginals with  $2 \leq k \leq d$  are copulas. The most important result from the theory of copulas is Sklar's Theorem ([Sklar, 1959](#)).

**Theorem 1 (Sklar 1959).** Consider a  $d$ -dimensional CDF  $\mathbf{\Pi}$  with marginals  $\mathbf{\Pi}_1, \dots, \mathbf{\Pi}_d$ . Then, there exists a copula  $C$  such that

$$\mathbf{\Pi}(x_1, \dots, x_d) = C(\mathbf{\Pi}_1(x_1), \dots, \mathbf{\Pi}_d(x_d)) \quad (18)$$

for all  $x_i \in [-\infty, \infty]$  and  $i = 1, \dots, d$ .

If  $\mathbf{\Pi}_i$  is continuous for all  $i = 1, \dots, d$ , then  $C$  is unique; otherwise  $C$  is uniquely determined only on  $\text{Ran}(\mathbf{\Pi}_1) \times \dots \times \text{Ran}(\mathbf{\Pi}_d)$ , where  $\text{Ran}(\mathbf{\Pi}_i)$  denotes the range of the CDF  $\mathbf{\Pi}_i$ .

Conversely, consider a copula  $C$  and univariate CDF's  $\mathbf{\Pi}_1, \dots, \mathbf{\Pi}_d$ . Then,  $\mathbf{\Pi}$  as defined in (18) is a multivariate CDF with marginals  $\mathbf{\Pi}_1, \dots, \mathbf{\Pi}_d$ .

A particularly important aspect of Sklar's Theorem in the context of this paper is that  $C$  is only uniquely determined on  $\text{Ran}(\mathbf{\Pi}_1) \times \dots \times \text{Ran}(\mathbf{\Pi}_d)$ . Because we are interested in applications with discrete state-spaces, this implies that there will be many copulas that lead to the same joint distribution  $\mathbf{\Pi}$ . It is for this reason that we prefer to work directly with the joint distribution of  $[O_{hx}]_{h,x}$  and  $[H_{hi}]_{h,i}$  (recall (4)). That said, we emphasize that specifying copulas for the exogenous vectors  $\mathbf{U}_t$  and  $\mathbf{V}_t$  is equivalent to specifying a particular structural causal model (SCM) in which any CQI can be computed.

The following important result was derived independently by Fréchet and Hoeffding and provides lower and upper bounds on copulas.

**Theorem 2 (The Fréchet-Hoeffding Bounds).** *Consider a copula  $C(\mathbf{u}) = C(u_1, \dots, u_d)$ . Then,*

$$\max \left\{ 1 - d + \sum_{i=1}^d u_i, 0 \right\} \leq C(\mathbf{u}) \leq \min\{u_1, \dots, u_d\}.$$

Three important copulas are the comonotonic, countermonotonic (only when  $d = 2$ ) and independence copulas which model extreme positive dependency, extreme negative dependency and (not surprisingly) independence. They are defined as follows.

**Comonotonic Copula.** The comonotonic copula is given by

$$C^P(\mathbf{u}) := \min\{u_1, \dots, u_d\}, \quad (19)$$

which coincides with the Fréchet-Hoeffding upper bound. It corresponds to the case of extreme positive dependence. For example, let  $\mathbf{U} = (U_1, \dots, U_d)$  with  $U_1 = U_2 = \dots = U_d \sim \text{Unif}[0, 1]$ . Then, clearly  $\min\{u_1, \dots, u_d\} = \mathbf{\Pi}(u_1, \dots, u_d)$  but by Sklar's Theorem  $F(u_1, \dots, u_d) = C(u_1, \dots, u_d)$  and so,  $C(u_1, \dots, u_d) = \min\{u_1, \dots, u_d\}$ .

**Countermonotonic Copula.** The countermonotonic copula is a 2-dimensional copula given by

$$C^N(\mathbf{u}) := \max\{u_1 + u_2 - 1, 0\}, \quad (20)$$

which coincides with the Fréchet-Hoeffding lower bound when  $d = 2$ . It corresponds to the case of extreme negative dependence. It is easy to check that (20) is the joint distribution of  $(U, 1 - U)$  where  $U \sim \text{Unif}[0, 1]$ . (The Fréchet-Hoeffding lower bound is only tight when  $d = 2$ . This is analogous to the fact that while a pairwise correlation can lie anywhere in  $[-1, 1]$ , the *average* pairwise correlation of  $d$  random variables is bounded below by  $-1/(d - 1)$ .)

**Independence Copula.** The independence copula satisfies

$$C^I(\mathbf{u}) := \prod_{i=1}^d u_i,$$

and it is easy to confirm using Sklar's Theorem that random variables are independent if and only if their copula is the independence copula.

A well known and important result regarding copulas is that they are invariant under monotonic transformations.

**Proposition 4 (Invariance Under Monotonic Transformations).** *Suppose the random variables  $X_1, \dots, X_d$  have continuous marginals and copula  $C_X$ . Let  $T_i : \mathbb{R} \rightarrow \mathbb{R}$ , for  $i = 1, \dots, d$  be strictly increasing functions. Then, the dependence structure of the random variables*

$$Y_1 := T_1(X_1), \dots, Y_d := T_d(X_d)$$

*is also given by the copula  $C_X$ .*

This leads immediately to the following result.

**Proposition 5.** *Let  $X_1, \dots, X_d$  be random variables with continuous marginals and suppose  $X_i = T_i(X_1)$  for  $i = 2, \dots, d$  where  $T_2, \dots, T_d$  are strictly increasing transformations. Then,  $X_1, \dots, X_d$  have the comonotonic copula.*

**Proof.** Apply the *invariance under monotonic transformations* proposition and observe that the copula of  $(X_1, X_1, \dots, X_1)$  is the comonotonic copula.  $\square$

Our optimization framework implicitly optimizes over the space of copulas by solving polynomial programs with possibly a large number of variables and constraints. (We saw in §4.3 that the number of variables and constraints is polynomial in  $|\mathbb{H}|$ ,

$|\mathbb{O}|$  and  $|\mathbb{X}|$  when calculating the probability of necessity (PN).) It may also be worthwhile, however, working explicitly with copulas. For example, the independence and comonotonic copulas are well understood and using these copulas to define SCMs may provide interesting benchmarks. Indeed, we estimate the PN for these benchmarks in our numerical results of §5. Towards this end, in §C.1 and §C.2, we explain how we can simulate our dynamic latent-state model to estimate the CQI under the independence (§C.1) and comonotonic (§C.2) copulas. Specifically, we assume each of the copulas for  $\mathbf{U}_t$  and  $\mathbf{V}_t$  are the independence copulas in §C.1, whereas in §C.2, we assume their copulas are the comonotonic copula.

There is no reason, however, why we couldn't combine them and assume, for example, that the copula for  $\mathbf{U}_t$  was the independence copula and the copula for  $\mathbf{V}_t$  was the comonotonic copula. More generally, we could use domain-specific knowledge to identify or narrow down sub-components of the copulas and leave the remaining components to be identified via the optimization problems. Since convex combinations of copulas are copulas, we could also optimize over such combinations. For example, suppose domain specific knowledge<sup>10</sup> tells us that the copula of  $\mathbf{V}_t$  is  $\lambda C^N + (1 - \lambda)C^I$ , i.e., a convex combination of the comonotonic and independence copulas, with  $\lambda \in [0, 1]$  unknown. Then, the optimization over  $\mathbf{V}_t$  would reduce to a single-variable ( $\lambda$ ) optimization with a linear constraint. Of course, the optimization over the copula of  $\mathbf{U}_t$  must also be included but domain-specific knowledge may also help to simplify and constrain that component of the optimization. Properties such as pathwise monotonicity (PM) and counterfactual stability (CS) can also be expressed in copula terms. Indeed, PM can be expressed via the comonotonic copula, as we discuss in §C.2.

### C.1. Counterfactual Simulations Under the Independence Copula

For convenience, we copy Figure 2 from the main text, which is now labelled as Figure 4. Furthermore, recall that  $(o_{1:T}, x_{1:T})$  is the observed data and  $\tilde{x}_{1:T}$  is the intervention policy that was applied.




Figure 4. The SCM underlying the dynamic latent-state model.

As in §4, we start with the posterior samples  $[h_{1:T}(b)]_{b=1}^B$  corresponding to the random path  $H_{1:T} \mid (o_{1:T}, x_{1:T})$ . These samples can be generated efficiently (cf. §B). For each sample  $b$ , our goal is to convert the sampled path  $h_{1:T}(b)$  into a counterfactual path  $\tilde{h}_{1:T}(b)$ . As noted in §4.2, irrespective of the copula choice, the counterfactual hidden state in period 1 equals the posterior sample, i.e.,

$$\tilde{h}_1(b) = h_1(b).$$

We next need to sample  $\tilde{h}_2(b)$ , but that first requires us to sample the counterfactual emission  $\tilde{o}_1(b)$  (cf. Figure 4). With the copula underlying  $\mathbf{V}_1$  being the independence copula, it follows that

$$\tilde{o}_1(b) = \begin{cases} o_1 & \text{if } x_1 = \tilde{x}_1 \text{ and } h_1(b) = \tilde{h}_1(b) \\ \text{sample from the emission distribution } [e_{\tilde{h}_1(b)\tilde{x}_1 i}]_i & \text{otherwise.} \end{cases}$$

The counterfactual emission  $\tilde{o}_1(b)$  allows us to sample the counterfactual state  $\tilde{h}_2(b)$ , which again leverages the fact that the copula underlying  $\mathbf{U}_2$  is the independence copula:

$$\tilde{h}_2(b) = \begin{cases} h_2(b) & \text{if } h_1(b) = \tilde{h}_1(b) \text{ and } o_1 = \tilde{o}_1(b) \\ \text{sample from the transition distribution } [q_{\tilde{h}_1(b)\tilde{o}_1(b)h'}]_{h'} & \text{otherwise.} \end{cases}$$

We then generate period 2 counterfactual emission  $\tilde{o}_2(b)$  in a similar manner and the process repeats until we hit the end of horizon. We summarize the procedure in Algorithm 2.

<sup>10</sup>It may be more likely that we only have domain specific knowledge over sub-components of the copulas (which are themselves copulas).

---

**Algorithm 2** Counterfactual simulations under the independence copula
 

---

**Require:**  $(\mathbf{E}, \mathbf{Q}), (o_{1:T}, x_{1:T}), [h_{1:T}(b)]_{b=1}^B, \tilde{x}_{1:T}$

- 1: **for**  $b = 1$  to  $B$  **do**
- 2:    $\tilde{h}_1(b) = h_1(b)$
- 3:   **for**  $t = 1$  to  $T - 1$  **do**
- 4:     **if**  $x_t = \tilde{x}_t$  and  $h_t(b) = \tilde{h}_t(b)$  **then**
- 5:        $\tilde{o}_t(b) = o_t$
- 6:     **else**
- 7:        $\tilde{o}_t(b) \sim \text{Categorical}([e_{\tilde{h}_t(b)\tilde{x}_t}^i]_i)$
- 8:     **end if**
- 9:     **if**  $h_t(b) = \tilde{h}_t(b)$  and  $o_t = \tilde{o}_t(b)$  **then**
- 10:        $\tilde{h}_{t+1}(b) = h_{t+1}(b)$
- 11:     **else**
- 12:        $\tilde{h}_{t+1}(b) \sim \text{Categorical}([q_{\tilde{h}_t(b)\tilde{o}_t(b)h'}]_{h'})$
- 13:     **end if**
- 14:   **end for**
- 15: **end for**
- 16: **return**  $[\tilde{h}_{1:T}(b)]_b$

---

## C.2. Counterfactual Simulations Under the Comonotonic Copula

Before the formal description (which involves non-trivial notation), we provide the intuition (which is relatively straightforward). We do so by revisiting Example 1, where we have the causal graph  $X \rightarrow Y$  with  $X \in \{0, 1\}$  (medical treatment) and  $Y \in \{\text{bad, better, best}\}$  (patient outcome). The outcome  $Y_x := Y \mid (X = x)$  obeys the following distribution:  $Y_0 \sim \{\text{bad, better, best}\}$  w.p.  $\{0.2, 0.3, 0.5\}$  and  $Y_1 \sim \{\text{bad, better, best}\}$  w.p.  $\{0.2, 0.2, 0.6\}$ . The underlying SCM is shown again in Figure 5.




Figure 5. SCM for Example 1 with the comonotonic copula and hence, the noise node is a scalar  $U \sim \text{Unif}[0, 1]$ , as opposed to a vector  $\mathbf{U}$ . The structural equation is  $Y = f(X, U)$ , which we denote by  $f_X(U)$ , the *inverse transform* function corresponding to the random variable  $Y_X$ . That is,  $f_0(u) = \text{bad, better, and best}$  if  $u \in [0, 0.2], u \in [0.2, 0.5]$ , and  $u \in [0.5, 1]$ , respectively. Similarly,  $f_1(u) = \text{bad, better, and best}$  if  $u \in [0, 0.2], u \in [0.2, 0.4]$ , and  $u \in [0.4, 1]$ , respectively.

Consider a patient whose outcome  $Y$  was “better” under no treatment ( $x = 0$ ). Given the prior  $U \sim \text{Unif}[0, 1]$ , we get the posterior  $U \mid (Y_0 = \text{better}) \sim \text{Unif}[0.2, 0.5]$ . Now, suppose we are interested in the understanding the counterfactual outcome under the intervention  $\tilde{x} = 1$ , i.e., the random variable  $\tilde{Y} := Y_1 \mid (Y_0 = \text{better})$ . Then, given the  $\text{Unif}[0.2, 0.5]$  belief over  $U$  and the functional form of  $f_1(\cdot)$  (as defined in the caption of Figure 5), we get that the  $[0.2, 0.4]$  region of  $U$  will map to “better” and the  $[0.4, 0.5]$  to “best”. Hence,  $\tilde{Y}$  equals “better” w.p.  $2/3$  and “best” w.p.  $1/3$ . This clearly obeys the pathwise monotonicity (PM) intuition we alluded to towards the end of Example 1 (“the counterfactual outcome  $\tilde{Y}$  should not be worse under treatment ( $\tilde{x} = 1$ ) than under no treatment ( $x = 0$ )”).

We now formalize this intuition to our dynamic latent-state model. As a prerequisite to discussing the notion of PM, one needs to define an ordering of the states (set  $\mathbb{H}$ ) and the emissions (set  $\mathbb{O}$ ), e.g., from “best” to “worst”. Denote by  $r_H(h)$  the rank of state  $h$  with respect to this ordering and by  $r_O(i)$  the rank of emission  $i$ . Furthermore, let  $r_H^{-1}(r)$  and  $r_O^{-1}(r)$  denote the inverse functions corresponding to  $r_H(h)$  and  $r_O(i)$ , respectively. That is,  $r_H^{-1}(r)$  returns the state with rank  $r$  and  $r_O^{-1}(r)$  returns the emission with rank  $r$ . Also, for each  $(h, i)$  pair, observe that  $[q_{hih'}]_{h'}$  denotes the transition distribution (which maps to the random variable  $H_{hi}$ ). Corresponding to this distribution, define the rank-ordered CDF as follows:

$$Q_{hih'} := \sum_{h'' : r_H(h'') \leq r_H(h')} q_{hih''} \quad \forall h'. \quad (21a)$$



Similarly, for each  $(h, x)$  pair, observe that  $[e_{hxi}]_i$  denotes the emission distribution (which maps to the random variable  $O_{hx}$ ). Corresponding to this distribution, define the rank-ordered CDF as follows:

$$E_{hxi} := \sum_{j:r_O(j) \leq r_O(i)} e_{hxj} \quad \forall i. \quad (21b)$$

Also, define  $Q_{hi0} = E_{hx0} = 0$  for all  $(h, i)$  and  $(h, x)$ . We discuss these orderings for the breast cancer application in §E.4.

As in §C.1, we start with the posterior samples  $[h_{1:T}(b)]_{b=1}^B$  corresponding to the random path  $H_{1:T} \mid (o_{1:T}, x_{1:T})$ . For each sample  $b$ , our goal is to convert the sampled path  $h_{1:T}(b)$  into a counterfactual path  $\tilde{h}_{1:T}(b)$ . As noted in §4.2, irrespective of the copula choice, the counterfactual hidden state in period 1 equals the posterior sample, i.e.,

$$\tilde{h}_1(b) = h_1(b).$$

To generate  $\tilde{o}_1(b)$ , we revisit the SCM in Figure 6, which now has the noise nodes as scalars (as opposed to vectors). This is a direct implication of the comonotonic copula - see the statement immediately below (19).




Figure 6. The SCM with the comonotonic copula. The key change is that the noise nodes are now scalars as opposed to vectors, i.e.,  $(U_t, V_t)$  as opposed to  $(\mathbf{U}_t, \mathbf{V}_t)$ .

By the structural equation (3a),  $\tilde{o}_1(b)$  equals

$$\tilde{o}_1(b) = f(\tilde{h}_1(b), \tilde{x}_1, V_1) = f_{\tilde{h}_1(b)\tilde{x}_1}(V_1), \quad (22a)$$

where  $f_{hx}(\cdot)$  is the inverse transform function corresponding to the rank-ordered CDF  $[E_{hxi}]_i$  (recall (21b)). Hence, all we need to sample  $\tilde{o}_1(b)$  is the posterior distribution of  $V_1$ , where the “posterior” corresponds to conditioning on  $O_{1h_1(b)x_1} = o_1$  (recall the notation  $O_{thx}$  from §4). Given the prior  $V_1 \sim \text{Unif}[0, 1]$ , we can compute the posterior in closed-form. In particular,

$$V_1 \mid (O_{1h_1(b)x_1} = o_1) \sim \text{Unif}[E_{h_1(b)x_1 o_1^-}, E_{h_1(b)x_1 o_1}], \quad (22b)$$

where  $o^- := r_O^{-1}(r_O(o) - 1)$  is the emission ranked just below  $o$ . Hence, we can efficiently sample  $V_1$  from its posterior, and this  $V_1$  sample can be used to generate  $\tilde{o}_1(b)$  (via (22a)). Given we encoded rank orderings in the CDF  $E_{hxi}$ , such sampling will naturally enforce pathwise monotonicity.

We can sample  $\tilde{h}_2(b)$  similarly. By the structural equation (3b),  $\tilde{h}_2(b)$  equals

$$\tilde{h}_2(b) = g(\tilde{h}_1(b), \tilde{o}_1, U_2) = g_{\tilde{h}_1(b)\tilde{o}_1}(U_2), \quad (23a)$$

where  $g_{hi}(\cdot)$  is the inverse transform function corresponding to the rank-ordered CDF  $[Q_{hih'}]_{h'}$  (recall (21a)). Hence, all we need to sample  $\tilde{h}_2(b)$  is the posterior distribution of  $U_2$ , where the “posterior” corresponds to conditioning on  $H_{2h_1(b)o_1} = h_2(b)$  (recall the notation  $H_{thi}$  from §4). Given the prior  $U_2 \sim \text{Unif}[0, 1]$ , we can compute the posterior in closed-form. In particular,

$$U_2 \mid (H_{2h_1(b)o_1} = h_2(b)) \sim \text{Unif}[Q_{h_1(b)o_1 h_2(b)^-}, Q_{h_1(b)o_1 h_2(b)}], \quad (23b)$$

where  $h^- := r_H^{-1}(r_H(h) - 1)$  is the state ranked just below  $h$ . Hence, we can efficiently sample  $U_2$  from its posterior, and this  $U_2$  sample can be used to generate  $\tilde{h}_2(b)$  (via (23a)). Given we encoded rank orderings in the CDF  $Q_{hih'}$ , such sampling will naturally enforce pathwise monotonicity.

We then generate period 2 counterfactual emission  $\tilde{o}_2(b)$  in a similar manner and the process repeats until we hit the end of horizon. We summarize the procedure in Algorithm 3.

---

**Algorithm 3** Counterfactual simulations under the comonotonic copula

---

**Require:**  $(\mathbf{E}, \mathbf{Q}), (o_{1:T}, x_{1:T}), [h_{1:T}(b)]_{b=1}^B, \tilde{x}_{1:T}, r_H(\cdot), r_O(\cdot)$

- 1: **for**  $b = 1$  to  $B$  **do**
- 2:    $\tilde{h}_1(b) = h_1(b)$
- 3:   **for**  $t = 1$  to  $T - 1$  **do**
- 4:      $v_t \sim \text{Unif}[E_{h_t(b)x_t o_t^-}, E_{h_t(b)x_t o_t}]$  % posterior sample of  $V_t$  (see (22b))
- 5:      $\tilde{o}_t(b) = f_{\tilde{h}_t(b)\tilde{x}_t}(v_t)$  % counterfactual emission (see (22a))
- 6:      $u_{t+1} \sim \text{Unif}[Q_{h_t(b)o_t h_{t+1}(b)^-}, Q_{h_t(b)o_t h_{t+1}(b)}]$  % posterior sample of  $U_{t+1}$  (see (23b))
- 7:      $\tilde{h}_{t+1}(b) = g_{\tilde{h}_t(b)\tilde{o}_t}(u_{t+1})$  % counterfactual state (see (23a))
- 8:   **end for**
- 9: **end for**
- 10: **return**  $[\tilde{h}_{1:T}(b)]_b$

---

## D. Enhancing the Scalability of the Polynomial Optimization

In this section, we discuss ways to enhance the scalability of the polynomial optimizations in (8). First, in §D.1, we show how the optimization can be reformulated to avoid the exponential dependence on  $T$  (recall the discussion towards the end of §4.3). Second, in §D.2, we discuss an approximate way to optimize our problem that drastically reduces the underlying dimensionality of the problem. Third, in §D.3, we combine our ideas from §D.1 and §D.2 and demonstrate (via numerics) that we can obtain high-quality solutions for  $T$  as large as 100 in just a few hours of compute time.

Related to scalability, we mention in passing that in each of our optimization problems, we added the constraint that the objective value (which is a probability) must lie in  $[0, 1]$ . Of course, this constraint is redundant but we found it helped speed up the solver convergence in a few instances, possibly because it shrunk the search space as the solver does not know a priori that the objective is a probability.

### D.1. Reformulating the Polynomial Optimization to Avoid the Exponential Dependence on $T$

Recall Lemmas 1 and 2, which characterize the objective function of our polynomial optimization problem. We repeat them here for the sake of convenience.

**Lemma 1.** *We have*

$$PN = 1 - \lim_{B \rightarrow \infty} \frac{1}{B} \sum_{b=1}^B \mathbb{P}_{\tilde{\mathbf{M}}(b)}(\tilde{H}_T = 7).$$

**Lemma 2.** *For  $t \in \{T, T - 1, \dots, 2\}$ ,  $\mathbb{P}_{\tilde{\mathbf{M}}(b)}(\tilde{h}_t) := \mathbb{P}_{\tilde{\mathbf{M}}(b)}(\tilde{H}_t = \tilde{h}_t)$  obeys the following recursion (over  $t$ ):*

$$\begin{aligned} \mathbb{P}_{\tilde{\mathbf{M}}(b)}(\tilde{h}_t) &= \sum_{\tilde{h}_{t-1}, \tilde{o}_{t-1}} \frac{\pi_{\tilde{h}_{t-1}\tilde{o}_{t-1}, h_{t-1}(b)o_{t-1}}(\tilde{h}_t, h_t(b))}{q_{h_{t-1}(b)o_{t-1}h_t(b)}} \times \\ &\quad \frac{\theta_{\tilde{h}_{t-1}\tilde{x}_{t-1}, h_{t-1}(b)x_{t-1}}(\tilde{o}_{t-1}, o_{t-1})}{e_{h_{t-1}(b)x_{t-1}o_{t-1}}} \times \\ &\quad \mathbb{P}_{\tilde{\mathbf{M}}(b)}(\tilde{h}_{t-1}). \end{aligned}$$

*The recursion breaks at  $t = 1$ :*

$$\mathbb{P}_{\tilde{\mathbf{M}}(b)}(\tilde{h}_1) = \begin{cases} 1 & \text{if } \tilde{h}_1 = h_1(b) \\ 0 & \text{otherwise.} \end{cases}$$

It is easy to see that a naive expansion of PN (as per Lemmas 1 and 2) results in a number of terms that is exponential in  $T$ . This is clearly undesirable since we end up running into memory issues for even a moderate value of  $T$ . For example,

such issues arise for  $T > 10$  in the breast cancer numerics of §5. It is possible to remove this exponential dependence, however, by a reformulation of the optimization, which we now discuss. (Note that the objective function remains the same irrespective of whether we optimize over the pairwise marginals (as discussed in §D.2) or the joint distribution (as presented in §4.3) and hence, the reformulation here is “universal”.)

The reformulation steps are as follows:

1. Define  $\mathbb{P}_{\widetilde{\mathbf{M}}(b)}(\widetilde{h}_t)$  from Lemma 2 as a decision variable for all  $(t, \widetilde{h}_t, b) \in [T] \times \mathbb{H} \times [B]$ .
2. Add the Lemma 2 equations as constraints in the optimization (for each  $(t, \widetilde{h}_t, b) \in [T] \times \mathbb{H} \times [B]$ ). Note that these are non-linear but polynomial constraints and hence, we remain within the class of polynomial programs. Furthermore, none of the constraints have an exponential number of terms since  $\mathbb{P}_{\widetilde{\mathbf{M}}(b)}(\widetilde{h}_t)$  are decision variables now.
3. The objective now is simply the expression in Lemma 1.

These steps result in the following<sup>11</sup> optimization, where we use the decision variable  $\gamma_{\widetilde{h}_t b}^t$  to denote the probability term<sup>12</sup>  $\mathbb{P}_{\widetilde{\mathbf{M}}(b)}(\widetilde{h}_t)$  in the LHS of Lemma 2 for all  $(t, \widetilde{h}_t, b) \in [T] \times \mathbb{H} \times [B]$ , with  $\boldsymbol{\gamma} := [\gamma_{\widetilde{h}_t b}^t]_{(t, \widetilde{h}_t, b)}$ :

$$\max_{(\boldsymbol{\theta}, \boldsymbol{\pi}) \in \mathcal{F}, \boldsymbol{\gamma}} \left\{ 1 - \frac{1}{B} \sum_{b=1}^B \gamma_{7B}^T \right\} \quad (24a)$$

$$\text{s.t. } \gamma_{hb}^t = \sum_{h' \in \mathbb{H}} \sum_{o' \in \mathbb{O}} \frac{\pi_{h'o', h_{t-1}(b) o_{t-1}}(h, h_t(b))}{q_{h_{t-1}(b) o_{t-1} h_t(b)}} \times \frac{\theta_{h' \widetilde{x}_{t-1}, h_{t-1}(b) x_{t-1}}(o', o_{t-1})}{e_{h_{t-1}(b) x_{t-1} o_{t-1}}} \times \gamma_{h'b}^{t-1} \quad \forall t > 1 \forall h \in \mathbb{H} \forall b \in [B] \quad (24b)$$

$$\gamma_{hb}^1 = 1 \quad \forall h = h_1(b) \forall b \in [B] \quad (24c)$$

$$\gamma_{hb}^1 = 0 \quad \forall h \neq h_1(b) \forall b \in [B]. \quad (24d)$$

As before (refer to §4), the feasibility set  $\mathcal{F}$  over  $(\boldsymbol{\theta}, \boldsymbol{\pi})$  can correspond to (10), (11), and (12). It can also include additional constraints such as CS and PM, or correspond to the lower-dimensional space over the pairwise marginals (as discussed in §D.2). Clearly, (24) has a linear objective and polynomial constraints, and is therefore also a polynomial program. The number of terms in the objective is no longer exponential in  $T$  but this has come at the cost of having to add a total of  $|\mathbb{H}|TB$  decision variables and (polynomial) constraints to the original formulation in (8). Though the size of our reformulation (number of variables and constraints) scales with both  $T$  and  $B$ , we found it to scale much more gracefully (with respect to  $T$ ) than the original formulation, as we discuss in §D.3 below.

Note that we do not necessarily need to add these  $|\mathbb{H}|TB$  variables and constraints to the optimization but for that, we need the ability to modify the source code of the optimization solver (BARON in our case). This is because even in the original formulation (8), we can actually evaluate the objective function in polynomial time and space rather than naively expanding it into exponentially many terms. To see this, consider a given sample number  $b \in [B]$ . We need to evaluate  $\mathbb{P}_{\widetilde{\mathbf{M}}(b)}(\widetilde{H}_T = 7)$  from Lemma 2. To do so, we start from period 1 and store  $\mathbb{P}_{\widetilde{\mathbf{M}}(b)}(\widetilde{h}_1)$  for all  $\widetilde{h}_1 \in \mathbb{H}$  (see Lemma 2’s base case). We then move to period 2 and store  $\mathbb{P}_{\widetilde{\mathbf{M}}(b)}(\widetilde{h}_2)$  for all  $\widetilde{h}_2 \in \mathbb{H}$  (see Lemma 2’s recursion). The key here is that when computing  $\mathbb{P}_{\widetilde{\mathbf{M}}(b)}(\widetilde{h}_2)$ , we make use of the stored values of  $\mathbb{P}_{\widetilde{\mathbf{M}}(b)}(\widetilde{h}_1)$ . We then move to period 3 evaluations, where we make use of the stored values of  $\mathbb{P}_{\widetilde{\mathbf{M}}(b)}(\widetilde{h}_2)$ . We repeat this procedure until we hit period  $T$ . Clearly, this procedure requires polynomial time and space. Furthermore, we can evaluate the gradient (and the Hessian) of  $\mathbb{P}_{\widetilde{\mathbf{M}}(b)}(\widetilde{H}_T = 7)$  in a similar manner (if needed by the optimization solver). We can therefore evaluate the objective and its gradient information at a given point in

<sup>11</sup>Note that we focus on the maximization problem from (8) but the same holds for the minimization counterpart. All we need to do is simply change the “max” to a “min” in the objective function (24a).

<sup>12</sup>To be pedantic, we could have added a “ $t$ ” super-script in  $\mathbb{P}_{\widetilde{\mathbf{M}}(b)}(\widetilde{h}_t)$  and used the notation  $\mathbb{P}_{\widetilde{\mathbf{M}}(b)}^t(\widetilde{h}_t)$  instead. However, we did not do so earlier since this dependence on  $t$  was implicitly understood to exist, and adding this extra super-script felt unnecessary.

polynomial time and space. These can then be used by the optimization solver. However, we are unable to modify the solver we use (BARON), and BARON by default does not exploit this structure but naively expands the objective into  $\exp(T)$  terms. As such, we use the reformulation presented in (24) instead.

## D.2. Approximating the Joint Optimization by the Pairwise Optimization

The problem (8) discussed in §4.3 optimizes over the joint PMFs (“joint optimization”). The challenge here lies in the dimensionality of the underlying joint distribution. As discussed towards the end of §4.3, the problem size (number of decision variables in particular) can grow exponentially in the primitives (e.g.,  $|\mathbb{H}|$ ,  $|\mathbb{X}|$ , and  $|\mathbb{O}|$ ). This is because the decision variables capture the entire joint distribution. Though we might be able to exploit application-specific sparsity to manage this blow-up (as we in fact do for the breast cancer application), it is worth exploring if there is a more tractable alternative in general (i.e., not specific to any application). We now show that this is possible.

Recall from §4.3 that we are interested in the following optimizations (repeating (8) for convenience):

$$\text{PN}^{\text{ub}}(B) := \max_{(\boldsymbol{\theta}, \boldsymbol{\pi}) \in \mathcal{F}} \text{PN}(\boldsymbol{\theta}, \boldsymbol{\pi} \mid [h_{1:T}(b)]_b) \quad (8a)$$

$$\text{PN}^{\text{lb}}(B) := \min_{(\boldsymbol{\theta}, \boldsymbol{\pi}) \in \mathcal{F}} \text{PN}(\boldsymbol{\theta}, \boldsymbol{\pi} \mid [h_{1:T}(b)]_b). \quad (8b)$$

The key observation here is that the objective function does not depend on the joint PMF of  $(\boldsymbol{\theta}, \boldsymbol{\pi})$  but only the corresponding pairwise marginals (recall Lemmas 1 and 2). We introduced the joint PMF decision variables to ensure the feasibility set  $\mathcal{F}$  is such that the pairwise marginals are valid. However, as an alternative, we can choose to not introduce the joint variables in the optimization and instead approximate  $\mathcal{F}$  by expressing it in terms of the pairwise variables. For example, since the pairwise variables correspond to the 2-dimensional PMFs, they must obey basic probability axioms. In particular, they must be non-negative and agree with their known 1-dimensional marginals so that

$$\sum_{h'} \pi_{\tilde{h}\tilde{i}, h\tilde{i}}(\tilde{h}', h') = q_{\tilde{h}\tilde{i}} \forall (\tilde{h}, \tilde{i}, \tilde{h}') \forall (h, i) \quad (25a)$$

$$\sum_{\tilde{h}'} \pi_{\tilde{h}\tilde{i}, h\tilde{i}}(\tilde{h}', h') = q_{h\tilde{i}} \forall (\tilde{h}, \tilde{i}) \forall (h, i, h') \quad (25b)$$

$$\sum_i \theta_{\tilde{h}\tilde{x}, h\tilde{x}}(\tilde{i}, i) = e_{\tilde{h}\tilde{x}} \forall (\tilde{h}, \tilde{x}, \tilde{i}) \forall (h, x) \quad (25c)$$

$$\sum_{\tilde{i}} \theta_{\tilde{h}\tilde{x}, h\tilde{x}}(\tilde{i}, i) = e_{h\tilde{x}} \forall (\tilde{h}, \tilde{x}) \forall (h, x, i). \quad (25d)$$

These constraints are analogous to (10) in §4.3. It is easy to see that if (10) is obeyed, then so is (25). However, the reverse implication does not hold, meaning the feasibility space defined by (25) and non-negativity (say  $\mathcal{F}'$ ) is a super-set of the feasibility space  $\mathcal{F}$  in §4.3. In other words, though the constraints in  $\mathcal{F}'$  are necessary, they are not sufficient to ensure the pairwise marginals correspond to a valid joint distribution. Hence, optimizing over  $\mathcal{F}'$  (“pairwise optimization”)<sup>13</sup> is a relaxation to the problem of optimizing over  $\mathcal{F}$ . In fact, as we show via a simple example next, this relaxation can be strict. (We thank an anonymous reviewer for this example.)

**Example 2.** Consider three random variables  $X$ ,  $Y$ , and  $Z$  with the following pairwise marginals:

$$(X, Y) = \begin{cases} (0, 0) \text{ w.p. } 1/2 \\ (1, 1) \text{ w.p. } 1/2 \end{cases} \quad (Y, Z) = \begin{cases} (0, 0) \text{ w.p. } 1/2 \\ (1, 1) \text{ w.p. } 1/2 \end{cases} \quad (X, Z) = \begin{cases} (1, 0) \text{ w.p. } 1/2 \\ (0, 1) \text{ w.p. } 1/2. \end{cases}$$

It is easy to verify these pairwise marginals obey (25) along with non-negativity. However, they do not correspond to any valid joint distribution over  $(X, Y, Z)$ . To see this, suppose  $(X, Y)$  realizes a value of  $(0, 0)$ . Then, the pairwise marginal of  $(Y, Z)$  implies  $(Y, Z)$  has to be  $(0, 0)$ , which implies  $(X, Z)$  must be  $(1, 0)$ , resulting in a contradiction. Therefore, the bivariate marginals are not consistent with any valid 3-dimensional joint distribution.

Despite the relaxation being strict<sup>14</sup>, we found it to produce high-quality solutions and be highly scalable (discussed in §D.3).

<sup>13</sup>Note that the pairwise optimization is identical to the joint optimization (8) but with the following two differences: (a)  $\mathcal{F}$  replaced by  $\mathcal{F}'$  and (b) joint decision variables not defined.

<sup>14</sup>Since the pairwise optimization is a relaxation of the joint optimization, it follows that Proposition 2 still holds for the bounds produced by the pairwise optimization.

The high scalability is primarily driven by the lower dimensionality of the decision variables. In particular, the pairwise optimization has *at most*  $|\mathbb{H}|^4|\mathbb{O}|^2 + |\mathbb{H}|^2|\mathbb{O}|^2|\mathbb{X}|^2$  decision variables (recall that the joint optimization has an additional  $|\mathbb{O}|^{|\mathbb{H}||\mathbb{X}|} + |\mathbb{H}|^{|\mathbb{H}||\mathbb{O}|}$  decision variables). In fact, after exploiting the sparsity in the breast cancer application (along with the variable elimination discussed in Footnote 5), the pairwise optimization has only 1082 decision variables. This is in contrast to the joint optimization which has 16,124 decision variables. Though the pairwise optimization has more constraints than the joint optimization, the difference is not that stark (2085 vs. 610). (The numbers reported here correspond to the objective formulation presented in §4.3 as opposed to the reformulation in §D.1. The reformulation adds a total of  $|\mathbb{H}|TB$  decision variables and  $|\mathbb{H}|TB$  constraints to both the pairwise and the joint optimizations.)

### D.3. Computational Performance

We now compute upper and lower bounds on PN by (a) using the reformulation discussed in §D.1 and (b) optimizing over the relaxed constraint set defined by the pairwise marginals as discussed in §D.2.<sup>15</sup> We focus on path 1 from §5 for brevity and note that the results for path 2 are similar. The implementation details remain the same as in §5 (i.e., we code in MATLAB-BARON with CPLEX as the LP / MIP solver, set absolute termination tolerance at 0.01, generate  $B = 100$  samples for SAA, and average over 20 seeds). To test the scalability of our approach, we now experiment with  $T \in \{5, 10, 15, 20, 25, 50, 75, 100\}$ . Note that  $T = 100$  is an order of magnitude larger than the longest horizon we have in §5, i.e.,  $T = 10$ . We next discuss the results that are shown in Figure 7, with Figure 7(a) showcasing scalability and Figure 7(b) quality.




Figure 7. Evaluating the computational performance of our ideas in §D.1 and §D.2 on path 1 from §5. All results are averaged over the 20 seeds we use. In sub-figure (b), the UB and LB curves for the “joint” optimization (black color) are the same as the ones in Figure 3(a), and only go as far as  $T = 10$  (since we run into memory issues for  $T > 10$ ). Furthermore, to avoid clutter, we do not show the standard deviation bars in sub-figure (b) and note that the maximum standard deviation value is less than 0.01. To be clear, the compute times in sub-figure (a) correspond to the blue curves in sub-figure (b) (“pairwise”), which are the focus of this section.

In Figure 7(a), we display the compute time as a function of  $T$ . Compute time refers to the total time taken to compute LB and UB. Note that we let the solver run until convergence to global optimality (on just one core with at most 16 GB RAM). We are able to solve for  $T = 100$  in 3 hours on average (over 20 seeds), with the minimum time being 1.2 hours and the maximum time being 8.4 hours. This demonstrates the scalability of our approach. (It is worth mentioning that after eliminating redundant variables and constraints, exploiting sparsity, using the reformulation of §D.1, and the pairwise approximation of §D.2, the  $T = 100$  and  $B = 100$  optimization has 71,082 decision variables, 2085 linear constraints, and 70,000 polynomial constraints.)

<sup>15</sup>We also experimented with other variations of these two approaches. If we use neither of them (as is the case in §5), then we run into memory issues for  $T > 10$ . In fact, even if we only use the second approach (optimizing over the relaxed constraint set), then we run into memory issues for  $T > 10$  since the objective still scales exponentially in  $T$ . Finally, if we only use the first approach, i.e. the reformulation of §D.1, and optimize over the joint, the BARON solver does not converge even for  $T$  as small as 5 in 24 hours of compute time. This is because keeping the joint variables while doing the reformulation results in an optimization with a very large number of variables and constraints (even after we exploit sparsity).

In Figure 7(b), we display the PN values as a function of  $T$ . The “joint” UB and LB curves are the same as the ones in Figure 3(a), and only go as far as  $T = 10$  because of the aforementioned memory issues for  $T > 10$ . The “pairwise” UB and LB curves are the focus of this section and our goal here is to evaluate the quality of the “pairwise” bounds (blue curves) and we do so in two ways. First, as we are able to solve the joint optimization for  $T \leq 10$ , we can use the “joint” bounds as benchmarks for the “pairwise” bounds. As may be seen from the figure, the joint and pairwise bounds are very close to each other (for values of  $T \leq 10$ ). In particular, the joint and pairwise lower bounds coincide and equal 0.8725 and 0.8884 for  $T = 5$  and  $T = 10$ , respectively. The pairwise and joint upper bounds also coincide and equal 0.9885 for  $T = 5$ . The only difference between the two is the upper bound for  $T = 10$  with values of 0.9904 and 0.9899, respectively. We therefore conclude that the pairwise bounds provide a very good approximation to the joint bounds, at least when  $T \leq 10$ .

Second, for  $T > 10$ , we use the fact that we can simulate the independence and comonotonic copulas, which by definition are feasible solutions to the joint optimization. We therefore know that maximizing (minimizing) over the joint distribution will yield an upper (lower) bound that is no lower (higher) than the independence (comonotonic) curves in the figure. As an example, the gap between the independence and the pairwise UB curves (for  $T > 10$ ) is never greater than 0.01, which means we lose at most 0.01 by restricting ourselves to the pairwise marginals. Similarly, the maximum gap between the comonotonic and the pairwise LB curves is  $\sim 0.02$ . Thus, the pairwise bounds provide a high quality approximation to the joint bounds even when  $T > 10$ .

We can also embed CS and PM constraints in the pairwise optimization (recall from §4.3 that both these constraints are over the pairwise variables) and we show the corresponding bounds in Figure 8. Naturally, the bounds we obtain are tighter than the pairwise bounds in Figure 7(b). In particular, the UB gets much tighter while the LB does not change much.

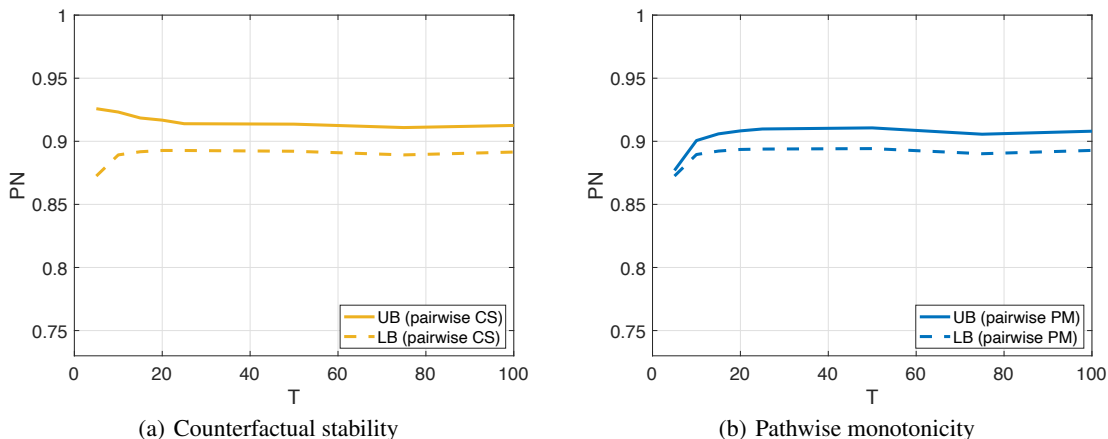


Figure 8. PN bounds obtained when we embed CS and PM constraints in the pairwise optimization. All results are computed as the average of bounds obtained from 20 different seeds with each seed being used to generate  $B = 100$  paths. Furthermore, to avoid clutter, we do not show the standard deviation bars and note that the maximum standard deviation value is less than 0.01.

## E. Further Details on the Breast Cancer Case Study

We discuss the breast cancer model primitives and their calibration in §E.1, followed by showing how we exploit sparsity to reduce the number of decision variables (§E.2). We then provide details on the PM constraints and the comonotonic copula in §E.3 and §E.4, respectively. Finally, in §E.5, we show the results for path 2.

### E.1. Model Primitives and Their Calibration

As discussed in §2, the breast cancer application has  $|\mathbb{H}| = 7$  states,  $|\mathbb{O}| = 7$  emissions, and  $|\mathbb{X}| = 2$  actions. To be consistent with the literature (Ayer et al., 2012), we treat each period as corresponding to 6 months. The model comprises of three primitives:  $\mathbf{p}$ ,  $\mathbf{Q}$ , and  $\mathbf{E}$ . We discuss their (sparse) structure and the calibration to real-data in §E.1.1, §E.1.2, and §E.1.3, respectively.

### E.1.1. INITIAL STATE DISTRIBUTION $\mathbf{p}$

We have  $\mathbf{p} := (p_1, \dots, p_7)$  where  $p_h := \mathbb{P}(H_1 = h)$  for all  $h$ . Usually, breast cancer screening starts around the age of 40 and the prevalence among females aged 40-49 is 1.0183% (Table 4.24 of NIH (2020), all races, females):

$$p_2 + p_3 = 0.010183.$$

Since in-situ cancer comprises 20% of new breast cancer diagnoses (Sprague & Trentham-Dietz, 2009), we get

$$\begin{aligned} p_2 &= 0.2 \times 0.010183 \\ p_3 &= 0.8 \times 0.010183. \end{aligned}$$

It is natural to set

$$p_4 = p_5 = p_6 = p_7 = 0$$

and hence,

$$p_1 = 1 - p_2 - p_3 = 1 - 0.010183.$$

### E.1.2. TRANSITION DISTRIBUTION $\mathbf{Q}$

We have  $\mathbf{Q} := [q_{hih'}]_{h,i,h'}$  with  $q_{hih'} := \mathbb{P}(H_{t+1} = h' \mid H_t = h, O_t = i)$ . Before discussing the calibration, we discuss the sparse structure of  $\mathbf{Q}$ . To do so, we define the transition matrix  $\mathbf{Q}(i) := [q_{hih'}]_{hh'}$  for each emission  $i$  (so that each row sums to 1) and observe that we have the following structure:

$$\begin{aligned} \mathbf{Q}(1) &= \begin{bmatrix} q_{11} & q_{12} & q_{13} & & & & \\ & q_{22} & q_{23} & & & & q_{27} \\ & & q_{33} & & & & q_{37} \\ & & & q_{44} & q_{45} & q_{46} & q_{47} \\ & & & & q_{55} & q_{56} & q_{57} \\ & & & & & 1 & \\ & & & & & & 1 \end{bmatrix} & \mathbf{Q}(2) &= \begin{bmatrix} q_{11} & q_{12} & q_{13} & & & & \\ & q_{22} & q_{23} & & & & q_{27} \\ & & q_{33} & & & & q_{37} \\ & & & & & & \\ & & & & & & \\ & & & & & & \\ & & & & & & \end{bmatrix} \\ \\ \mathbf{Q}(3) &= \begin{bmatrix} q_{11} & q_{12} & q_{13} \\ & & \\ & & \\ & & \\ & & \\ & & \\ & & \end{bmatrix} & \mathbf{Q}(4) &= \begin{bmatrix} & q_{24} & q_{25} & q_{26} & \bar{q}_{27} \\ & q_{44} & q_{45} & q_{46} & q_{47} \\ & & & & \\ & & & & \\ & & & & \\ & & & & \\ & & & & \end{bmatrix} \\ \\ \mathbf{Q}(5) &= \begin{bmatrix} & & & & & & \\ & & & & & & \\ & q_{35} & q_{36} & \bar{q}_{37} & & & \\ & q_{55} & q_{56} & q_{57} & & & \\ & & & & & & \\ & & & & & & \\ & & & & & & \end{bmatrix} & \mathbf{Q}(6) &= \begin{bmatrix} & & & & & & \\ & & & & & & \\ & & & & & & \\ & & & & & & \\ & & & & & & \\ & & & & & & \\ & & & & & & \end{bmatrix} & \mathbf{Q}(7) &= \begin{bmatrix} & & & & & & \\ & & & & & & \\ & & & & & & \\ & & & & & & \\ & & & & & & \\ & & & & & & \\ & & & & & & 1 \end{bmatrix}. \end{aligned}$$

A few comments are in order. First, an empty row means it is an impossible  $(h, i)$  combination. For example, if the we observe an emission  $i = 3$  (i.e., a negative biopsy), then the underlying patient state has to be healthy, i.e.,  $h \notin \{2, 3, 4, 5, 6, 7\}$ . Thus, rows 2 to 7 are empty in  $\mathbf{Q}(3)$ .

Second, observe that there is a decent amount of overlap across  $[\mathbf{Q}(i)]_i$  in terms of the underlying parameters. For example,  $q_{11}$  corresponds to the probability a healthy patient stays healthy, which is independent of the emission being 1 (no test), 2 (negative test), or 3 (positive test but negative biopsy). Hence,  $q_{11}$  appears in all three matrices  $\mathbf{Q}(1)$ ,  $\mathbf{Q}(2)$ , and  $\mathbf{Q}(3)$ . Of course, if the emission is 4, 5, 6, or 7, then the patient can not be healthy and hence, the corresponding entry in matrices  $\mathbf{Q}(4)$ ,

$\mathbf{Q}(5)$ ,  $\mathbf{Q}(6)$ , and  $\mathbf{Q}(7)$  is absent (in fact, the entire first row is empty, which means it is an impossible  $(h, i)$  combination as discussed above).

Third, some rows have only a partial set of entries, which means that the other entries equal 0. For example, if a patient is healthy (state 1), then her state can not transition to 4 (diagnosed in-situ with treatment started), 5 (diagnosed invasive with treatment started), 6 (recovered), or 7 (death) and hence,  $q_{14} = q_{15} = q_{16} = q_{17} = 0$ . Hence, we do not show  $q_{14}, q_{15}, q_{16}, q_{17}$  in  $\mathbf{Q}(1)$ ,  $\mathbf{Q}(2)$ , or  $\mathbf{Q}(3)$ .

Fourth, observe that we have a “bar” over  $\bar{q}_{27}$  (in  $\mathbf{Q}(4)$ ) and  $\bar{q}_{37}$  (in  $\mathbf{Q}(5)$ ). This is done to recognize them being different from  $q_{27}$  (in  $\mathbf{Q}(1)$  and  $\mathbf{Q}(2)$ ) and  $q_{37}$  (in  $\mathbf{Q}(1)$  and  $\mathbf{Q}(2)$ ). To see the difference, consider  $q_{37}$  versus  $\bar{q}_{37}$ .  $q_{37}$  corresponds to the patient state transitioning from invasive cancer to death when the cancer was not detected (and hence, no treatment). On the other hand,  $\bar{q}_{37}$  corresponds to the patient state transitioning from invasive cancer to death when the cancer was detected (and hence, treatment was provided). Naturally, we expect  $\bar{q}_{37} \leq q_{37}$ .

Finally, since states 6 (recovery) and 7 (death) are absorbing, we have  $q_{66} = q_{77} = 1$ .

Having discussed the structure of  $\mathbf{Q}$ , we now calibrate it to real-data. We iterate over each state  $h \in \{1, \dots, 7\}$  in a sequential manner.

**State 1 (healthy).** For state 1, we are interested in  $(q_{11}, q_{12}, q_{13})$ . These probabilities can depend on a woman’s age but we ignore that and work with averages. Let’s focus on  $(q_{12}, q_{13})$  since

$$q_{11} = 1 - q_{12} - q_{13}.$$

For  $q_{12}$ , we use the in-situ incidence rates from Table 4.12 of [NIH \(2020\)](#) (all races, females). For  $q_{13}$ , we use the invasive incidence rates from Table 4.11 of [NIH \(2020\)](#) (all races, females). The reported numbers are per year and we should divide by 2 to convert to a 6-month scale:

$$q_{12} = \frac{1}{2} \times \frac{33.0}{100000}$$

$$q_{13} = \frac{1}{2} \times \frac{128.5}{100000}.$$

Note that consistent with the 20-80 split in  $(p_2, p_3)$ , we have  $q_{13} \approx 4q_{12}$ .

**State 2 (undiagnosed in-situ cancer).** We are interested in  $(q_{22}, q_{23}, q_{27})$  (if cancer is not detected) and  $(q_{24}, q_{25}, q_{26}, \bar{q}_{27})$  (if cancer is detected). First, consider  $(q_{22}, q_{23}, q_{27})$ . Table 4.13 of [NIH \(2020\)](#) and Page 26 of [UWBCS \(2013\)](#) imply there is no death from in-situ cancer:

$$q_{27} = 0.$$

[Haugh & Lacedelli \(2019\)](#) assumed  $q_{23}$  to equal the invasive incidence rate  $q_{13}$  and so do we:

$$q_{23} = q_{13}$$

$$q_{22} = 1 - q_{23} - q_{27} = 1 - q_{13}.$$

Second, consider  $(q_{24}, q_{25}, q_{26}, \bar{q}_{27})$ . As  $q_{27} = 0$  and we expect  $\bar{q}_{27} \leq q_{27}$  (recall comment #4 above), we set

$$\bar{q}_{27} = 0.$$

As all in-situ cancer patients survive (if treated), no one transitions to invasive (if in-situ detected):

$$q_{25} = 0.$$

Finally, we have

$$q_{24} + q_{26} = 1.$$

The split between  $q_{24}$  and  $q_{26}$  is irrelevant in terms of the patient dying or not (all will survive as there is no positive probability path from state 4 to death; this will become clear when we discuss state 4 below).



**State 3 (undiagnosed invasive cancer).** We are interested in  $(q_{33}, q_{37})$  (if cancer is not detected) and  $(q_{35}, q_{36}, \bar{q}_{37})$  (if cancer is detected). First, consider  $(q_{33}, q_{37})$ .  $q_{37}$  is the probability of dying from invasive breast cancer (under no treatment). According to [Johnstone et al. \(2000\)](#), the 5-year and 10-year survival rates for invasive breast cancer patients (under no treatment) are 18.4% and 3.6%, respectively. On calibrating to 5-year rate, we get  $(1 - q_{37})^{10} = 0.184$ , which implies  $q_{37} \approx 15.6\%$  (note that we use “10” in the exponent since our time periods correspond to 6 months and hence, 5 years correspond to 10 periods). Similarly, on calibrating to 10-year rate, we get  $(1 - q_{37})^{20} = 0.036$  implies  $q_{37} \approx 15.3\%$ . The two calibrations are consistent with each other (lending evidence to time-invariance). Minimizing sum of squared errors over the two data points, i.e.,  $\min_{q_{37} \in [0,1]} \{((1 - q_{37})^{10} - 0.184)^2 + ((1 - q_{37})^{20} - 0.036)^2\}$ , gives the following estimate:

$$q_{37} = 0.1554.$$

Naturally, we have

$$q_{33} = 1 - q_{37}.$$

Second, consider  $(q_{35}, q_{36}, \bar{q}_{37})$ .  $q_{36}$  and  $\bar{q}_{37}$  are the probabilities of recovering and dying from invasive breast cancer (under treatment). Table 4.14 of [NIH \(2020\)](#) has various survival rates we can use to calibrate. We calibrate using the 10 data points corresponding to the year 2007 (see Figure 9):

$$\begin{aligned} q_{36} &= 0.0459 \\ \bar{q}_{37} &= 0.0113. \end{aligned}$$

As a sanity check, note that  $\bar{q}_{37} < q_{37}$ . Finally,

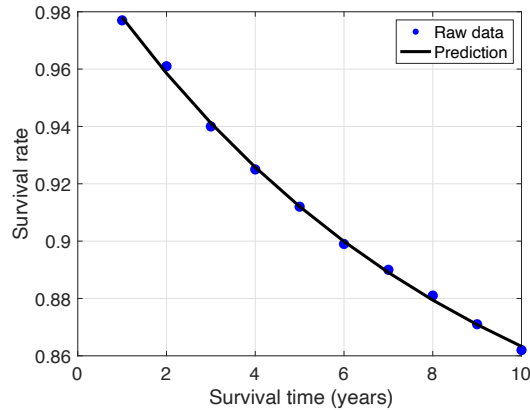


Figure 9. Calibration of  $(q_{36}, \bar{q}_{37})$ . Under our (time-invariant) Markov model, with a starting state of invasive breast cancer (under treatment), the survival rate after  $x$  years equals  $1 - \bar{q}_{37} \sum_{i=0}^{2x-1} (1 - q_{36} - \bar{q}_{37})^i$ . Minimizing the sum of squared errors (on the 10 blue data points in the plot) over  $(q_{36}, \bar{q}_{37})$  gives us an estimate of  $(0.0459, 0.0113)$ . The prediction using our fit is shown via the black curve.

$$q_{35} = 1 - q_{36} - \bar{q}_{37}.$$

**State 4 (diagnosed in-situ cancer).** We are interested in  $(q_{44}, q_{45}, q_{46}, q_{47})$ . Under our Markov model (which by definition is “memoryless”), it seems reasonable to set

$$(q_{44}, q_{45}, q_{46}, q_{47}) = (q_{24}, q_{25}, q_{26}, \bar{q}_{27}).$$

**State 5 (diagnosed invasive cancer).** We are interested in  $(q_{55}, q_{56}, q_{57})$ . Under our Markov model, it seems reasonable to set

$$(q_{55}, q_{56}, q_{57}) = (q_{35}, q_{36}, \bar{q}_{37}).$$

**States 6 (recovery) and 7 (death).** These two states are absorbing and hence,

$$q_{66} = q_{77} = 1.$$

### E.1.3. EMISSION DISTRIBUTION $\mathbf{E}$

We have  $\mathbf{E} := [e_{hxi}]_{h,x,i}$  with  $e_{hxi} := \mathbb{P}(O_t = i \mid H_t = h, X_t = x)$ . Before discussing the calibration, we discuss the sparse structure of  $\mathbf{E}$ . To do so, we define the matrix  $\mathbf{E}(x) := [e_{hxi}]_{hi}$  for each action  $x$  (so that each row sums to 1) and observe that we have the following structure:

$$\mathbf{E}(0) = \begin{bmatrix} 0 & e_{12} & 1 - e_{12} & & & & & & & & \\ 0 & 1 - e_{24} & & e_{24} & & & & & & & \\ 0 & 1 - e_{35} & & & e_{35} & & & & & & \\ 0 & & & 1 & & & & & & & \\ 0 & & & & & 1 & & & & & \\ 0 & & & & & & 1 & & & & \\ 0 & & & & & & & 1 & & & \end{bmatrix} \quad \mathbf{E}(1) = \begin{bmatrix} 1 & 0 & 0 & & & & & & & & \\ 1 & 0 & 0 & & & & & & & & \\ 1 & 0 & 0 & & & & & & & & \\ 0 & 0 & 0 & 1 & & & & & & & \\ 0 & 0 & 0 & & 1 & & & & & & \\ 0 & 0 & 0 & & & 1 & & & & & \\ 0 & 0 & 0 & & & & 1 & & & & \end{bmatrix}.$$

A few comments are in order. First, for  $x = 1$  (no mammogram screening), the emission matrix  $\mathbf{E}(1)$  is extremely sparse with entries in  $\{0, 1\}$ . For instance, when hidden state equals 1 (healthy), 2 (undiagnosed in-situ), or 3 (undiagnosed invasive), we observe no signal (emission equals 1) w.p. 1. When hidden state equals 4, 5, 6, or 7, we naturally observe the same emission w.p. 1.

Second, for  $x = 0$  (screening), the emission matrix  $\mathbf{E}(0)$  is quite sparse as well. If the patient is healthy (row 1), then the test result is negative (true negative) w.p.  $e_{12}$  and positive (false positive) w.p.  $1 - e_{12}$ . When the patient has undiagnosed in-situ cancer (row 2), it is detected (true positive) w.p.  $e_{24}$  and missed (false negative) w.p.  $1 - e_{24}$ . The parameter  $e_{35}$  has the same interpretation as  $e_{24}$  but for invasive cancer. As for  $x = 1$ , when hidden state equals 4, 5, 6, or 7, we observe the same emission w.p. 1.

Having discussed the structure of  $\mathbf{E}$ , we now calibrate it to real-data. There are 3 parameters:  $e_{12}$ ,  $e_{24}$ , and  $e_{35}$ . All of them can be age specific but we ignore that.  $e_{12}$  is the *specificity* of the mammogram screening (i.e., probability of a true negative) and we calibrate it using Table 3 of [Ayer et al. \(2012\)](#):

$$e_{12} = 0.9.$$

$e_{24}$  is the in-situ *sensitivity* (i.e., probability of a true positive) and we calibrate it using Table 3 of [Ayer et al. \(2012\)](#):

$$e_{24} = 0.8.$$

Finally,  $e_{35}$  is the invasive sensitivity and following [Ayer et al. \(2012\)](#), we set

$$e_{35} = e_{24}.$$

## E.2. Reducing the Number of Joint Decision Variables by Exploiting Sparsity

Recall from §4.3 the following setup, which we repeat for convenience. Let  $k \equiv (h, x)$  and  $m \equiv (h, i)$  so that

$$\begin{aligned} O_k &\equiv O_{hx}, \quad e_{ki} \equiv e_{hxi} \\ H_m &\equiv H_{hi}, \quad q_{mh'} \equiv q_{hih'}. \end{aligned}$$

We have  $k \in [K]$  and  $m \in [M]$ , where  $K := |\mathbb{H}||\mathbb{X}|$  and  $M := |\mathbb{H}||\mathbb{O}|$ . The  $K$  and  $M$  dimensional joint PMFs for all  $i_1, \dots, i_K \in \mathbb{O}$  and  $h_1, \dots, h_M \in \mathbb{H}$  are defined as

$$\theta_{1,\dots,K}(i_1, \dots, i_K) := \mathbb{P}(O_1 = i_1, \dots, O_K = i_K) \tag{9a}$$

$$\pi_{1,\dots,M}(h_1, \dots, h_M) := \mathbb{P}(H_1 = h_1, \dots, H_M = h_M). \tag{9b}$$

As discussed towards the end of §4.3, it follows from (9) that we have *at most*  $|\mathbb{O}|^{|\mathbb{H}||\mathbb{X}|} + |\mathbb{H}|^{|\mathbb{H}||\mathbb{O}|}$  joint variables. We now show that these are merely upper bounds and we can exploit the sparsity inherent in the underlying application to drastically reduce these numbers.

Consider the  $\theta_{1,\dots,K}$  decision variables for now. Since  $\theta_{1,\dots,K}$  represents the joint PMF of the random variables  $[O_k]_k$  where  $k \equiv (h, x)$ , we first understand which  $(h, x)$  pairs are valid (as opposed to naively considering all  $(h, x) \in \mathbb{H} \times \mathbb{X}$ ). Recall that the state  $h$  has the following encoding:

1. healthy
2. undiagnosed in-situ cancer
3. undiagnosed invasive cancer
4. diagnosed in-situ cancer
5. diagnosed invasive cancer
6. recovery
7. death.

Furthermore,  $x$  equals 0 maps to mammogram being performed and 1 to it not being performed. It is easy to see that all 14 combinations of  $(h, x) \in \mathbb{H} \times \mathbb{X}$  are valid so none of the corresponding decision variables can be set to zero (and therefore removed). Turning now to the observations, we recall that they are encoded as follows:

1. no screening took place
2. negative screening result (possibly a false negative)
3. positive mammogram result, but followed by a negative biopsy
4. diagnosed in-situ cancer
5. diagnosed invasive cancer
6. recovery
7. death.

Given this, Table 1 documents the range of all  $7 \times 2 = 14$  random variables  $[O_{hx}]_{h,x}$ .

Table 1. Range of the 14 random variables  $[O_{hx}]_{h,x}$  corresponding to  $\theta_{1,\dots,K}$ .

State $h$	Policy $x$	Range of $O_{hx}$	Range cardinality
1	0	{2, 3}	2
1	1	{1}	1
2	0	{2, 4}	2
2	1	{1}	1
3	0	{2, 5}	2
3	1	{1}	1
4	0	{4}	1
4	1	{4}	1
5	0	{5}	1
5	1	{5}	1
6	0	{6}	1
6	1	{6}	1
7	0	{7}	1
7	1	{7}	1

Multiplying all of the 14 cardinalities (last column in Table 1) implies that there are only eight  $\theta_{1,\dots,K}$  decision variables that need to be considered. This is in contrast to the upper bound of  $|\mathbb{O}|^{|\mathbb{H}||\mathbb{X}|} = 7^{14}$ .

The same logic applies to the  $\pi_{1,\dots,M}$  decision variables. In fact, for the  $\pi_{1,\dots,M}$  decision variables, even the first step proves useful since not all  $(h, i)$  pairs are valid. For instance, if  $h = 1$ , then  $i \notin \{4, 5, 6, 7\}$ . In particular, the first step allows us to trim down the number of  $[H_{hi}]_{h,i}$  random variables from  $|\mathbb{H}||\mathbb{O}| = 49$  to 13. The second step trims down the range of each of the 13 random variables. We document this in Table 2 and are able to reduce the number of  $\pi_{1,\dots,M}$  decision variables from  $|\mathbb{H}|^{|\mathbb{H}||\mathbb{O}|} = 7^{49}$  to 15, 552 (which equals the product of the cardinalities presented in the last column).

Table 2. Range of the 13 random variables  $[H_{hi}]_{h,i}$  corresponding to  $\pi_{1,\dots,M}$ . Only 13  $(h, i)$  pairs are shown as the other 36 are not valid.

State $h$	Observation $i$	Range of $H_{hi}$	Range cardinality
1	1	{1, 2, 3}	3
1	2	{1, 2, 3}	3
1	3	{1, 2, 3}	3
2	1	{2, 3}	2
2	2	{2, 3}	2
2	4	{4, 6}	2
3	1	{3, 7}	2
3	2	{3, 7}	2
3	5	{5, 6, 7}	3
4	4	{4, 6}	2
5	5	{5, 6, 7}	3
6	6	{6}	1
7	7	{7}	1

### E.3. Details on the Pathwise Monotonicity (PM) Constraints

PM can be enforced via linear constraints. We briefly discussed this in §4.3 and now discuss all underlying PM constraints we embedded in our breast cancer numerics.

Recalling our §4.3 discussion for convenience, suppose the patient has in-situ cancer in period  $t$  which is not detected but the patient's state remains at in-situ in period  $t + 1$ . Then, in the counterfactual world, if the cancer is detected in period  $t$ , then PM would require that the cancer can not be worse than in-situ in period  $t + 1$ , i.e.,

$$\mathbb{P}(H_{\tilde{h}\tilde{i}} = \tilde{h}' \mid H_{hi} = h') = 0$$

for  $h = 2, i \in \{1, 2\}, h' = 2, \tilde{h} \in \{2, 4\}, \tilde{i} = 4, \tilde{h}' \in \{5, 7\}$ . There can be multiple such cases to consider and we can enforce all the PM constraints by setting the corresponding  $\pi_{\tilde{h}\tilde{i}, h'}(\tilde{h}', h')$  variables equal to 0 as  $\pi_{\tilde{h}\tilde{i}, h'}(\tilde{h}', h') = \mathbb{P}(H_{hi} = h' \mid H_{\tilde{h}\tilde{i}} = \tilde{h}' \mid H_{hi} = h')$ .

Hence, to provide details on which all PM constraints we enforce, it suffices to enumerate the  $(h, i, h', \tilde{h}, \tilde{i}, \tilde{h}')$  combinations for which we set the  $\pi_{\tilde{h}\tilde{i}, h'}(\tilde{h}', h')$  variables equal to 0. To do so, we iterate over each state  $h \in \{1, \dots, 7\}$ . (Note that for PM, there are no  $(h, x, i, \tilde{h}, \tilde{x}, \tilde{i})$  combinations for which we set the  $\theta_{\tilde{h}\tilde{x}, h_x}(\tilde{i}, i)$  variables equal to 0.)

**State  $h = 1$  (healthy).** We enforce PM for the following combinations:

- If  $(h, i, h')$  equals (healthy, whatever emission, healthy), then the counterfactual state  $\tilde{h}'$  can not be in-situ, invasive, or death if  $\tilde{h}$  is healthy. That is,  $h = 1, i \in \mathbb{O}, h' = 1, \tilde{h} = 1, \tilde{i} \in \mathbb{O}$ , and  $\tilde{h}' \in \{2, 3, 4, 5, 7\}$ .
- If  $(h, i, h')$  equals (healthy, whatever emission, in-situ), then the counterfactual state  $\tilde{h}'$  can not be healthy, invasive, or death if  $\tilde{h}$  is healthy. That is,  $h = 1, i \in \mathbb{O}, h' = 2, \tilde{h} = 1, \tilde{i} \in \mathbb{O}$ , and  $\tilde{h}' \in \{1, 3, 5, 7\}$ .
- If  $(h, i, h')$  equals (healthy, whatever emission, invasive), then the counterfactual state  $\tilde{h}'$  can not be healthy, in-situ, or death if  $\tilde{h}$  is healthy. That is,  $h = 1, i \in \mathbb{O}, h' = 3, \tilde{h} = 1, \tilde{i} \in \mathbb{O}$ , and  $\tilde{h}' \in \{1, 2, 4, 7\}$ .
- If  $(h, i, h')$  equals (healthy, whatever emission, death), then the counterfactual state  $\tilde{h}'$  can not be healthy, in-situ, or invasive if  $\tilde{h}$  is healthy. That is,  $h = 1, i \in \mathbb{O}, h' = 7, \tilde{h} = 1, \tilde{i} \in \mathbb{O}$ , and  $\tilde{h}' \in \{1, 2, 3, 4, 5\}$ .

**State  $h = 2$  (undiagnosed in-situ).** We enforce PM for the following combinations:

- If  $(h, i, h')$  equals (in-situ, undetected, in-situ), then the counterfactual state  $\tilde{h}'$  can not be invasive or death if  $\tilde{h}$  is healthy or in-situ. That is,  $h = 2, i \in \{1, 2, 3\}, h' = 2, \tilde{h} \in \{1, 2, 4\}, \tilde{i} \in \mathbb{O}$ , and  $\tilde{h}' \in \{3, 5, 7\}$ .

- If  $(h, i, h')$  equals (in-situ, detected, in-situ), then the counterfactual state  $\tilde{h}'$  can not be invasive or death if  $\tilde{h}$  is in-situ and detected. That is,  $h = 2, i = 4, h' = 4, \tilde{h} \in \{2, 4\}, \tilde{i} = 4$ , and  $\tilde{h}' \in \{3, 5, 7\}$ .
- If  $(h, i, h')$  equals (in-situ, undetected, invasive), then the counterfactual state  $\tilde{h}'$  can not be death if  $\tilde{h}$  is healthy or in-situ. That is,  $h = 2, i \in \{1, 2, 3\}, h' = 3, \tilde{h} \in \{1, 2, 4\}, \tilde{i} \in \mathbb{O}$ , and  $\tilde{h}' = 7$ .
- If  $(h, i, h')$  equals (in-situ, detected, invasive), then the counterfactual state  $\tilde{h}'$  can not be death if  $\tilde{h}$  is in-situ and detected. That is,  $h = 2, i = 4, h' = 5, \tilde{h} \in \{2, 4\}, \tilde{i} = 4$ , and  $\tilde{h}' = 7$ .
- If  $(h, i, h')$  equals (in-situ, detected, recovered), then the counterfactual state  $\tilde{h}'$  can not be in-situ, invasive, or death if  $\tilde{h}$  is in-situ and detected. That is,  $h = 2, i = 4, h' = 6, \tilde{h} \in \{2, 4\}, \tilde{i} = 4$ , and  $\tilde{h}' \in \{2, 3, 4, 5, 7\}$ .
- If  $(h, i, h')$  equals (in-situ, undetected, death), then the counterfactual state  $\tilde{h}'$  can not be in-situ, invasive, or recovered if  $\tilde{h}$  is in-situ and undetected. That is,  $h = 2, i \in \{1, 2, 3\}, h' = 7, \tilde{h} = 2, \tilde{i} \in \{1, 2, 3\}$ , and  $\tilde{h}' \in \{2, 3, 4, 5, 6\}$ .
- If  $(h, i, h')$  equals (in-situ, detected, death), then the counterfactual state  $\tilde{h}'$  can not be in-situ, invasive, or recovered if  $\tilde{h}$  is in-situ and detected. That is,  $h = 2, i = 4, h' = 7, \tilde{h} \in \{2, 4\}, \tilde{i} = 4$ , and  $\tilde{h}' \in \{2, 3, 4, 5, 6\}$ .

**State  $h = 3$  (undiagnosed invasive).** We enforce PM for the following combinations:

- If  $(h, i, h')$  equals (invasive, undetected, invasive), then the counterfactual state  $\tilde{h}'$  can not be death if  $\tilde{h}$  is healthy, in-situ, or invasive. That is,  $h = 3, i \in \{1, 2, 3\}, h' = 3, \tilde{h} \in \{1, 2, 3, 4, 5\}, \tilde{i} \in \mathbb{O}$ , and  $\tilde{h}' = 7$ .
- If  $(h, i, h')$  equals (invasive, detected, invasive), then the counterfactual state  $\tilde{h}'$  can not be death if  $\tilde{h}$  is invasive and detected. That is,  $h = 3, i = 5, h' = 5, \tilde{h} \in \{3, 5\}, \tilde{i} = 5$ , and  $\tilde{h}' = 7$ .
- If  $(h, i, h')$  equals (invasive, detected, recovered), then the counterfactual state  $\tilde{h}'$  can not be invasive or death if  $\tilde{h}$  is invasive and detected. That is,  $h = 3, i = 5, h' = 6, \tilde{h} \in \{3, 5\}, \tilde{i} = 5$ , and  $\tilde{h}' \in \{3, 5, 7\}$ .
- If  $(h, i, h')$  equals (invasive, undetected, death), then the counterfactual state  $\tilde{h}'$  can not be invasive or recovered if  $\tilde{h}$  is invasive and undetected. That is,  $h = 3, i \in \{1, 2, 3\}, h' = 7, \tilde{h} \in \{3, 5\}, \tilde{i} \in \{1, 2, 3\}$ , and  $\tilde{h}' \in \{3, 5, 6\}$ .
- If  $(h, i, h')$  equals (invasive, detected, death), then the counterfactual state  $\tilde{h}'$  can not be invasive or recovered if  $\tilde{h}$  is invasive and detected. That is,  $h = 3, i = 5, h' = 7, \tilde{h} \in \{3, 5\}, \tilde{i} = 5$ , and  $\tilde{h}' \in \{3, 5, 6\}$ .

**State  $h = 4$  (diagnosed in-situ).** We enforce PM for the following combinations:

- If  $(h, i, h')$  equals (in-situ, detected, in-situ), then the counterfactual state  $\tilde{h}'$  can not be invasive, recovered, or death if  $\tilde{h}$  is in-situ and detected. That is,  $h = 4, i = 4, h' = 4, \tilde{h} \in \{2, 4\}, \tilde{i} = 4$ , and  $\tilde{h}' \in \{5, 6, 7\}$ .
- If  $(h, i, h')$  equals (in-situ, detected, invasive), then the counterfactual state  $\tilde{h}'$  can not be in-situ, recovered, or death if  $\tilde{h}$  is in-situ and detected. That is,  $h = 4, i = 4, h' = 5, \tilde{h} \in \{2, 4\}, \tilde{i} = 4$ , and  $\tilde{h}' \in \{4, 6, 7\}$ .
- If  $(h, i, h')$  equals (in-situ, detected, recovery), then the counterfactual state  $\tilde{h}'$  can not be in-situ, invasive, or death if  $\tilde{h}$  is in-situ and detected. That is,  $h = 4, i = 4, h' = 6, \tilde{h} \in \{2, 4\}, \tilde{i} = 4$ , and  $\tilde{h}' \in \{4, 5, 7\}$ .
- If  $(h, i, h')$  equals (in-situ, detected, death), then the counterfactual state  $\tilde{h}'$  can not be in-situ, invasive, or recovered if  $\tilde{h}$  is in-situ and detected. That is,  $h = 4, i = 4, h' = 7, \tilde{h} \in \{2, 4\}, \tilde{i} = 4$ , and  $\tilde{h}' \in \{4, 5, 6\}$ .

**State  $h = 5$  (diagnosed invasive).** We enforce PM for the following combinations:

- If  $(h, i, h')$  equals (invasive, detected, invasive), then the counterfactual state  $\tilde{h}'$  can not be recovered or death if  $\tilde{h}$  is invasive and detected. That is,  $h = 5, i = 5, h' = 5, \tilde{h} \in \{3, 5\}, \tilde{i} = 5$ , and  $\tilde{h}' \in \{6, 7\}$ .

- If  $(h, i, h')$  equals (invasive, detected, recovery), then the counterfactual state  $\tilde{h}'$  can not be invasive or death if  $\tilde{h}$  is invasive and detected. That is,  $h = 5, i = 5, h' = 6, \tilde{h} \in \{3, 5\}, \tilde{i} = 5$ , and  $\tilde{h}' \in \{5, 7\}$ .
- If  $(h, i, h')$  equals (invasive, detected, death), then the counterfactual state  $\tilde{h}'$  can not be invasive or recovery if  $\tilde{h}$  is invasive and detected. That is,  $h = 5, i = 5, h' = 7, \tilde{h} \in \{3, 5\}, \tilde{i} = 5$ , and  $\tilde{h}' \in \{5, 6\}$ .

**State  $h = 6$  (recovery).** No combination for which we enforce PM.

**State  $h = 7$  (death).** No combination for which we enforce PM.

#### E.4. Details on the Comonotonic Copula

We discussed the counterfactual simulation under the comonotonic copula for a general dynamic latent-state model in §C.2. In this section, we connect that discussion to the breast cancer application. To do so, it suffices to define the rank functions  $r_H(\cdot)$  (for states) and  $r_O(\cdot)$  (for emissions). For states, there are two possible orderings that seem “natural” (from “best” to “worst”):

- (1, 6, 4, **2**, 5, 3, 7)
- (1, 6, 4, **5**, 2, 3, 7).

Recalling the  $\mathbf{Q}(i)$  notation from §E.1, observe that columns 2 and 5 are never “active” simultaneously in any row of  $\mathbf{Q}(i)$  (for any  $i$ ). Hence, the choice of ordering (between the two orderings above) will not matter and we can pick any one. Suppose we pick the first one. Then, this ordering defines the rank function. For example,  $r_H(6) = 2$ , i.e., rank of state 6 equals 2. For the inverse function,  $r_H^{-1}(2) = 6$ .

It is unclear how to define  $r_O(\cdot)$  for the breast cancer application but as it turns out, it does not matter. To see why, consider the generic path of interest (from (1)):

$$\underbrace{o_1, \dots, o_{\tau_s-1}}_{\in \{2,3\}}, \underbrace{o_{\tau_s}, \dots, o_{\tau_e}}_{=1}, \underbrace{o_{\tau_e+1}, \dots, o_{\tau_d-1}}_{\in \{4,5\}}, \underbrace{o_{\tau_d:T}}_{=7}.$$

For the first  $\tau_s - 1$  periods, observe that the counterfactual emission  $\tilde{o}_{1:\tau_s-1}(b)$  equals the observed emission  $o_{1:\tau_s-1}$  for each  $b$ . This is because the intervention policy  $\tilde{x}_{1:\tau_s-1}$  equals the observed policy  $x_{1:\tau_s-1}$ . Now, consider periods  $\tau_s$  to  $\tau_e$ , during which the screening was not done, i.e.,  $x_{\tau_s:\tau_e} = 1$ . Hence, the emissions  $o_{\tau_s:\tau_e} = 1$  w.p. 1 (see the matrix  $\mathbf{E}(1)$  in §E.1.3). This means that the emissions does not contain any information regarding the underlying noise variables  $V_{\tau_s:\tau_e}$  (see Figure 6) and hence, their posterior equals their prior, which is  $\text{Unif}[0, 1]$ . As such, for  $t \in \{\tau_s, \dots, \tau_e\}$ , we can sample  $\tilde{o}_t(b)$  using the categorical distribution over the probability vector  $[e_{\tilde{h}_t(b)\tilde{x}_t i}]_i$ . Note that we can use  $\tilde{o}_{\tau_s-1}(b)$  to sample  $\tilde{h}_{\tau_s}(b)$ , which we can use to sample  $\tilde{o}_{\tau_s}(b)$ , and so on (until we have sampled  $\tilde{h}_{\tau_e+1}(b)$ ). Now, consider  $t = \tau_e + 1$ . We know  $o_t \in \{4, 5\}$ :

- If  $o_t = 4$ , then  $h_t(b) = 2$  and  $\tilde{h}_t(b) \in \{2, 4, 6\}$  (cf. pathwise monotonicity).
  - If  $\tilde{h}_t(b) \in \{4, 6\}$ , then  $\tilde{o}_t(b) = \tilde{h}_t(b)$  (since rows 4 and 6 of  $\mathbf{E}(0)$  have 1 on the diagonal).
  - Else, if  $\tilde{h}_t(b) = 2 (= h_t(b))$ , then  $\tilde{o}_t(b) = o_t = 4$ .
- Else, if  $o_t = 5$ , then  $h_t(b) = 3$  and  $\tilde{h}_t(b) \in \{3, 4, 5, 6\}$  (cf. pathwise monotonicity).
  - If  $\tilde{h}_t(b) \in \{4, 5, 6\}$ , then  $\tilde{o}_t(b) = \tilde{h}_t(b)$  (since rows 4, 5, 6 of  $\mathbf{E}(0)$  have 1 on the diagonal).
  - If  $\tilde{h}_t(b) = 3 (= h_t(b))$ , then  $\tilde{o}_t(b) = o_t = 5$ .

Finally, for  $t \geq \tau_e + 2$ , we know  $h_t(b) \geq 4$  and that the corresponding rows in  $\mathbf{E}(0)$  are 0-1. Hence, the posterior of  $V_t$  equals the prior and we can sample  $\tilde{o}_t(b)$  using the categorical distribution over the probability vector  $[e_{\tilde{h}_t(b)\tilde{x}_t i}]_i$ . By construction, the comonotonic copula will obey pathwise monotonicity and hence, will ensure that in the counterfactual world, patient does not die before period  $T$ , i.e.,  $H_{T-1} \neq 7$  w.p. 1.

E.5. Results for Path 2




Figure 10. PN results for path 2 as we vary  $T \in \{4, \dots, 10\}$  (analogous to Figure 3 for path 1). As in Figure 3, observe that the LB, LB (CS), and LB (PM) curves coincide (the lowest curve in each figure). Further, the UB (CS) and UB (PM) curves coincide for path 2.