Behavioral Analytics for Myopic Agents

Phil Kaminsky

Anil Aswani, Yonatan Mintz, Mo Zhou, Elena Flowers, Yoshimi Fukuoka, Ken Goldberg

Support from NSF Award CMMI-1450963, NIH R01 HL104147, and NIH K24 NR015812

BERKELEY IEOR INDUSTRIAL ENGINEERING & OPERATIONS RESEARCH

Behavioral Analytics Framework

Goal: Adaptive data-driven design of incentives

→ **Repeated games (Radner 1985, Fudenberg et al. 1994, Laffont and Martimort 2002)**

→ **Leveraging optimization, estimation to develop practical tools**

→ **Initial work in a healthcare setting**

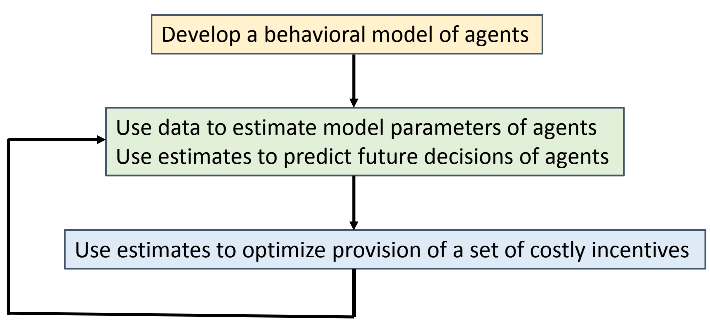

2

Outline

- Behavioral Analytics Framework
- Potential Supply Chain Applications
- Behavioral Analytics for Myopic Agents
  - Healthcare application
    - Prediction and assessment
      - Paper: Behavioral modeling in weight loss interventions
    - Experimental results
      - Paper: Personalizing Mobile Fitness Apps using Reinforcement Learning
  - Generalizations + Optimization
      - Paper: Behavioral analytics for myopic agents
- Extensions Needed for Supply Chain Applications

3

Setting

- More specifically, at time *t*
  - A pool of agents
    - Each with a set of evolving utility function parameters (**motivational state**)
    - Each with a **system state**
    - Each makes a decision by optimizing a **myopic** utility function
  - A single coordinator
    - Noisy o**bservations** of system states, decisions of each agent
    - **Assigns incentives** (behavioral, financial) to (a subset of) agents
  - Incentives + **dynamics** change states at time t+1
  - Process repeats, t advances
- **Coordinator's Decision**
  - **Assign** incentives subject to a budget
  - **Goal**: minimize cost function

Motivational states/utility function parameters unknown
Motivational states not measured
Measurements are noisy, missing
Budget is fixed

4

In A Supply Chain…

Decentralized Supply Chain With Sales Effort

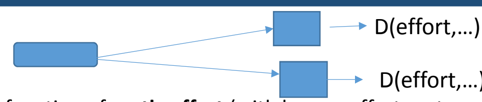

- **Demand** is a (known, stationary, stochastic) function of **costly effort** (with known effort cost function)
- **Goal**: coordinate supply chain, maximize profit, equilibrium prices, **coordinating contracts** (rebates, returns, etc)
- **Extensions:** free rider on effort in omni-channel, endogenous pricing, alternative contract structures, partially observable/verifiable effort, varying risk preferences

- **How can a coordinator use new, noisy data streams (demand, time allocated to selling, survey results, etc.) to maximize incentive allocation strategy?**
  - Seller might want to incentivize quality, customer service, sales
  - How can buyers/agents utility function be estimated?
  - Incentives could be: Training, advertising, events, agent sales efforts, etc.?

Franchise Systems

Key: New (Noisy) Data Streams to Maximize Incentive Allocation Efficiency

- Both the coordinator and the agents have more complex objectives/utility functions.
- The coordinator has many different options of data to monitor (sales, quality, ….)
- Effort is more challenging to estimate, is likely to evolve over time, and can be a function of many different incentives; literature highlights expensive data acquisition, difficult to estimate utilities

Chu and Desai(1995), Lariviere and Padmanabhan(97), Netessine and Rudi (2000), Taylor(2002), Krishnan at al (2004) Mukhopadhyay at al. (2008), Cachon (2003), Xing and Liu (2012), Gal-Or (1995), Li (1997), Xie et al. 2016 5

In A Supply Chain…

Sales Force Compensation

- Firm maximizes profit (some perhaps stochastic function of effort, price)
- Salespeople maximize utility (known function of incentives, effort)
- Effort observed, functions known

- Effort not observed, firm maximizes expected profit
- Demand stochastic
- Distributional assumptions on utility (risk preferences), estimation of utility, relationship to demand

- Commercial software vendors sell software that using ML techniques to try to determine relationship between incentive efforts, sales (but not in an interpretable way)

Key: New (Noisy) Data Streams to Maximize Incentive Allocation Efficiency

- **Question: Can we model utility, optimize incentives? How does our framework need to be extended?**

Coughlan (1993) 6

Our Motivating Problem - Obesity in United States

- Epidemic of obesity
  - 70% of adults are overweight/obese
  - Health costs estimated at $147 billion annually
  - Risk factor for type 2 diabetes and other disease
- Clinical weight loss programs – most effective approach
  - Utilize in-person sessions/goals – effectiveness of sessions decreases w/time
  - Reduce diabetes risk without side-effects
  - Expensive to implement and maintain

Healthcare Challenge:

Innovative approaches are needed to lower costs and increase efficacy of weight loss interventions.

(Flegal et al., 2012); (Finkelstein et al., 2009); (Eyre et al., 2004); (DPP, 2002, 2003, 2009)

7

Personalized Adaptive Physical Activity Programs

Mobile Health for Physical Activity:

How can we design algorithms that customize the components of a program that encourages physical activity + schedules visits?

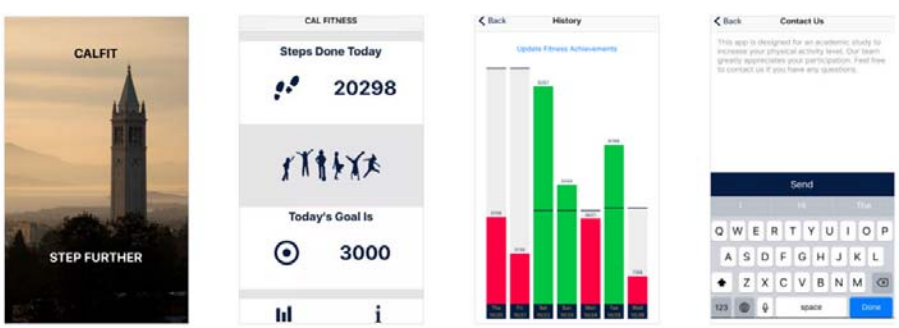

(a) Splash Screen (b) Home Tab (c) History Tab (d) Contact Tab

8

Behavioral Analytics Framework for Weight Loss

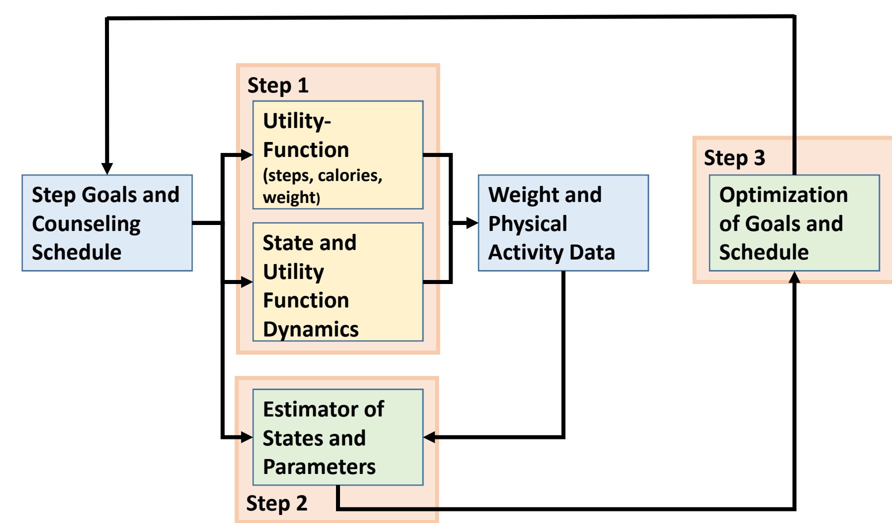

MC: Ayer et al.(2012), Mason et al.(2013), Deo et al.(203); Bandit: Negoescu(2014) Plan design: Bertsimas/O'Hair(20130 9

Step 1: Utility-Maximizing Behavioral Model

- Individual decision-making without exercise goals

weight — number of steps — caloric intake

$$(u_t, f_t) = \arg\max_{u,f} \; -w_{t+1}^2 - r u_t^2 + q u_t - r f_t^2 + s_t f_t$$
$$\text{s.t. } w_{t+1} = a \cdot w_t + b \cdot u_t + f_t + k.$$

$\theta_t = (a, b, k, r, q, s_0, s_t, p_t, \mu, \gamma)$

- Individual decision-making with exercise goals

goal

$$(u_t, f_t) = \arg\max_{u,f} \; -w_{t+1}^2 - r u_t^2 + q u_t - r f_t^2 + s_t f_t + p_t \cdot (u_t - g_t)^-$$
$$\text{s.t. } w_{t+1} = a \cdot w_t + b \cdot u_t + f_t + k.$$

- Temporal evolution for weight and type dynamics

pre-computed — random fluctuations — meeting effectiveness — meeting indicator — step goal

$$w_{t+1} = a \cdot w_t + b \cdot u_t + f_t + k + z_t$$
$$s_{t+1} = \gamma \cdot (s_t - s_0) + s_0 - \beta_t \cdot d_t$$
$$p_{t+1} = \gamma \cdot p_t + \delta_t \cdot d_t + \mu \cdot \mathbf{1}(u_t \geq g_t)$$

Bandura(2001) – Social Cognitive Theory 10

Maximum Likelihood Estimation

- **Goal:** for a specific individual, with limited weight and step data, + goals and visits, estimate all of the parameters
- **Challenge:** Significant amount of missing data and noise
- **Approach:** Jointly estimate parameters and noisy/missing data -- challenging problem

Bilevel Optimization

$$\mathbf{P}_{\text{mle}} \quad \min \sum_{i=1}^{n_w} -\log\psi_\nu(\tilde{w}_{t_i} - w_{t_i}) + \sum_{i=1}^{n_u} -\log\psi_\omega(\tilde{u}_{\tau_i} - u_{\tau_i}) + \sum_{t=1}^{n} -\log\psi_z(z_t)$$
$$\text{s.t. } \mathbf{U}_{\text{no goals}}, (1) \text{ for } t = 1, \ldots, m-1; \quad \mathbf{U}_{\text{goals}}, (1), (2), (3) \text{ for } t = m, \ldots, n.$$

- Difficult since (u_t, f_t) arg max – properties → MILP formulation

$$\mathbf{P}_{\text{mle-milp}} \quad \min \sigma_1^{-1/2} \sum_{i=1}^{n_w} |\tilde{w}_{t_i} - w_{t_i}| + \sigma_2^{-1/2} \sum_{i=1}^{n_u} |\tilde{u}_{\tau_i} - u_{\tau_i}| + \sigma_3^{-1/2} \sum_{t=1}^{n} |z_t|$$
$$\text{s.t. } (1), (4), \text{ for } t = 1, \ldots, m-1; \quad (1), (2), (5), (6), (7), \text{ for } t = m, \ldots, n.$$

Bertsimas et al. 2014 11

Maximum Likelihood Estimation

- **Goal:** for a specific individual, with limited weight and step data, estimate all of the parameters of this model
- **Challenge:** Significant amount of missing data and noise
- **Approach:** Jointly estimate parameters and noisy/missing

PROPOSITION 3. *Necessary and sufficient optimality conditions for* $\mathbf{U}_{\text{no goals}}$ ...

PROPOSITION 4. *Neccesary and sufficient optimality conditions for* $\mathbf{U}_{\text{goals}}$ ...

PROPOSITION 5. *For fixed* g_t, *the logical constraint* $(u_t < g_t, g_{t+1} \geq g_t, d_{t+1} = 0) \Rightarrow (u_{t+1} < g_{t+1})$ ...

PROPOSITION 6. *The dynamics on* p_t *(3) can be represented by the linear constraints:* ...

mization

$\mathbf{P}$ ... $\log\psi_z(z_t)$

... $m, \ldots, n.$

- Difficult since (u_t, f_t) arg max – properties → MILP

$$\mathbf{P}_{\text{mle-milp}} \quad \min \sigma_1^{-1/2} \sum_{i=1}^{n_w} |\tilde{w}_{t_i} - w_{t_i}| + \sigma_2^{-1/2} \sum_{i=1}^{n_u} |\tilde{u}_{\tau_i} - u_{\tau_i}| + \sigma_3^{-1/2} \sum_{t=1}^{n} |z_t|$$
$$\text{s.t. } (1), (4), \text{ for } t = 1, \ldots, m-1; \quad (1), (2), (5), (6), (7), \text{ for } t = m, \ldots, n.$$

Bertsimas et al. 2014 12

Maximum Likelihood Estimation

Appendix A: Complete MILP Formulation for MLE

$$
\begin{aligned}
\min\ & \sigma_1^{-1/2}\textstyle\sum_{i=1}^{n_w}\xi_{w,i}+\sigma_2^{-1/2}\sum_{i=1}^{n_u}\xi_{u,i}+\sigma_3^{-1/2}\sum_{t=1}^{n}\xi_{z,t} & \\
\text{s.t. } & 2b(aw_t+bu_t+f_t+k)+2ru_t-q=0, & \text{for } t=1,\dots,m-1\\
& 2(aw_t+bu_t+f_t+k)+2rf_t-s_0=0, & \text{for } t=1,\dots,m-1\\
& 2b(aw_t+bu_t+f_t+k)+2ru_t-q-\lambda_t^2=0, & \text{for } t=m,\dots,n\\
& 2(aw_t+bu_t+f_t+k)+2rf_t-s_t=0, & \text{for } t=m,\dots,n\\
& g_t-\epsilon-(g_t-\epsilon)\cdot x_t^1\le u_t\le M+(g_t-\epsilon-M)\cdot x_t^1, & \text{for } t=m,\dots,n\\
& (g_t-\epsilon)\cdot x_t^2\le u_t\le M+(g_t+\epsilon-M)\cdot x_t^2, & \text{for } t=m,\dots,n\\
& (g_t+\epsilon)\cdot x_t^3\le u_t\le g_t+\epsilon+(M-g_t-\epsilon)\cdot x_t^3, & \text{for } t=m,\dots,n\\
& 0\le\lambda_t^2\le p_t, & \text{for } t=m,\dots,n\\
& p_t-M\cdot(1-x_t^1)\le\lambda_t^2\le M\cdot(1-x_t^3), & \text{for } t=m,\dots,n\\
& x_t^1,x_t^2,x_t^3\in\{0,1\}, & \text{for } t=m,\dots,n\\
& x_t^1+x_t^2+x_t^3=1, & \text{for } t=m,\dots,n\\
& w_{t+1}=a\cdot w_t+b\cdot u_t+f_t+k+z_t, & \text{for } t=1,\dots,n-1\\
& s_{t+1}=\gamma\cdot(s_t-s_0)+s_0-\beta_{t+1}\cdot d_{t+1}, & \text{for } t=m,\dots,n-1\\
& p_{t+1}\ge\gamma\cdot p_t+\delta_{t+1}\cdot d_{t+1}, & \text{for } t=m,\dots,n-1\\
& p_{t+1}\le\gamma\cdot p_t+\delta_{t+1}\cdot d_{t+1}+M\cdot(1-x_t^1), & \text{for } t=m,\dots,n-1\\
& p_{t+1}\ge\gamma\cdot p_t+\delta_{t+1}\cdot d_{t+1}+\mu-Mx_t^1, & \text{for } t=m,\dots,n-1\\
& p_{t+1}\le\gamma\cdot p_t+\delta_{t+1}\cdot d_{t+1}+\mu, & \text{for } t=m,\dots,n-1\\
& x_{t+1}^1\ge x_t^1-d_{t+1}-1(g_{t+1}-g_t<0), & \text{for } t=m,\dots,n-1\\
& x_{t+1}^2\le x_t^2+d_{t+1}+1(g_{t+1}-g_t<0), & \text{for } t=m,\dots,n-1\\
& x_{t+1}^3\le x_t^3+d_{t+1}+1(g_{t+1}-g_t<0), & \text{for } t=m,\dots,n-1\\
& -\xi_{w,i}\le\tilde{w}_{t_i}-w_{t_i}\le\xi_{w,i}, & \text{for } i=1,\dots,n_w\\
& -\xi_{u,i}\le\tilde{u}_{r_i}-u_{r_i}\le\xi_{u,i}, & \text{for } i=1,\dots,n_u\\
& -\xi_{z,t}\le z_t\le\xi_{z,t}, & \text{for } t=1,\dots,n
\end{aligned}
$$

- **Goal:** for a … ight and step data, estim… odel
- **Challenge:** … and noise
- **Approach:** … oisy/missing data -- cha…

Bilevel Optimization

$\mathbf{P}_{\text{mle}}$ min … $) + \sum_{t=1}^{n} -\log\psi_z(z_t)$

s.t. … $), (3)$ for $t = m, \dots, n.$

- Difficult sin…

$\mathbf{P}_{\text{mle-milp}}$ mi… $\sigma_3^{-1/2}\sum_{t=1}^{n}|z_t|$

s.t … $(7),$ for $t = m, \dots, n.$

Bertsimas et al. 2014

13

Mobile Diabetes Prevention Program (mDPP) Trial

- mDPP is mobile phone delivered weight loss intervention
- 5 month duration of program
- Control group
  - Pedometer only (31 adults)
  - Average weight change 0.3 ± 3.0 kg
- Treatment group
  - Pedometer + Mobile app + 7 On-site meetings (30 adults)
  - Average weight change –6.2 ± 5.9 kg
  - Step goals increase 20% weekly after baselining
  - Asked to input weight measurements into app

(Fukuoka, et al., 2011, 2015)

14

Model-Estimated Trajectories

- Representative example:

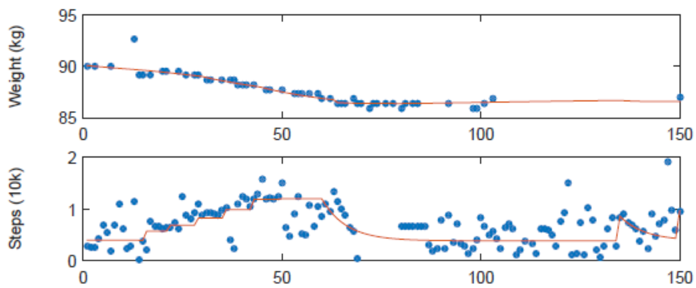

Bayesian Predictions

Participants arrive sequentially to the program

Challenge
Use all data from past participants and short span of data from new individual to make predictions of future states

- Solution approach
  - Use histograms of past participants' parameters (MLE) as prior
  - Approximate Bayesian posterior with profile likelihood

$$\psi(w_{t_f} \mid C, \tilde{W}, \tilde{U}) = \int \psi(W, U, F, \Theta \mid C, \tilde{W}, \tilde{U}) \cdot dW_{-t_f} \cdot dU \cdot dF \cdot d\Theta$$

$$\approx \max_{W_{-t_f}, U, F, \Theta} \psi(W, U, F, \Theta \mid C, \tilde{W}, \tilde{U})$$

weights, step

visits,goals

Posterior distrib of final weight

Difficult to compute

Severini 1999

MILP Formulation – Similar approach for MAP estimate of parameters

(Severini, 1999)

Bayesian Predictions

Challenge

Use all da … ta from new individua …

PROPOSITION 7. *The objective of* $\mathbf{P}_{\mathrm{pl}}$ *(after computing its negative logarithm) is*

$$\sigma_1^{-1/2}\sum_{i=1}^{n_w}|\tilde{w}_{t_i}-w_{t_i}|+\sigma_2^{-1/2}\sum_{i=1}^{n_u}|\tilde{u}_{\tau_i}-u_{\tau_i}|+\sigma_3^{-1/2}\sum_{t=1}^{n}|z_t|+$$
$$2^{-1/2}\sum_{X\in\{\mu,q,s_0,\beta_0,\delta_0\}}\sum_{i=1}^{m_x}\log\pi_i^x\cdot y_i^x+2^{-1/2}\sum_{X\in\{\beta,\delta\}}\sum_{k=0}^{n_d-1}\sum_{i=1}^{m_x}\sum_{j=1}^{\eta_x}\log\pi_{i,j}^x\cdot y_{i,j}^{x,k},$$

subject to constraints for one-dimensional histograms

$$\sum_{i=1}^{m_x}h_i^x\cdot y_i^x\le X\le\sum_{i=1}^{m_x}h_{i+1}^x\cdot y_i^x;\quad \sum_{i=1}^{m_x}y_i^x=1;\quad y_i^x\in\{0,1\},\ \forall i=1,\dots,m_x,$$

for all $X\in\{\mu,q,s_0,\beta_0,\delta_0\}$, *and constraints for conditional histograms*

$$\sum_{i=1}^{m_x}\sum_{j=1}^{\eta_x}h_{i,j}^x\cdot y_{i,j}^{x,k}\le X_{k+1}\le\sum_{i=1}^{m_x}\sum_{j=1}^{\eta_x}h_{i+1}^{x,k}\cdot y_{i,j}^x$$
$$\sum_{i=1}^{m_x}\sum_{j=1}^{\eta_x}\phi_{i,j}^x\cdot y_{i,j}^{x,k}\le X_k\le\sum_{i=1}^{m_x}\sum_{j=1}^{\eta_x}\phi_{i+1}^x\cdot y_{i,j}^{x,k}$$
$$y_{i,j}^{x,k}\in\{0,1\},\ \forall i=1,\dots,m_x,\ j=1,\dots,\eta_x;\quad \sum_{i=1}^{m_x}\sum_{j=1}^{\eta_x}y_{i,j}^{x,k}=1,$$

for all $X\in\{\beta,\delta\}$ *and* $k=0,\dots,n_d-1$.

- Solution …
  - Use h … LE) as prior
  - Appro … od

$\psi(w_{t_f}$ … $dU\cdot dF\cdot d\Theta$

Posterior distrib… of final weight

MILP Formulation – Similar approach for MAP estimate of parameters

Theorem (Consistency of Approximation)

Under mild conditions, our solution approach generates statistically consistent estimates of the Bayesian posterior distribution

(Severini, 1999)

Bayesian Predictions

Challenge

Use all da … ta from new individua …

$$\ell(w_{t_f}=\omega)=$$
$$\begin{aligned}
\min\ & \sigma_1^{-1/2}\sum_{i=1}^{n_w}\xi_{w,i}+\sigma_2^{-1/2}\sum_{i=1}^{n_u}\xi_{u,i}+\sigma_3^{-1/2}\sum_{t=1}^{n}\xi_{z,t}+\\
& 2^{-1/2}\sum_{X\in\{\mu,q,s_0,\beta_0,\delta_0\}}\sum_{i=1}^{m_x}\log\pi_i^x\cdot y_i^x+\\
& 2^{-1/2}\sum_{X\in\{\beta,\delta\}}\sum_{k=0}^{n_d-1}\sum_{i=1}^{m_x}\sum_{j=1}^{\eta_x}\log\pi_{i,j}^x\cdot y_{i,j}^{x,k}\\
\text{s.t.}\ & 2b(aw_t+bu_t+f_t+k)+2ru_t-q=0, && \text{for } t=1,\dots,m-1\\
& 2(aw_t+bu_t+f_t+k)+2rf_t-s_0=0, && \text{for } t=1,\dots,m-1\\
& 2b(aw_t+bu_t+f_t+k)+2ru_t-q-\lambda_t^2=0, && \text{for } t=m,\dots,n\\
& 2(aw_t+bu_t+f_t+k)+2rf_t-s_t=0, && \text{for } t=m,\dots,n\\
& g_t-\epsilon-(g_t-\epsilon)\cdot x_t^1\le u_t\le M+(g_t-\epsilon-M)\cdot x_t^1, && \text{for } t=m,\dots,n\\
& (g_t-\epsilon)\cdot x_t^2\le u_t\le M+(g_t+\epsilon-M)\cdot x_t^2, && \text{for } t=m,\dots,n\\
& (g_t+\epsilon)\cdot x_t^3\le u_t\le g_t+\epsilon+(M-g_t-\epsilon)\cdot x_t^3, && \text{for } t=m,\dots,n\\
& 0\le\lambda_t^2\le p_t, && \text{for } t=m,\dots,n\\
& p_t-M\cdot(1-x_t^1)\le\lambda_t^2\le M\cdot(1-x_t^3), && \text{for } t=m,\dots,n\\
& x_t^1,x_t^2,x_t^3\in\{0,1\}, && \text{for } t=m,\dots,n\\
& x_t^1+x_t^2+x_t^3=1, && \text{for } t=m,\dots,n\\
& w_{t+1}=a\cdot w_t+b\cdot u_t+f_t+k+z_t, && \text{for } t=1,\dots,n-1\\
& s_{t+1}=\gamma\cdot(s_t-s_0)+s_0-\beta_{t+1}\cdot d_{t+1}, && \text{for } t=m,\dots,n-1\\
& p_{t+1}\ge\gamma\cdot p_t+\delta_{t+1}\cdot d_{t+1}, && \text{for } t=m,\dots,n-1\\
& p_{t+1}\le\gamma\cdot p_t+\delta_{t+1}\cdot d_{t+1}+M\cdot(1-x_t^1), && \text{for } t=m,\dots,n-1\\
& p_{t+1}\ge\gamma\cdot p_t+\delta_{t+1}\cdot d_{t+1}+\mu-Mx_t^1, && \text{for } t=m,\dots,n-1\\
& p_{t+1}\le\gamma\cdot p_t+\delta_{t+1}\cdot d_{t+1}+\mu, && \text{for } t=m,\dots,n-1\\
& x_{t+1}^1\ge x_t^1-d_{t+1}-\mathbb{1}(g_{t+1}-g_t<0), && \text{for } t=m,\dots,n-1\\
& x_{t+1}^2\le x_t^2+d_{t+1}+\mathbb{1}(g_{t+1}-g_t<0), && \text{for } t=m,\dots,n-1\\
& x_{t+1}^3\le x_t^3+d_{t+1}+\mathbb{1}(g_{t+1}-g_t<0), && \text{for } t=m,\dots,n-1\\
& -\xi_{w,i}\le\tilde{w}_{t_i}-w_{t_i}\le\xi_{w,i}, && \text{for } i=1,\dots,n_w\\
& -\xi_{u,i}\le\tilde{u}_{\tau_i}-u_{\tau_i}\le\xi_{u,i}, && \text{for } i=1,\dots,n_u\\
& -\xi_{z,t}\le z_t\le\xi_{z,t}, && \text{for } t=1,\dots,n\\
& \sum_{i=1}^{m_x}h_i^x\cdot y_i^x\le X\le\sum_{i=1}^{m_x}h_{i+1}^x\cdot y_i^x, && \text{for } X\in\{\mu,q,s_0,\beta_0,\delta_0\}\\
& y_i^x\in\{0,1\},\ \forall i=1,\dots,m_x, && \text{for } X\in\{\mu,q,s_0,\beta_0,\delta_0\}\\
& \sum_{i=1}^{m_x}y_i^x=1, && \text{for } X\in\{\mu,q,s_0,\beta_0,\delta_0\}\\
& \sum_{i=1}^{m_x}\sum_{j=1}^{\eta_x}h_{i,j}^x\cdot y_{i,j}^{x,k}\le X_{k+1}\le\sum_{i=1}^{m_x}\sum_{j=1}^{\eta_x}h_{i+1}^{x,k}\cdot y_{i,j}^x,
\end{aligned}$$

- Solution …
  - Use h … LE) as prior
  - Appro … od

$\psi(w_{t_f}$ … $dU\cdot dF\cdot d\Theta$

Posterior distrib… of final weight

… ate of parameters

Theorem (…

Under mild … ates statistically consistent … bution

(Severini, 1999)

Mobile Diabetes Prevention Program (mDPP) Trial

- mDPP is mobile phone delivered weight loss intervention
- 5 month duration of program
- Control group
  - Pedometer only (31 adults)
  - Average weight change 0.3 ± 3.0 kg
- Treatment group
  - Pedometer + Mobile app + 7 On-site meetings (30 adults)
  - Average weight change –6.2 ± 5.9 kg
  - Step goals increase 20% weekly after baselining
  - Asked to input weight measurements into app

(Fukuoka, et al., 2011, 2015)

Computational Results of Bayesian Prediction

- Representative example:

Posterior Likelihood

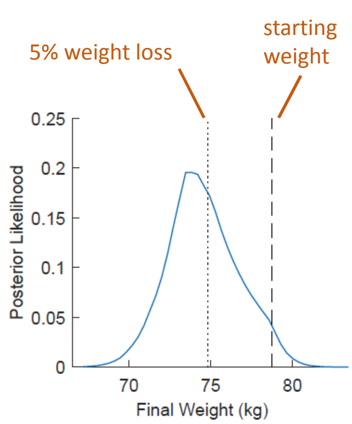

Maximum A Posteriori (MAP) Estimate

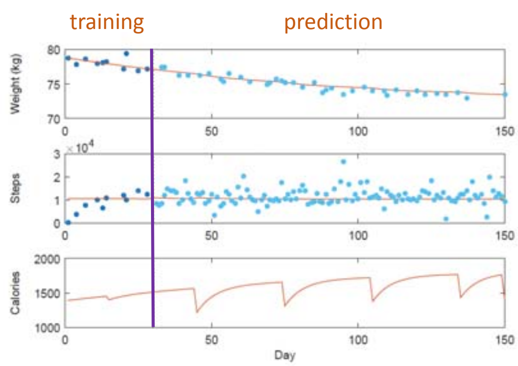

Model Validation

- Compared against standard machine learning algorithms

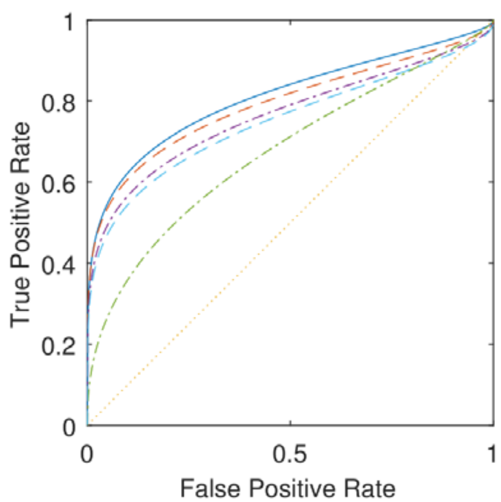

Legend:
Bayesian IOP
IOP
Linear SVM
Logistic Regression
Decision Tree

21

Policy Opt -- Optimal Policy for Single Agent

- Weight loss policy optimization – minimize expected posterior loss:

Challenging to calculate since no closed form for distrib…

noisy observations

$$\min\left\{\mathbb{E}\left[\ell(\{x_t,u_t\}_{t=T+1}^{T+n})\,\middle|\,\{\tilde{x}_{t_i}\}_{i=0}^{n_x},\{\tilde{y}_{\tau_i}\}_{i=0}^{n_u},\{\pi_i\}_{i=0}^{T}\right]\,\middle|\,\{\pi_i\}_{i=T+1}^{T+n}\in\Pi^n\right\}$$

weight   decisions   incentives

Class of functions   Bickel and Doksum 2006   **Value function of MIP**

$$\min\left\{\frac{\sum_{i=1}^{M}\varphi(x_{i,0},\theta_{i,0},\pi)\exp(\psi_T(x_{i,0},\theta_{i,0};\pi))}{\sum_{i=1}^{M}\exp(\psi_T(x_{i,0},\theta_{i,0};\pi))}\,\middle|\,\{\pi_i\}_{i=T+1}^{T+n}\in\Pi^n\right\}$$

Kaut and Wallace 2003

- Can be approximated as bi-level MILP, lower level MILP
  - Hard to solve, even with 1 constraint
  - Existing algorithms only solve small instances (Moore and Bard, 1992, 1990; DeNegre and Ralphs, 2009)

22

Two Stage Adaptive Algorithm (2SSA)

Algorithm 1 Two Stage Single Agent Algorithm (2SSA)

1: Input: $\{\bar{x}_{t_i}\}_{i=0}^{T_x}, \{\bar{u}_{r_i}\}_{i=0}^{T_u}, \pi^{init}$

2: Compute $(\hat{x}_{0,T}, \hat{\theta}_{0,T}) = \arg\max_{(x_0,\theta_0)} \psi_T(x_0, \theta_0; \pi^{init})$ — MAP estimate of "type"

3: Output $\pi_{2SSA}(T) = \arg\min_{\pi \in \Pi} \varphi(\hat{x}_{0,T}, \hat{\theta}_{0,T}, \pi)$ — Estimate of Policy given type

Can be formulated, solved as MILP

Two MILP's – asymptotically optimal

Theorem 1 *Note that* $\arg\min\{\varphi(x_0^*, \theta_0^*, \{\pi_i\}_{i=0}^{T+n}) | \{\pi_i\}_{i=T+1}^{T+n} \in \Pi^n\}$ *is the set of optimal solutions under the agent's true initial conditions* (x_0^*, θ_0^*)*. If some techinical assumptions hold, then we have that*

$$\text{dist}\Big(\pi_{2SSA}(T), \arg\min\{\varphi(x_0^*, \theta_0^*, \{\pi_i\}_{i=0}^{T+n}) \mid \{\pi_i\}_{i=T+1}^{T+n} \in \Pi^n\}\Big) \xrightarrow{p} 0$$

as $T \to \infty$*, for any* $\pi_{2SSA}(T)$ *returned by 2SSA.* $\text{dist}(x, B) = \inf_{y \in B} \|x - y\|$.

23

Multiple Agent Policy Computation

- Need to account for joint budget constraints
- Solve single patient problem for fixed budget $v \in \{0, \ldots, \beta\}$
- $y_v^a \in \mathbb{B}$: Equals 1 if a policy with budget *v* is chosen for patient *a*

$$\begin{aligned} y \in \arg\min \ & \textstyle\sum_{a \in \mathcal{A}} \sum_{v \in V} \phi_v^a \cdot y_v^a \\ \text{s.t.} \ & \textstyle\sum_{a \in \mathcal{A}} \sum_{v \in V} v \cdot y_v^a \leq \beta \\ & \textstyle\sum_{v \in V} y_v^a = 1 \text{ for } a \in \mathcal{A} \\ & y_v^a \in \{0, 1\} \end{aligned}$$

24

Overall Approach

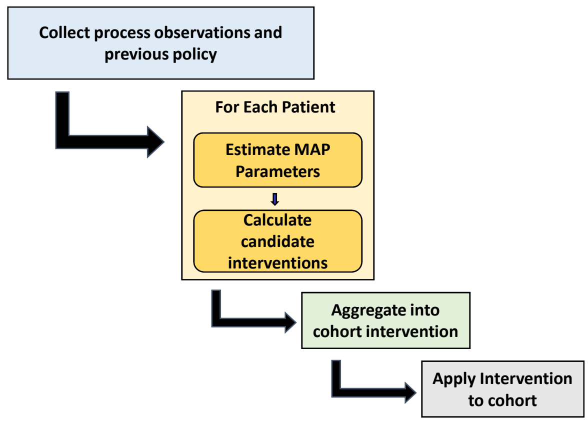

25

Steps 2/3: ABMA Algorithm

Algorithm 2 Adaptive Behavioral Multi-Agent Algorithm (ABMA)

Require: $\{\tilde{x}^a_{t_i}\}_{i=0}^{n^a_x}, \{\tilde{u}^a_{\tau_i}\}_{i=0}^{n^a_u}, \{\pi^a_i\}_{i=0}^{T}$ for $a \in \mathcal{A}$

for all $a \in \mathcal{A}$ **do** — For each patient

  compute $(\hat{x}^a_{0,T}, \hat{\theta}^a_{0,T}) = \arg\max_{(x_0,\theta_0)} \psi_T(x_0, \theta_0)$

  for all $v \in V$ **do**

    set $\pi^a_v \in \arg\min\{\varphi^a(\hat{x}_{0,T}, \hat{\theta}_{0,T}, \{\pi_i\}_{i=0}^{T+n}) \mid \{\pi_i\}_{i=T+1}^{T+n} \in S_v\}$

    set $\phi^a_v = \varphi^a(\hat{x}^a_{0,T}, \hat{\theta}^a_{0,T}, \pi^a_v)$

  end for

end for

compute:

$$\begin{aligned} y \in \arg\min\ & \textstyle\sum_{a\in\mathcal{A}}\sum_{v\in V} \phi^a_v \cdot y^a_v \\ \text{s.t. } & \textstyle\sum_{a\in\mathcal{A}}\sum_{v\in V} v \cdot y^a_v \le \beta \\ & \textstyle\sum_{v\in V} y^a_v = 1 \text{ for } a \in \mathcal{A} \\ & y^a_v \in \{0,1\} \end{aligned}$$

Knapsack

for all $a \in \mathcal{A}$ and $v \in V$ **do**

  set $\pi^a_{ABMA}(T) = \pi^a_v$ **if** $y^a_v = 1$

end for

return $\pi^a_{ABMA}(T)$ for $a \in \mathcal{A}$

26

Steps 2/3: (Asymptotically Optimal) ABMA Algorithm

Algorithm 2 Adaptive Behavioral Multi-Agent Algorithm (ABMA)

Require: $\{\tilde{x}^a_{t_i}\}_{i=0}^{n_x^a}, \{\tilde{u}^a_{\tau_i}\}_{i=0}^{n_u^a}, \{\pi^a_i\}_{i=0}^{T}$ for $a \in \mathcal{A}$

for all $a \in \mathcal{A}$ **do**
  compute $(\hat{x}^a_{0,T}, \hat{\theta}^a_{0,T}) = \arg\max_{(x_0,\theta_0)} \psi_T(x_0, \theta_0)$
  for all $v \in V$ **do**
    set $\pi^a_v \in \arg\min\{\varphi^a(\hat{x}_{0,T}, \hat{\theta}_{0,T}, \{\pi_i\}_{i=0}^{T+n}) \mid \{\pi_i\}_{i=T+1}^{T+n} \in S_v\}$
    set $\phi^a_v = \varphi^a(\hat{x}^a_{0,T}, \hat{\theta}^a_{0,T}, \pi^a_v)$
  end for
end for

THEOREM 2. *Note that* $\arg\min\{\Phi(x_0^{*,a}, \theta_0^{*,a}, \{\pi_i^{*,a}\}_{i=0}^{T+n}$ *for* $a \in \mathcal{A}) \mid \{\{\pi_i^a\}_{i=T+1}^{T+n}$ *for* $a \in \mathcal{A}\} \in \Omega\}$ *is the set of optimal solutions under the agents' true initial conditions* $(x_0^{*,a}, \theta_0^{*,a})$*. If Assumptions 1–8 hold, then we have that*

$$\operatorname{dist}\Big(\{\pi^a_{ABMA}(T) \text{ for } a \in \mathcal{A}\}, \arg\min\{\Phi(x_0^{*,a}, \theta_0^{*,a}, \{\pi_i^a\}_{i=0}^{T+n} \text{ for } a \in \mathcal{A}) \mid \{\{\pi_i^a\}_{i=T+1}^{T+n} \text{ for } a \in \mathcal{A}\} \in \Omega\}\Big) \xrightarrow{p} 0 \quad (32)$$

as $T \to \infty$*, for any* $\pi^a_{ABMA}(T)$ *returned by ABMA. Recall that* $\operatorname{dist}(x, B) = \inf_{y \in B} \|x - y\|$.

return $\pi^a_{ABMA}(T)$ for $a \in \mathcal{A}$

Mobile Diabetes Prevention Program (mDPP) Trial

- mDPP is mobile phone delivered weight loss intervention
- 5 month duration of program
- Control group
  - Pedometer only (31 adults)
  - Average weight change 0.3 ± 3.0 kg
- Treatment group
  - Pedometer + Mobile app + 7 On-site meetings (30 adults)
  - Average weight change –6.2 ± 5.9 kg
  - Step goals increase 20% weekly after baselining
  - Asked to input weight measurements into app

(Fukuoka, et al., 2011, 2015)

Simulation of ABMA Algorithm

- Simulate 150 day clinical trial
- Three loss functions: step, hinge, time-varying hinge → step

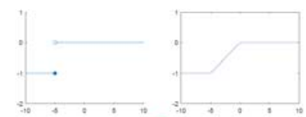

- Compare algorithm to simple heuristics:
  - Only one visit and no exercise goals
  - Adaptive: Set goals as 10% increase of moving average, schedule visits at end for patients closest to meeting weight loss goal
  - Original mDPP trial: seven visits, goals increased 20% each week starting with average steps during "no goal" period
- Recalculate the intervention design at the beginning of each month of the trial
- Using behavioral model with estimated individual parameters from trial (MLE)

29

Simulation of ABMA Algorithm

- Weight loss interventions
  - Daily quantitative exercise goals
  - Clinical counseling (constrained or budgeted)
  - Recalculated monthly
- Behavioral Analytics (ABMA) for personalizing treatment

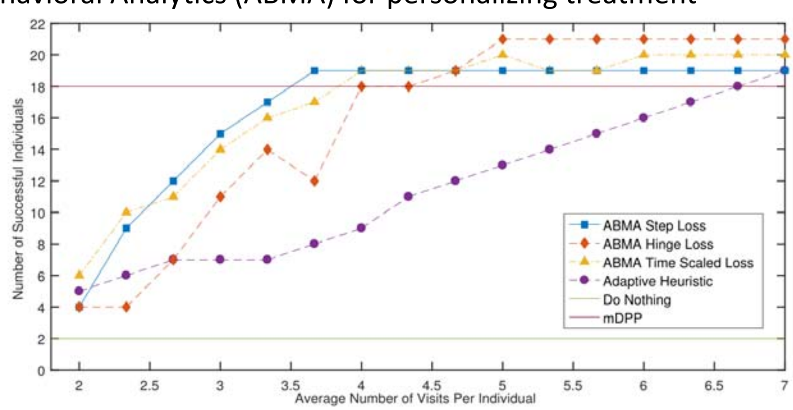

30

Simulation of ABMA Algorithm

- For unsuccessful individuals…

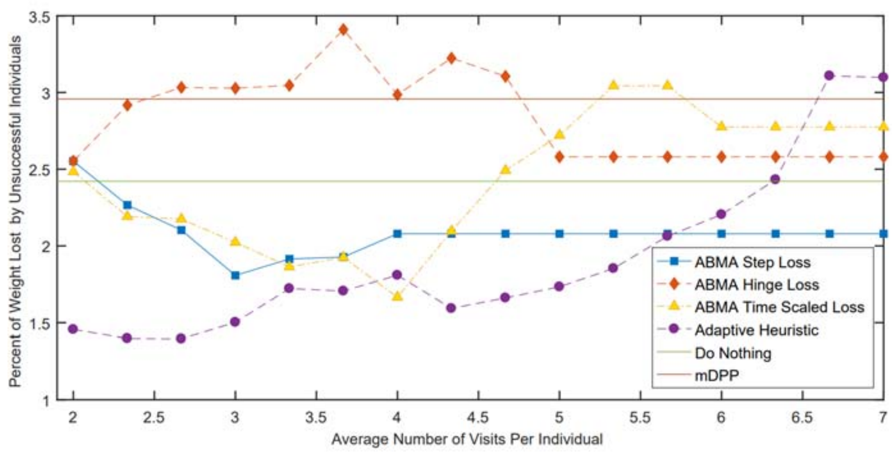

mSTAR Study

- Randomized pilot study with UC Berkeley students (64)
- Just focusing on steps…
- 10 weeks
- Control group
  - Received constant 10,000 steps/day goals
- Treatment group
  - Received personalized step goals

(a) Splash Screen (b) Home Tab (c) History Tab (d) Contact Tab

mSTAR Study Results

- Control had decrease of 1520 steps/day after 10 weeks
- Treatment had increase of 700 steps/day after 10 weeks
- Treatment did 2200 steps/day (1 mile/day) more walking ($p = 0.039$)

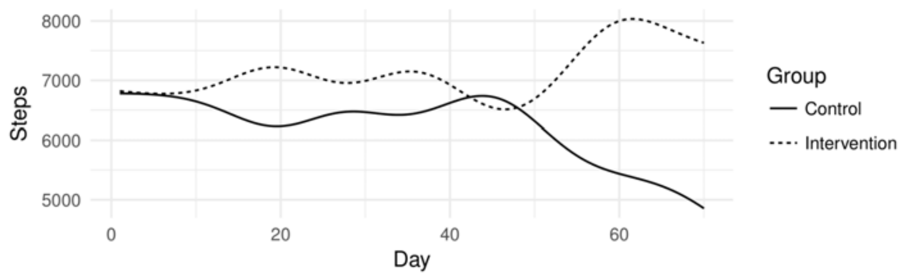

- Similar results in our CalFitness (larger randomized controlled) trial

33

Summary

Step I: Model of a Single Agent
- Myopic agents
- Deterministic dynamics
- Bounded states and inputs
- Concave utility functions

Step II: Estimate Parameters
- Conditions for MILP for MLE
- Conditions for MAP estimation using set of MILPs – **consistent estimates**

Step III: Incentive Optimization
- Optimal incentive design for a single agent using MAP estimate of parameters to minimize loss function
- MILP to estimate incentives for single agent - **Asymptotically Optimal as data set grows**
- For multiple agents:
  - Approach built around decomposing problems per agent and per incentive level
  - Knapsack problem to determine optimal incentive set
  - **Asymptotically Optimal as data set grows**

34

Outline

- Behavioral Analytics Framework
- Potential Supply Chain Applications
- Behavioral Analytics for Myopic Agents
  - Healthcare application
    - Behavioral analytics algorithm
      - Paper: Behavioral modeling in weight loss interventions
      - Paper: Behavioral analytics for Myopic Agents
    - Experimental results
      - Paper: Personalizing Mobile Fitness Apps using Reinforcement Learning
  - Some Generalizations
- Extensions Needed for Supply Chain Applications

In A Supply Chain…

Decentralized Supply Chain With Sales Effort

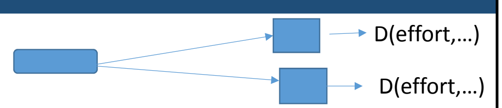

- **Demand** is a (known, stationary, stochastic) function of **costly effort** (with known effort cost function)
- **Goal**: equilibrium prices, **coordinating contracts** (rebates, returns, etc)
- **Extensions:** free rider on effort in omni-channel supply chain, endogenous pricing, alternative contract structures, partially observable/verifiable effort

KNOWN MOTIVATIONAL DYNAMICS?

MYOPIC?

- **How can a coordinator use new, noisy data streams to maximize incentive allocation strategy?**
  - How can effort be estimated?
  - What about other kinds of incentives? Training, advertising, events, agent sales efforts, etc.?
  - What about more complex utility functions?

SIZE?

Franchise Logistics

KNOWN SYSTEM DYNAMICS?

Key: New (Noisy) Data Streams to Maximize Incentive Allocation Efficiency

- Both the coordinator and the agents have more complex objectives/utility functions.
- The coordinator has many different options of data to monitor
- Effort is more challenging to estimate, is likely to evolve over time, and can be a function of many different incentives; literature highlights expensive data acquisition, difficult to estimate utilities

Chu and Desai(1995), Lariviere and Padmanabhan(97), Netessine and Rudi (2000), Taylor(2002), Krishnan at al (2004) Mukhopadhyay at al. (2008), Cachon (2003), Xing and Liu (2012), Gal-Or (1995), Li (1997), Xie et al. 2016

In A Supply Chain…

Sales Force Compensation

KNOWN MOTIVATIONAL DYNAMICS

- Firm maximizes profit (some perhaps stochastic function of effort, price)
- Salespeople maximize utility (known function of incentives, effort)
- Effort observed, functions known SIZE? KNOWN SYSTEM DYNAMICS?

MYOPIC?

- Effort not observed, firm maximizes expected profit
- Demand stochastic
- Distributional assumptions on utility, estimation of utility, relationship to demand

- Commercial software vendors sell software that using ML techniques to try to determine relationship between incentive efforts, sales (but not in an interpretable way)

Key: New (Noisy) Data Streams to Maximize Incentive Allocation Efficiency

- **Question: Can we model utility, optimize incentives?**

Coughlan (1993) 37

Limitations/Opportunities

To effectively model supply chain problems, we need to relax the following:

- System dynamics known Some ideas – theoretically challenging
- Motivational state dynamics known – agent "type" evolves deterministically
- Myopic agents Some ideas – computationally challenging
- Perhaps as problems get bigger, we'll need to approximate MIPS

38

Thank you