Hypertension Management: A Value of Information Perspective

Research Team:

Manaf Zargoush[1], Mehmet Gumus[2], Vedat Verter[3], & Stella Daskalopoulou[4]

[1] DeGroote School of Business, McMaster University, Canada

[2] Desautels Faculty of Management, McGill University, Canada

[3] Eli Board College of Business, Michigan State University, USA

[4] Faculty of Medicine, McGill University, Canada

Presentation:

Manaf Zargoush

Assistant Professor; Health Policy & Management

DeGroote School of Business

McMaster University

zargoush@mcmaster.ca

Agenda

✓Background

✓Problem Statement

✓Methods (Prediction and Optimization)

✓Results and Discussions

Hypertension (HTN)

- A chronic medical condition in which the blood pressure (BP) in vessels elevates to a level higher than its normal range.
  - ✓ major Public health issue worldwide;
  - ✓ highly prevalent with serious consequences (One billion in the world and $\frac{1}{3}$ in the US)

- **Mechanism**

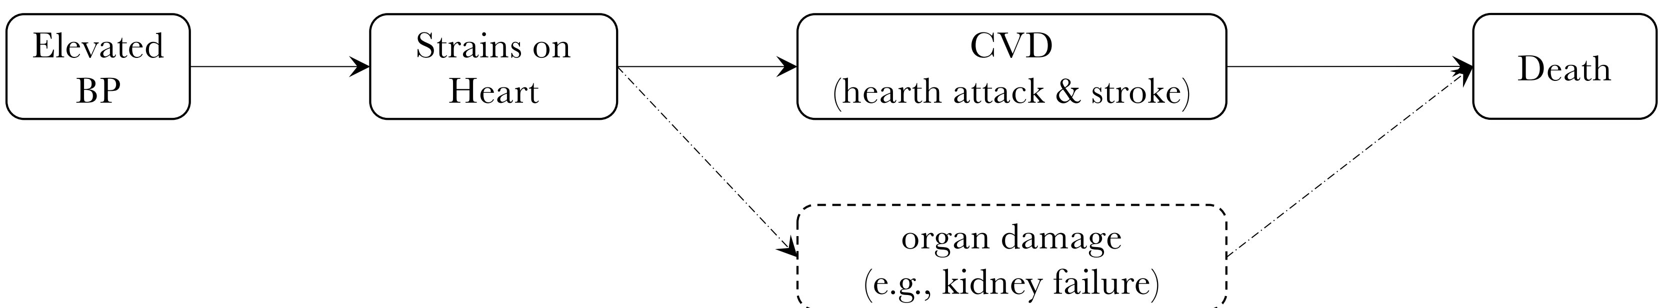

Importance of HTN Control

The **leading risk for death** in North America (WHO)

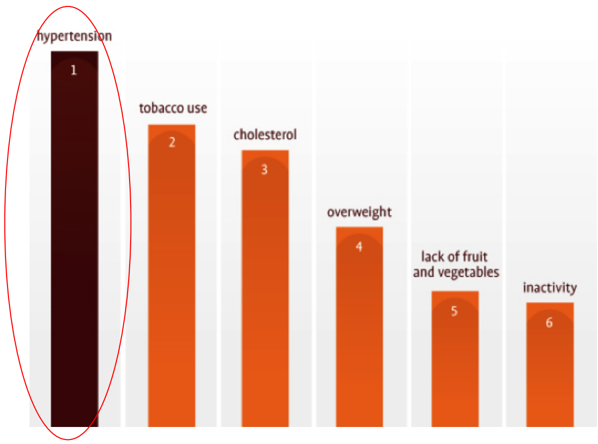

77% of first stroke events occur among patients with HTN

treating **HTN:**

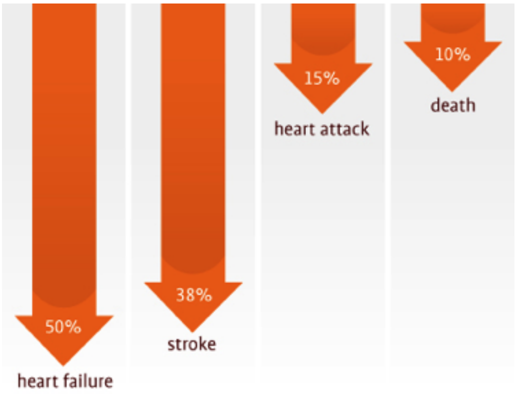

HTN Control

- **Good news**
  - ✓HTN can be controlled with promising benefits

- **Bad news**
  - ✓Only a few hypertensive patients have their BP under control
    - In the US, less than half of patients with HTN have it controlled!

Reasons for The Poor Control of HTN?

1. HTN is **asymptomatic**
    ✓**silent killer**
    ✓solution: keep track of BP

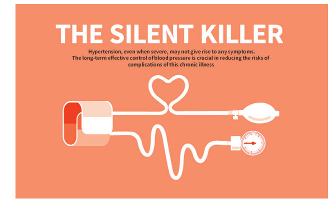

2. BP is complex; fluctuates both in the short- and long-terms.
    ✓very difficult consistent and reliable measurement of BP.

3. Traditional BP measurement is **noisy**
    ✓Obscuring the true underlying BP.

Solution:
increase the **accuracy** of:
- measurement
- predictions

4. 1,2,3 => profound **subjectivity** in clinical decision-making!
    ✓ physician inertia:
    - *Physician's failure to adequately adjust treatment (i.e., add medication) in response to elevated BP*

med. literature

"**humanistic**" issue related to "**physician behavior**" or **judgment bias**
"to err is human!"

BP Measurements

➢ Measurement

1. Traditional approach (gold standard)

- ✓ peripheral BP
- ✓ noisy : ***inaccurate***

2. New technologies

- ✓ e.g., tech. based on ultrasound
- ✓ Applanation Tonometry or Automated Office BP (AOBP)
- ✓ ***noise-free or at least less noisy***
- ✓ more costly (staff, time, technology, etc.)

low adoption of these technologies
→ uncertainty over their benefits vs cost!

→

Value of Information (VOI)
comparing our best decisions:
in the presence and absence of information

HTN Control

1. **Measurement**
   1. Systolic BP (SBP): usually on quarterly basis

2. Treatment

- Medication therapy through a class of medications called ***antihypertensives***
- Usually combination therapy (i.e., multiple medications)

Five common classes of antihypertensives:

1. **Beta Blockers**
2. **ACEI** (Angiotensin-Converting-Enzyme Inhibitor)
3. **ARBs:** Angiotensin II Receptor Blockers
4. **Diuretic** (aka. thiazide)
5. **CCBs** (Calcium Channel Blockers)

Problem Statement

➢ How HTN can be controlled considering:

✓ **Measurement uncertainty**

- **underestimation vs. overestimation**

Learning (prediction)

✓ **Intervention trade-off**

- **Optimal course actions (medication therapy)**
- **too early (unnecessary medication side-effects) vs. too late (risk of CVD)**

Optimization

➢ **From analytics perspective**

➢ How to effectively marry **predictive analytics** and **prescriptive analytic** → **VOI?**

↓

focus of today's presentation

How BP Evolves in the short- and long-term?

1. Everyone has a mean BP (θ_t): changes over time and **is unobservable** → basis for physician's medication decision
2. In the short-term (e.g., daily), one's BP observation (b_t) varies according to a Normal distribution with
   - ✓ mean= θ_t
   - ✓ variances= person's short-term BP variability (σ_b^2) **+** measurement noise (τ^2)
3. In the long-term (e.g., quarterly), θ_t changes according to a Normal distribution such that:
   - ✓ mean at $t+1$= mean at t **+** known/deterministic changes (such as change due to aging and medications)
   - ✓ variance= person's long-term BP variability (σ_θ^2)

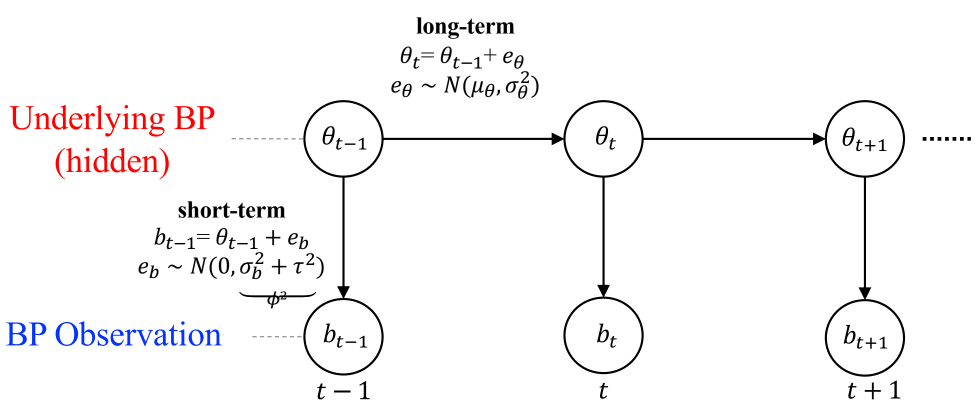

The Framework for Prediction and Decision-Making

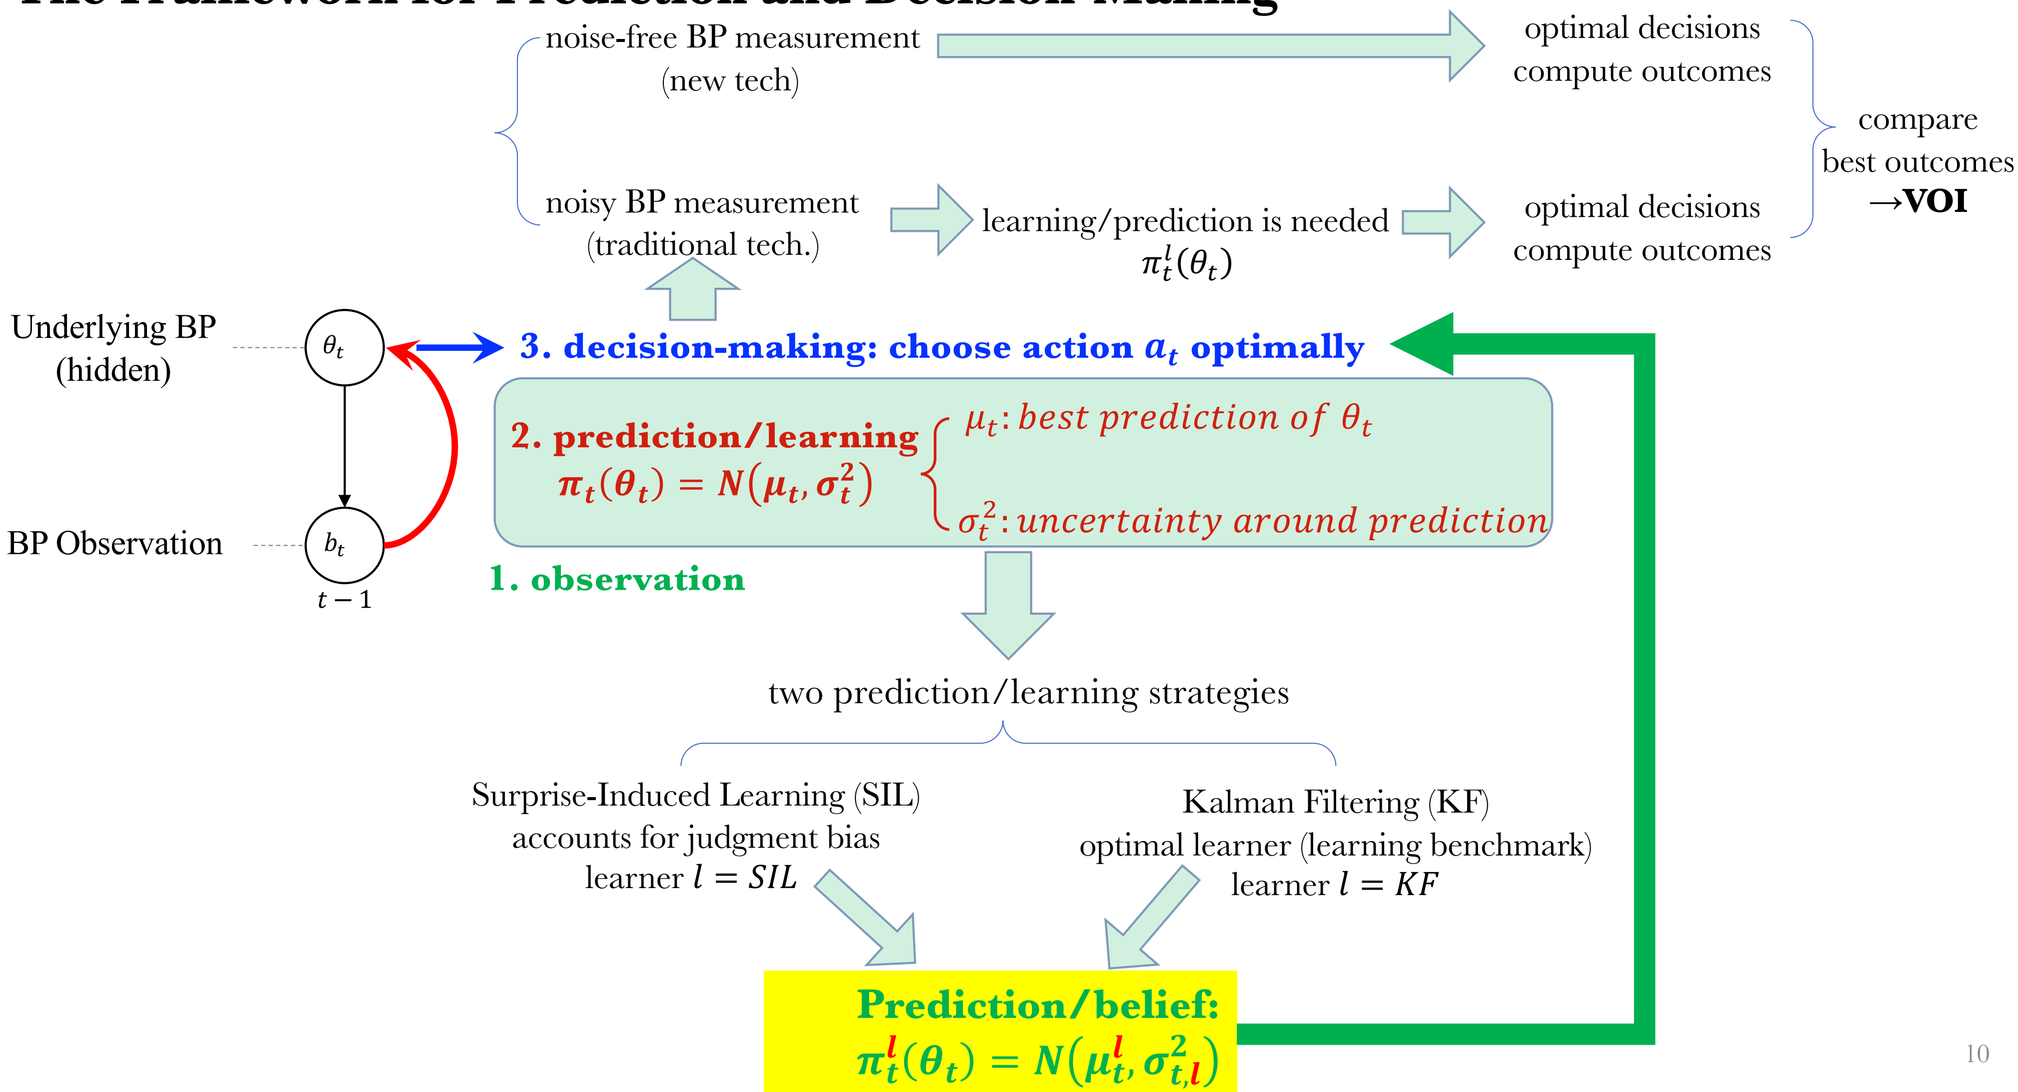

The Problem

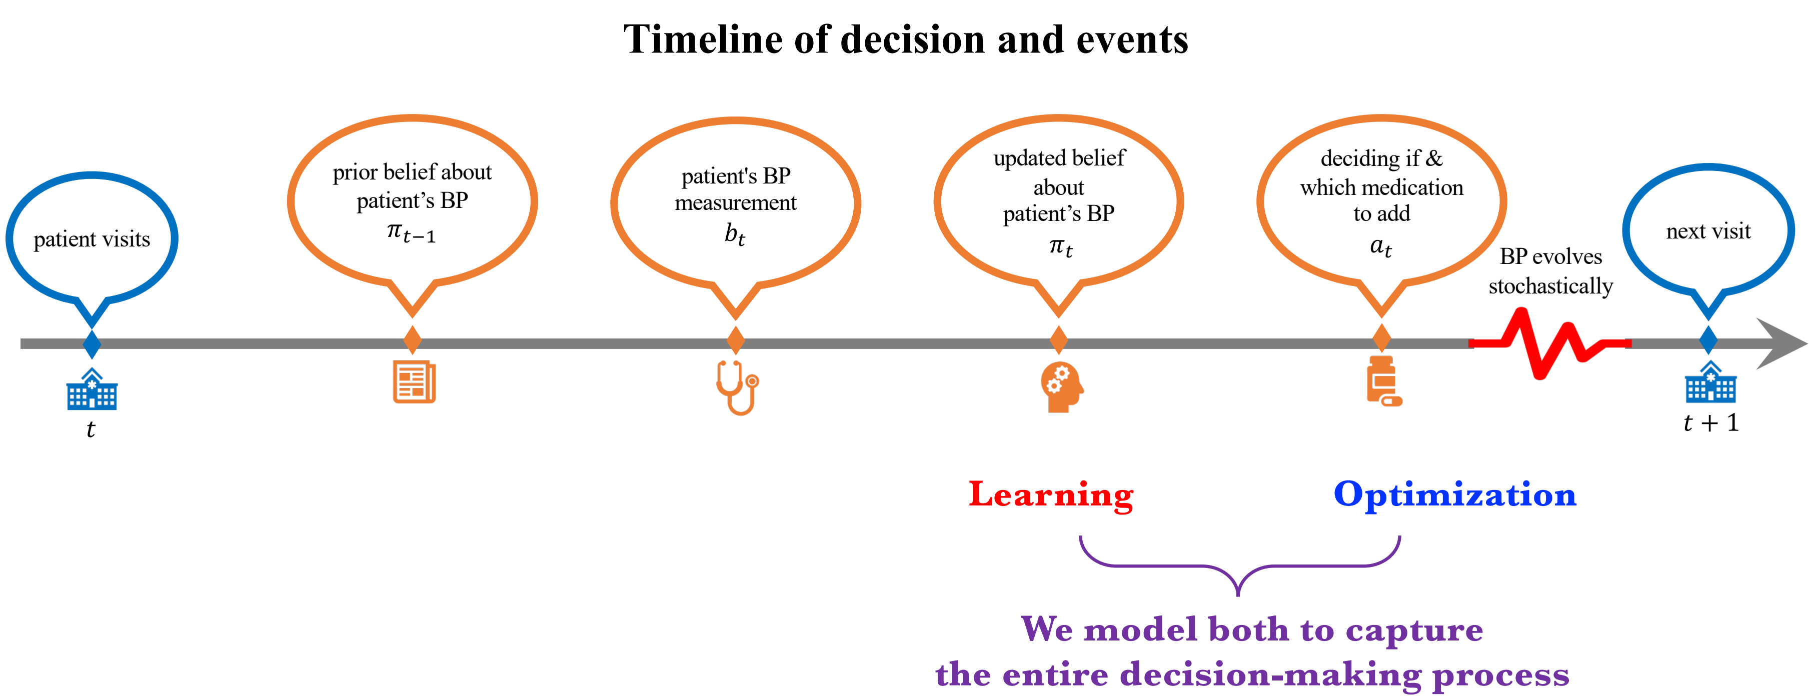
Timeline of decision and events
patient visits
t
prior belief about patient's BP π_{t-1}
patient's BP measurement b_t
updated belief about patient's BP π_t
deciding if & which medication to add a_t
BP evolves stochastically
next visit
$t+1$
Learning
Optimization
We model both to capture the entire decision-making process

KF Learning

- KF characterized the parameters of belief about θ_t, i.e., $\pi_t^{KF}(\theta)$ as follows:

$$\begin{cases} \mu_t^{KF} = K_t \boxed{b_t} + (1 - K_t)\boxed{\mu_{t-1}^{KF}} \\ \sigma_{t,KF}^2 = (1 - K_t)\varsigma_t \end{cases}$$

$$K_t = \frac{\varsigma_t}{\phi^2 + \varsigma_t}$$

The weight we should assign to the new observation/evidence as opposed to the history/prior belief

$$\varsigma_t = \sigma_{t-1,KF}^2 + \sigma_\theta^2$$

- $K_t \in [0,1]$ is called *Kalman gain* which identifies the relative contribution of the new evidence b_t in building the new belief.
- new belief/prediction= convex combination of old prediction and the new observation
- does not account for any subjectivity (hence bias) in predictions

Surprise Induced Learning (SIL): a modified Bayesian updating!

➢Conventional Bayesian Updating → very similar to KF!

$$\pi_t^B(\theta_t) \equiv \pi_t^B(\theta) = \mathcal{N}\left(\mu_t^B, \sigma_{t,B}^2\right)$$

$$\begin{cases} \boldsymbol{\mu_t^B = \rho_t b_t + (1 - \rho_t)\mu_{t-1}^B} \\ \sigma_{t,B}^2 = (1 - \rho_t)\sigma_{t-1,B}^2 \end{cases}$$

$$\rho_t = \frac{\sigma_{t-1,B}^2}{\sigma_{t-1,B}^2 + \phi^2}$$

➢The **issue**:

✓Conventional Bayesian updating assumes stationary mean → belief convergence!

Not reacting to the new observation

SIL Strategy: a modified traditional Bayesian Updating!

- To resolve the **issue**:
  - ✓We introduce the notion of **surprise**, as:

difference between the prior belief and the new evidence

$$s_t = \begin{cases} 1, & \text{if } \boxed{|\mu_{t-1} - b_t|} \geq \Delta \\ 0, & \text{otherwise} \end{cases}$$

threshold for the expectedness of events

surprise ~ observing unexpected events

Signal for BP regimen change

When surprise occurs, <u>we impose a *shock* to the belief/prediction system</u> by resetting the belief updating mechanism such that:

- **maximum belief uncertainty (we are in a surprise state)**
- **maximum weight to new observations → <u>*surprise triggers attention*</u>**
- **minimum weight to the prior belief**

SIL Strategy

surprise-free event

parameters

$$\pi_t^{SIL}(\theta) = \mathcal{N}\left(\mu_t^{SIL}, \sigma_{t,SIL}^2\right)$$

$$\begin{cases} \mu_t^{SIL} = \rho_t b_t + (1-\rho_t)\mu_{t-1}^{SIL}; & \sigma_{t,SIL}^2 = (1-\rho_t)\sigma_{t-1,SIL}^2 \quad \text{if } s_t = 0 \\ \mu_t^{SIL} = \rho_1 b_t + (1-\rho_1)\mu_{t-1}^{SIL}; & \sigma_{t,SIL}^2 = (1-\rho_1)\sigma_{0,SIL}^2 \quad \text{if } s_t = 1 \end{cases}$$

$$\rho_t = \frac{\sigma_{t-1,SIL}^2}{\sigma_{t-1,SIL}^2 + \phi^2} \Rightarrow \rho_1 = \frac{\sigma_{0,SIL}^2}{\sigma_{0,SIL}^2 + \phi^2}$$

surprising event

SIL Strategy

- Δ: characterizes the physician's learning behavior → physician's characteristics → **commitment to belief**
  - ✓ physician with **low Δ** → **undercommitment to belief**
    - experiences more frequent surprises,
    - lower expectation for change → perceiving even small changes as unexpected,
    - assigns larger weights to the new observations → **overestimating evidence**
  - ✓ physician with **high Δ** → **overcommitment to belief**
    - becomes surprised less frequently,
    - higher expectations for change, → ignoring larger changes
    - with time, she assigns higher weights to her own belief → **underestimating evidence**
    - Captures the so-called **physician inertia** (a key obstacle in HTN treatment recently mentioned in the EU Guideline for HTN control)
- Both cases indicate sub-optimal learning behaviors.

Question: Is there an optimal Δ? →
Answer: Yes! The one which minimizes prediction error or maximizes outcomes!

Data Setting at the Montreal General Hospital

➢Two sets of data

1. Noise-Free Environment

- ✓Patients undergoing meticulous BP measurements in the clinic
- ✓Quarterly visits
- ✓Using Automated Office Blood Pressure (AOBP) technology

2. Noisy Environment

- ✓**For the same patient**
- ✓At the same day of clinic visit
- ✓Undergoing 24hr BP measurement, every 20-30min
  - Called 24hr BP measurement or Ambulatory BP Monitoring (ABPM)

Characterizing Optimal Decision Making Through Optimization

- Markov Decision Procession (MDP)
  - ✓Choosing optimal medication decisions to maximize the expected quality adjusted life years of patients over the problem horizon
- Key component:
  - ✓**States**: information needed for making decisions and characterizing the evolution of system → patient's BP mean (either we know it or we learn it)

- Both learning strategies used in our study are <u>Markovian,</u> <u>sequential,</u> and <u>recursive</u> → ideal for MDP

- States in SIL strategy:
  - ✓Best prediction about patient's BP mean → μ_t^{SIL}
  - ✓Number of office visits since last surprise $n_t = \{\mathbf{0}, 1, \ldots, N\}$
    - one-to-one relationship to σ_t^2!
    - measures belief strength

surprise state

Optimization Framework

➢ We develop three MDP models:

Under Noise-Free Measurement:

1. Under noise-free measurement, called $\mathbf{MDP^0}$.

Under Noisy Measurement:

2. Under noisy measurement but KF learning strategy, called $\mathbf{MDP^{KF}}$.
3. Under noisy measurement but SIL learning strategy, called $\mathbf{MDP^{SIL}}$.

Optimal Policies for $MDP^{SIL}(\Delta^*)$

Theorem 5. *Suppose that **As.1-4 - As4-4** hold for $t = 1, 2, ..., T$. Then, at each period t and for fixed levels of μ_t and $\mathbf{m_t}$, there exists an optimal policy $a_t^*(\mu_t, n_t, \mathbf{m_t})$ which is nondecreasing in n_t. In other words, there is a threshold n_t^* such that:*

$$a_t^*(\mu_t, n_t, \mathbf{m_t}) = \begin{cases} i^-, & n_t < n_t^* \\ i^+, & n_t \geq n_t^* \end{cases} \tag{4.25}$$

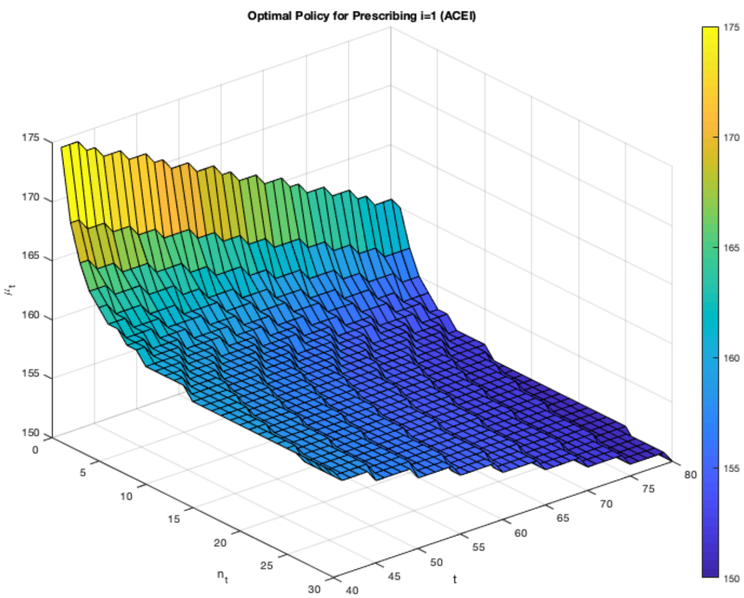

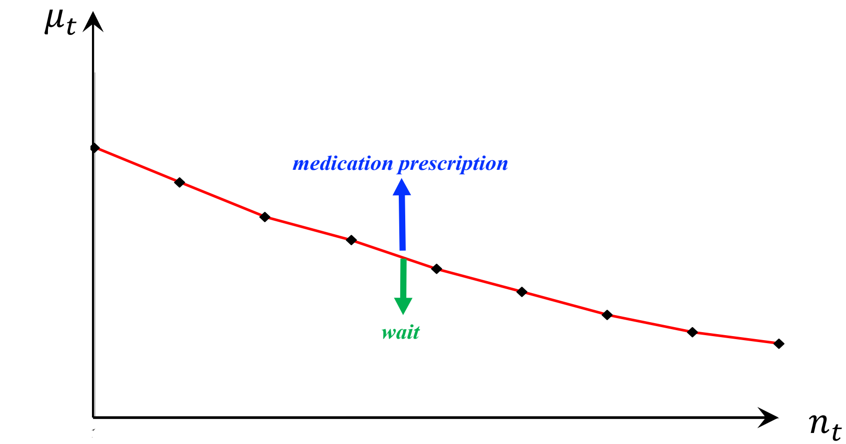

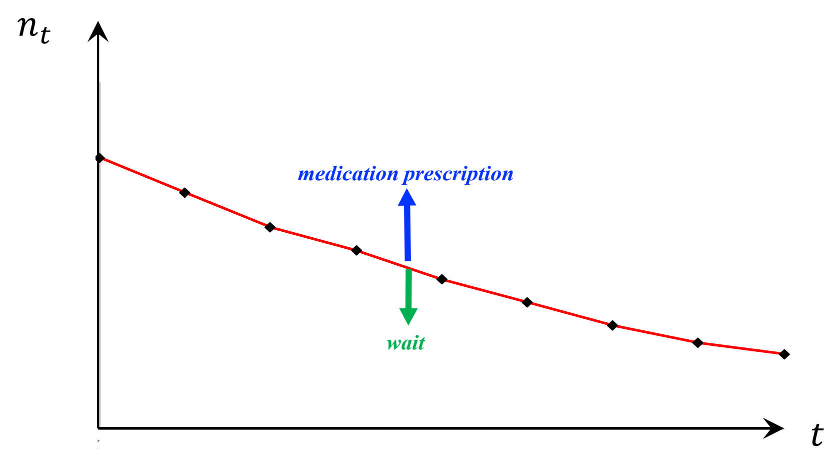

Value of Information (VOI) Analysis

MDP

➢ Definitions:

✓ v^0: value function under perfect information (i.e., noise-free)

✓ v^l: value function under imperfection information, **learner $\boldsymbol{l}$** (i.e., noisy)

$$\boldsymbol{VOI = v^0 - v^l}\textbf{: in terms of Total QALY gained}$$

Improvement in outcomes as a result of reducing uncertainty over θ_t

[max] price paid for reducing uncertainty over θ_t

✓ Ratio of Value of Information (RVOI)

$$\boldsymbol{RVOI = \frac{v^0 - v^l}{v^l}}\textbf{: in terms of \%Total QALY gained}$$

VOI Decomposition : Important Lessons

- More specifically:
  - ✓ v^0: the value function under perfect information.
  - ✓ v^{KF}: the value function under imperfect information, yet KF-learner (learning benchmark)
  - ✓ v^*: the value function under imperfect information, yet Δ^*-learner
  - ✓ v^l: the value function under imperfect information, yet $\Delta^{l\neq *}$-learner
- Therefore:

$$VOI^l = v^0 - v^l = \overbrace{\left(v^0 - v^{KF}\right)}^{KF\to\theta} + \overbrace{\left(v^{KF} - v^*\right)}^{\Delta^*\to KF} + \overbrace{\left(v^* - v^l\right)}^{\Delta^l\to\Delta^*}$$

price paid for **information**
(net value of information)
value of new technology

price paid for
optimal learning <u>strategy</u>

price paid for
optimal learning <u>behavior</u>

Conclusion:

- not all the price we pay is because of not knowing the truth (which can be learned/predicted),
- we also pay for our sub-optimal learning strategy (predictive models) or suboptimal learning behavior!

In our study: $VOI = VOI^l(\sigma_\theta, \sigma_b, \tau, j)$; **l:learner/physician, j:patient's baseline risk**

Conclusions:

- ➢ The new technology is valuable!
- ➢ Its value depends on:
  - ✓ **Patient:** risk profile and her BP variability
  - ✓ **Measurement technology**: current traditional devices
  - ✓ **Physician:** those who are not good learners pay more!

➢ Not all the price we pay for information (because of lack of knowledge) is because of the information itself (that we tend to know or predict); we also pay for our suboptimal learning strategies (predictive models) or suboptimal learning behaviors!

not choosing the best predictive models/methods

not using the predictive models in the best possible way

Thank You!