Single Agent Dynamics: Dynamic Discrete Choice Models

Part II: Estimation

Günter J. Hitsch
The University of Chicago Booth School of Business

2013

Overview

Overview

Data

- We observe the choice behavior of N individuals
- For each individual i, define a vector of states and choices for periods $t = 0, \ldots, T_i : Q_i = (x_{it}, a_{it})_{i=1}^{T_i}$
- The full data vector is $Q = (Q_1, \ldots, Q_N)$
- We will initially assume that all individuals are identical and that the components of the state x are observed to us
  - We will discuss the case of (permanent) heterogeneity and unobserved state variables later

Data-generating process

- We assume that the choice data are generated based on the choice-specific value functions

$$v_j(x) = u_j(x) + \beta \int w(x') f(x'|x, j)\, dx'$$

- We observe action a_{it} conditional on the state x_{it} if and only if

$$v_k(x_{it}) + \epsilon_{kit} \geq v_j(x_{it}) + \epsilon_{jit} \quad \text{for all } j \in \mathcal{A},\ j \neq k$$

- The CCP $\sigma_k(x_{it})$ is the probability that the inequality above holds, given the distribution of the latent utility components $g(\epsilon)$

Static vs dynamic discrete choice models

- Using the concept of choice-specific value functions, the predictions of a dynamic discrete choice model can be expressed in the same manner as the predictions of a static discrete choice model
- Therefore, it appears that we can simply estimate each $v_j(x)$ by approximating it using a flexible functional form ψ, e.g. a polynomial or a linear combination of basis functions:

$$v_j(x) \approx \psi_j(x; \theta)$$

- In a static discrete choice model, we typically start directly with a parametric specification of each choice-specific utility function, $u_j(x; \theta)$
- Unlike $u_j(x)$, $v_j(x)$ is not a structural object, but the solution of a dynamic decision process that depends on the model primitives, $u_j(x)$, $f(x'|x, a)$, β, and $g(\epsilon)$.
- Why is this important? — Because it affects the questions that the estimated model can answer

Example: The storable goods demand problem

- Remember that $x_t \equiv (i_t, P_t)$
  - $i_t \in \{0, 1, \ldots, I\}$, and $P_t \in \{P^{(1)}, \ldots, P^{(L)}\}$
- The random utility components are Type I Extreme Value, hence the discrete values $v_j(x)$ can be estimated just as in a standard multinomial logit model
- Suppose we estimate $v_j(x)$ based on data generated from the high/low promotion process discussed in the example in Part I
- We want to evaluate a policy where the price is permanently set at the low price level
- Will knowledge of $v_j(x)$ allow us to answer this question?

CCP's — Pricing with promotions

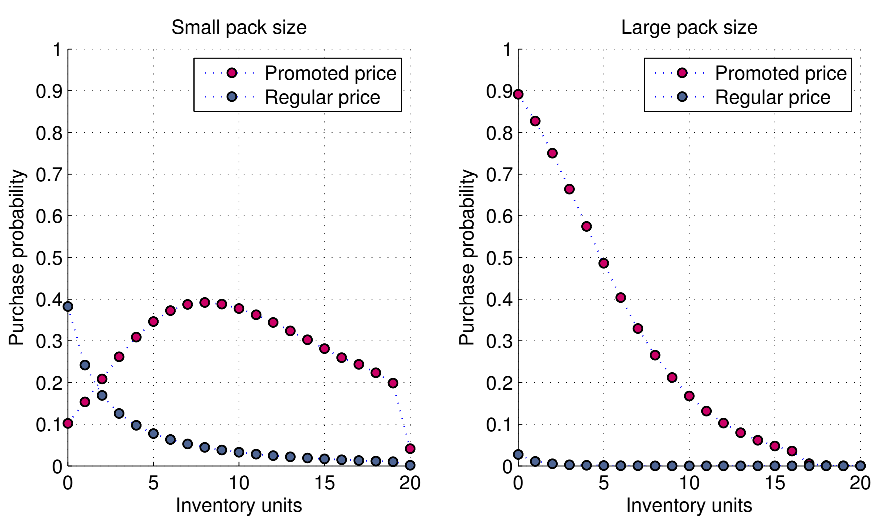

CCP's — only low price

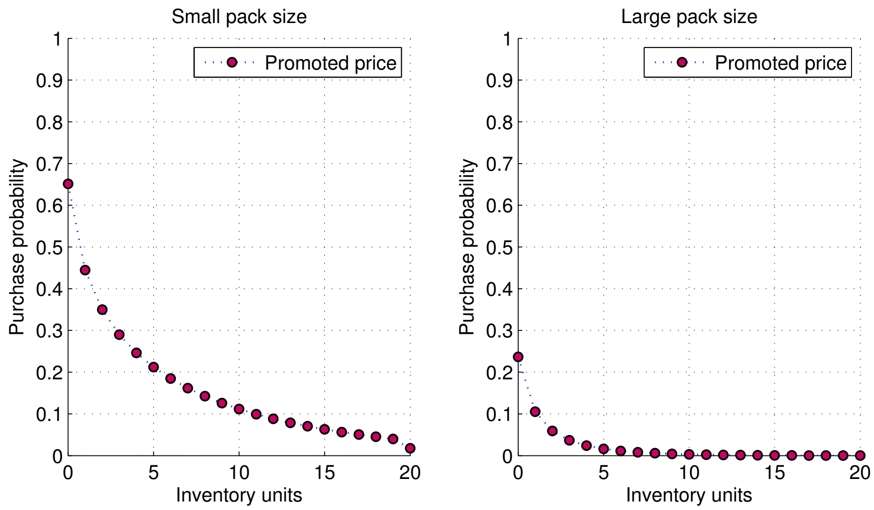

Example: The storable goods demand problem

- Based on the data generated under the price promotion process, the sales volume ratio between promoted and regular price periods is $3.782/0.358 = 10.6$
- However, when we permanently lower the price the sales ratio is only $0.991/0.358 = 2.8$

Example: The storable goods demand problem

- Evaluation of a policy where the price is permanently set at the low price level:
  - We are not just changing a component of x (the price), but also the price expectations, $f(x'|x, a)$
  - Hence, $v_j(x)$ will also change to $\tilde{v}_j(x)$
  - Unless consumer choices under the alternative price process are observed in our data, we cannot estimate $\tilde{v}_j(x)$
  - Instead, we must predict $\tilde{v}_j(x)$ using knowledge of the model primitives, $u_j(x)$, $f(x'|x, a)$, β, and $g(\epsilon)$
- The same problem does not arise in a static discrete choice model (if the static discrete choice model accurately describes consumer behavior)

Structural and reduced form of the model

- $(u_j(x), f(x'|x, a), \beta, g(\epsilon))$ is the structural form of the dynamic discrete choice model
- $(v_j(x), g(\epsilon))$ is the reduced form of the model
- The reduced form describes the joint distribution of the data, but typically cannot predict the causal effect of a marketing policy intervention
- Policy predictions are therefore not only dependent on the statistical properties of the model and model parameters, but also on the behavioral assumptions we make about how decisions are made
- Good background reading on structural estimation and the concept of inferring causality from data that are non-experimentally generated: Reiss and Wolak (2007)

Overview

Identification

- All behavioral implications of the model are given by the collection of CCP's $\{\sigma_j(x) : j \in \mathcal{A}, x \in \mathbb{X}\}$
- Suppose we have "infinitely many" data points, so that we can observe the CCP's $\{\sigma_j(x)\}$
  - For example, if $\mathbb{X}$ is finite we can estimate $\sigma_j(x)$ as the frequency of observing j conditional on x
- We assume we know the distribution of the random utility components, $g(\epsilon)$
- Could we then uniquely infer the model primitives describing behavior from the data:

$$u_j(x),\ f(x'|x,j),\ \beta\ ?$$

- I.e., is the structural form of the model identified (in a non-parametric sense)?

Identification

- Hotz and Miller (1993): If ϵ has a density with respect to the Lebesgue measure on $\mathbb{R}^{K+1}$ and is strictly positive, then we can invert the observed CCP's to infer the choice-specific value function differences:

$$v_j(x) - v_0(x) = \Psi_j^{-1}(\sigma(x)) \quad \text{for all } j \in \mathcal{A}$$

- If ϵ_j is Type I Extreme Value distributed, this inversion has a closed form:

$$v_j(x) - v_0(x) = \log(\sigma_j(x)) - \log(\sigma_0(x))$$

- We see that $\beta = 0$, $u_0(x) \equiv 0$, and

$$u_j(x) \equiv \log(\sigma_j(x)) - \log(\sigma_0(x))$$

completely rationalize the data!

Identification

Theorem

Suppose we know the distribution of the random utility components, $g(\epsilon)$, the consumers' beliefs about the evolution of the state vector, $f(x'|x,j)$, and the discount factor β. Assume that $u_0(x) \equiv 0$. Let the CCP's, $\sigma_j(x)$, be given for all x and $j \in \mathcal{A}$. Then:

(i) We can infer the unique choice-specific value functions, $v_j(x)$, consistent with the consumer decision model.

(ii) The utilities $u_j(x)$ are identified for all states x and choices j.

(Non)identification proof

- Define the difference in choice-specific value functions:

$$\tilde{v}_j(x) \equiv v_j(x) - v_0(x)$$

- Hotz-Miller (1993) inversion theorem:

$$\tilde{v}_j(x) = \Psi_j^{-1}(\sigma(x)) \quad \text{for all } j \in \mathcal{A}$$

- The expected value function can be expressed as a function of the data and the reference alternative:

$$\begin{aligned} w(x) &= \int \max_{k\in\mathcal{A}}\{v_k(x) + \epsilon_k\} g(\epsilon) d\epsilon \\ &= \int \max_{k\in\mathcal{A}}\{\tilde{v}_k(x) + \epsilon_k\} g(\epsilon) d\epsilon + v_0(x) \end{aligned}$$

(Non)identification proof

- Choice-specific value of reference alternative:

$$\begin{aligned} v_0(x) &= u_0(x) + \beta \int w(x') f(x'|x,0) dx' \\ &= \beta \int \max_{k\in\mathcal{A}}\{\tilde{v}_k(x') + \epsilon_k\} g(\epsilon) f(x'|x,0) d\epsilon dx' \\ &\quad + \beta \int v_0(x') f(x'|x,0) dx' \end{aligned}$$

  - Defines a contraction mapping and hence has a unique solution

- Recover all choice-specific value functions:

$$v_j(x) = \Psi_j(\sigma(x)) + v_0(x)$$

- Recover the utility functions:

$$u_j(x) = v_j(x) - \beta \int w(x') f(x'|x,j) dx$$

(Non)identification

- The proof of the proposition shows how to calculate the utilities, $u_j(x)$, from the data and knowledge of β and $f(x'|x,j)$
- Proof shows that if $\beta' \neq \beta$ or $f'(x'|x,j) \neq f(x'|x,j) \Rightarrow$ $u'_j(x) \neq u_j(x)$ in general
- Implications
  - If either the discount factor or the consumer's belief about the state evolution is unknown, the utility function is not identified
  - I.e., the model primitives $u_j(x),\ \beta,\ f(x'|x,j)$ are non-parametrically unidentified

Implications of (non)identification result

- In practice: Researchers assume a given discount factor calibrated from some "overall" interest rate r, such that $\delta = 1/(1+r)$

$$u'(c_t) = \delta \mathbb{E}_t \left[(1+r)u'(c_{t+1})\right]$$

  - Typically, discount factor corresponding to 5%-10% interest rate is used
- Assume rational expectations: the consumer's subjective belief $f(x'|x,j)$ coincides with the actual transition process of x_t
  - Allows us to estimate $f(x'|x,j)$ from the data

Identifying β based on exclusion restrictions

- A "folk theorem": β is identified if there are states that do not affect the current utility but the transition probability of x
- More formally: Suppose there are two states x_1 and x_2 such that $u_j(x_1) = u_j(x_2)$ but $f(x'|x_1, j) \neq f(x'|x_2, j)$
- Intuition: Variation in x does not change the current utility but the future expected value, and thus β is identified:

$$v_j(x) = u_j(x) + \beta \int w(x') f(x'|x, j) \, dx'$$

Identifying β based on exclusion restrictions

- The "folk theorem" above is usually attributed to Magnac and Thesmar (2002), but the actual statement provided in their paper is more complicated
- Define the *current value function*, the expected difference between (i) choosing action j today, action 0 tomorrow, and then behaving optimally afterwards, and (ii) choosing action 0 today and tomorrow and behaving optimally afterwards

$$\mathcal{U}_j(x) \equiv \left(u_j(x) + \beta \int v_0(x') p(x'|x, j) dx' \right) - \left(u_0(x) + \beta \int v_0(x') p(x'|x, 0) dx' \right)$$

- Also, define

$$\mathcal{V}_j(x) \equiv \int \left(w(x') - v_0(x') \right) p(x'|x, j) dx'$$

Identifying β based on exclusion restrictions

- The proposition proved in Magnac and Thesmar (2002):

Theorem

Suppose there are states x_1 *and* x_2, $x_1 \neq x_2$, *such that* $\mathcal{U}_j(x_1) = \mathcal{U}_j(x_2)$ *for some action* j. *Furthermore, suppose that*

$$(\mathcal{V}_j(x_1) - \mathcal{V}_0(x_1)) - (\mathcal{V}_j(x_2) + \mathcal{V}_0(x_2)) \neq 0.$$

Then the discount factor β *is identified.*

- Note: The assumptions of the theorem require knowledge of the solution of the decision process and are thus difficult to verify

Identifying β based on exclusion restrictions

- Is the "folk theorem" true?
- Although widely credited to Magnac and Thesmar (2002), I believe this attribution is false
- However, a recent paper by Fang and Wang (2013), "Estimating Dynamic Discrete Choice Models with Hyperbolic Discounting, with an Application to Mammography Decisions," seems to prove the claim in the "folk theorem"

Identification of β from data on static and dynamic decisions

- Yao et al. (2012)
- Observe consumer choice data across two scenarios
  - Static: Current choice does not affect future payoffs
  - Dynamic
- Examples:
  - Cell phone customers are initially on a linear usage plan, and are then switched to a three-part tariff
    - Under the three-part tariff current cell phone usage affects future per-minute rate
  - Any finite horizon problem
- Yao et al. prove identification of β for continuous controls
  - No proof provided for discrete choices, but (my guess) statement is true more generally

Identification of β using stated choice data

- Dubé, Hitsch, and Jindal (2013)
- Conjoint design to infer product adoption choices
  - Present subjects with forecasts of future states (e.g. prices)
  - Collect stated choice data on adoption timing
- Identification assumption:
  - Subjects take the forecast of future states as given
- Allows identification of discount factor β, or more generally a discount function $\rho(t)$
- Intuition
  - Treatments: Manipulations of states that change current period utilities by the same amount in period $t = 0$ and $t > 0$
  - Relative effect on choice probabilities and hence choice-specific value differences at $t > 0$ versus $t = 0$ identifies $\rho(t)$

Overview

Goal of estimation

- We would like to recover the structural form of the dynamic discrete choice model, $(u_j(x), f(x'|x,a), \beta, g(\epsilon))$
- We assume:
  - $g(\epsilon)$ is known
  - The decision makers have rational expectations, and thus $f(x'|x,a)$ is the true transition density of the data
  - Typically we also assume that β is "known"
- Assume that the utility functions and transition densities are parametric functions indexed by θ : $u_j(x;\theta)$ and $f(x'|x,a;\theta)$
- Goal: Develop an estimator for θ

Likelihood function

- The likelihood contribution of individual i is

$$l_i(Q_i;\theta) = \left(\prod_{t=0}^{T_i} \Pr\{a_{it}|x_{it};\theta\} \cdot f(x_{it}|x_{i,t-1},a_{i,t-1};\theta)\right) \times \dots$$
$$\dots \times \Pr\{a_{i0}|x_{i0};\theta\} \cdot \phi(x_{i0};\theta)$$

- The likelihood function is

$$l(Q;\theta) = \prod_{i=1}^{N} l_i(Q_i;\theta)$$

- Define the maximum likelihood estimator,

$$\theta^{NFP} = \arg\max_{\theta\in\Theta}\{\log(l(Q;\theta))\}$$

The nested fixed point estimator

- θ^{NFP} is a *nested fixed point estimator*
- We employ a maximization algorithm that searches over possible values of θ
  - Given θ, we first solve for $w(x;\theta)$ as the fixed point of the integrated Bellman equation
  - Given $w(x;\theta)$, we calculate the choice-specific value functions and then the CCP's $\sigma_j(x;\theta)$
  - Allows us to assemble $l(Q;\theta)$
- The solution of the fixed point $w(x;\theta)$ is nested in the maximization algorithm
- This estimator is computationally intensive, as we need to solve for the expected value function at each θ!

Estimating θ in two steps

- Separate $\theta = (\theta_u, \theta_f)$ into components that affect the utility functions, $u_j(x;\theta_u)$, and the transition densities $f(x'|x,a;\theta_f)$
- Note the log-likelihood contribution of individual i:

$$\log(l_i(Q_i;\theta)) = \sum_{t=1}^{T_i} \log(\Pr\{a_{it}|x_{it};\theta_u,\theta_f\}) + \dots$$
$$\sum_{t=2}^{T_i} \log(f(x_{it}|x_{i,t-1},a_{i,t-1};\theta_f))$$

- This expression suggests that we can estimate θ in two steps:
  1. Find a consistent estimator for θ_f (need not be a ML estimator)
  2. Conditional on $\hat{\theta}_f$, maximize the sum in the first row above to find $\hat{\theta}_u$ (adjust standard errors to account for sampling error in first step)

Overview

Bayesian learning models in marketing

- Large literature on demand for experience goods
  - Following Erdem and Keane (1996)
  - Product (or service) needs to be consumed or used to fully ascertain its utility
  - Examples
    - Product with unknown flavor or texture
    - Pharmaceutical drug with unknown match value, e.g. effectiveness or side effects
- Importance: Consumer learning may cause inertia in brand choices
  - Inertia = state dependence in a purely statistical sense
  - If true, has implications for pricing and other marketing actions
- Optimal sequential learning about demand
  - Hitsch (2006)

Bayesian learning: General model structure

- Learning about an unknown parameter vector $\vartheta \in \mathbb{R}^N$
- Information acquisition: Sampling (choosing) $j \in \mathcal{A}$ yields a signal $\xi_j \sim f_j(\cdot|\vartheta)$
- Knowledge: Prior $\pi_t(\vartheta)$
  - Knowledge about ϑ before any additional information in period t is sampled
- Learning through Bayesian updating:

$$\pi_{t+1}(\vartheta) \equiv \pi_t(\vartheta|\xi_{jt}) \propto f_j(\xi_{jt}|\vartheta) \cdot \pi_t(\vartheta)$$

  - Posterior at end of period t is prior at the beginning of period $t+1$

General model structure

- Notation: $\pi_t \equiv \pi_t(\vartheta)$
- State vector $x_t = (\pi_t, z_t)$
  - Ignore for now that π_t is infinite-dimensional in general
- Utility:

$$u_j(x_t) = \mathbb{E}(\mathfrak{u}_j(z_t, \xi_{jt}, \vartheta)|x_t) = \int \mathfrak{u}_j(z_t, \xi, \vartheta) f_j(\xi|\vartheta)\pi_t(\vartheta)\, d\xi d\vartheta$$

  - Expected utility given z_t and the agents' belief $\pi_t(\vartheta)$ about ϑ

General model structure

- We assume the decision maker is able to anticipate how her knowledge evolves, conditional on the potential information that she may receive in this period
- Allows to define a corresponding Markov transition probability

$$f(\pi_{t+1}|\pi_t, a_t)$$

- $\Rightarrow$ learning model is a special case of the dynamic discrete choice framework

Review: Normal linear regression model with conjugate priors

- Sample $\boldsymbol{y} \sim N(\boldsymbol{X}\boldsymbol{\beta}, \boldsymbol{\Omega})$, $\boldsymbol{y} \in \mathbb{R}^m, \boldsymbol{\beta} \in \mathbb{R}^k$
- Suppose $\boldsymbol{\Omega}$ is known
  - Can be generalized with inverse-Wishart prior on $\boldsymbol{\Omega}$, but rarely (never?) used in extant consumer learning literature
- Goal: Inference about $\boldsymbol{\beta}$
- Prior: $\boldsymbol{\beta} \sim N(\boldsymbol{\beta}_0, \boldsymbol{\Sigma}_0)$,
- Then the posterior is also normal:

$$
\begin{aligned}
p(\boldsymbol{\beta}|\boldsymbol{y}, \boldsymbol{X}, \boldsymbol{\Omega}) &= N(\boldsymbol{\beta}_n, \boldsymbol{\Sigma}_n) \\
\boldsymbol{\Sigma}_n &= (\boldsymbol{\Sigma}_0^{-1} + \boldsymbol{X}^T\boldsymbol{\Omega}^{-1}\boldsymbol{X})^{-1} \\
\boldsymbol{\beta}_n &= (\boldsymbol{\Sigma}_0^{-1} + \boldsymbol{X}^T\boldsymbol{\Omega}^{-1}\boldsymbol{X})^{-1}(\boldsymbol{\Sigma}_0^{-1}\boldsymbol{\beta}_0 + \boldsymbol{X}^T\boldsymbol{\Omega}^{-1}\boldsymbol{y})
\end{aligned}
$$

The Erdem and Keane (1996) learning model

- Consumers make purchase (= consumption) decisions over time, $t = 0, 1, \ldots$
- ϑ_j is the mean attribute level (quality) of product j
  - Affects utility
  - Consumers face uncertainty over ϑ_j
- Realized utility from consumption in period t affected by realized attribute level:

$$\xi_{jt} = \vartheta_j + \nu_{jt}, \quad \nu_{jt} \sim N(0, \sigma_\nu^2)$$

  - Inherent variability in attribute level
  - Variability in consumer's perception of attribute level
- Utility:

$$\mathfrak{u}_j(z_t, \xi_{jt}) = \gamma\left(\xi_{jt} - r\xi_{jt}^2\right) - \alpha P_{jt}$$

  - $r > 0$: risk aversion

Information sources

- Assume for notational simplicity that there is only one product with unknown attribute level, and hence we can drop the j index
- Learning from consumption:

$$\xi_t = \vartheta + \nu_t, \quad \nu_t \sim N(0, \sigma_\nu^2)$$

- Let $t_1, t_2, \ldots$ denote the time periods when a consumer receives a consumption signal
- Let $\mathcal{H}_t = \{t_k : t_k < t\}$ be the time periods prior to t when a consumption signal was received, and $N_t = |\mathcal{H}_t|$ be the corresponding number of signals
- Mean consumption signal prior to t :

$$\bar{\xi}_t = \frac{1}{N_t} \sum_{\tau \in \mathcal{H}_t} \xi_\tau$$

Prior

- Prior in the first decision period: $\pi_0(\vartheta) = N(\mu_0, \sigma_0^2) = N(\bar{\vartheta}, \sigma_0^2)$
  - $\bar{\vartheta}$ is the "product class mean attribute level"

- Refer back to the slide on the normal linear regression model with conjugate priors, and define:

$$\begin{aligned}
\boldsymbol{\beta} &= \vartheta \\
\boldsymbol{X} &= (1, \ldots, 1) \qquad \text{(vector of length } N_t) \\
\boldsymbol{\Omega} &= \text{diagonal matrix with elements } \sigma_\nu^2 \\
\boldsymbol{\beta}_0 &= \bar{\vartheta} \\
\boldsymbol{\Sigma}_0 &= \sigma_0^2
\end{aligned}$$

  - Elements in $\boldsymbol{\Omega}$ correspond to the consumption and advertising signals in $\mathcal{H}_t$ (order does not matter)

Posterior

- Results on normal linear regression model show that prior $\pi_t =$ posterior given the signals ξ_τ for $\tau < t$ is also normal:

$$\pi_t(\vartheta) = N(\mu_t, \sigma_t^2)$$
$$\sigma_t^2 = \mathbf{\Sigma}_n = \left(\frac{1}{\sigma_0^2} + \frac{N_t}{\sigma_\nu^2}\right)^{-1}$$
$$\mu_t = \boldsymbol{\beta}_n = \left(\frac{1}{\sigma_0^2} + \frac{N_t}{\sigma_\nu^2}\right)^{-1}\left(\frac{1}{\sigma_0^2}\mu_0 + \frac{N_t}{\sigma_\nu^2}\bar{\xi}_t\right)$$

  - Remember: Inverse of covariance matrix also called the "precision matrix"
  - Precisions add up in the normal linear regression model with conjugate priors
- Shows that we can equate the prior in period t with the mean and variance of a normal distribution:

$$\pi_t = (\mu_t, \sigma_t^2)$$

Transition of prior

- Conditional on the current prior $\pi_t = (\mu_t, \sigma_t^2)$, and only in periods t when the consumer receives a consumption signal ξ_t :

$$\sigma_{t+1}^2 = \left(\frac{1}{\sigma_t^2} + \frac{1}{\sigma_\nu^2}\right)^{-1}$$
$$\mu_{t+1} = \left(\frac{1}{\sigma_t^2} + \frac{1}{\sigma_\nu^2}\right)^{-1}\left(\frac{1}{\sigma_t^2}\mu_t + \frac{1}{\sigma_\nu^2}\xi_t\right)$$

- Can correspondingly derive the Markov transition density $f(\pi_{t+1}|\pi_t, a_t)$
  - σ_t^2 evolves deterministically
  - μ_{t+1} is normally distributed with mean $\mathbb{E}(\mu_{t+1}|\pi_t, a_t) = \mu_t$
    - Note that var$(\mu_{t+1}|\pi_t, a_t) \neq \sigma_{t+1}^2$

Model solution

- Reintroduce j subscript
- Priors are independent, and $\pi_t = (\pi_{1t}, \dots, \pi_{Jt})$
- Recall that $\xi_{jt} = \vartheta_j + \nu_{jt}$
- Expected utility:

$$u_j(x_t) = \mathbb{E}(\mathrm{u}_j(z_t, \xi_{jt})|\pi_t, z_t) = \gamma\mu_{jt} - \gamma r \mu_{jt}^2 - \gamma r(\sigma_{jt}^2 + \sigma_\nu^2) - \alpha P_{jt}$$

- State transition:

$$f(x_{t+1}|x_t, a_t) = f(z_{t+1}|z_t, a_t) \cdot \prod_{j=1}^{J} f(\pi_{j,t+1}|\pi_{jt}, a_t)$$

- $\Rightarrow$ well-defined dynamic discrete choice model with unique choice-specific value functions characterizing optimal decisions

Estimation

- Conditional on $x_t = (z_t, \pi_t)$ we can calculate the CCP's $\Pr\{a_t = j|z_t, \pi_t\} = \sigma_j(z_t, \pi_t)$
- Data
  - $Q_t = a_t$, $\boldsymbol{Q}_t = (Q_0, \dots, Q_t)$, and $\boldsymbol{Q} = (Q_0, \dots, Q_T)$
  - $\boldsymbol{z}_t = (z_0, \dots, z_t)$ and $\boldsymbol{z} = (z_0, \dots, z_T)$
- Assume $f(z_{t+1}|z_t, a_t)$ is known
  - Estimated from data in a preliminary estimation step
- Difficulty in constructing the likelihood: π_t is an unobserved state
  - Need to integrate out the unobserved states from the likelihood function

Constructing the priors

- Define $\nu_t = (\nu_{1t}, \dots, \nu_{Jt})$ and $\boldsymbol{\nu} = (\nu_0, \dots, \nu_{T-1})$
- Conditional on ϑ_j (an estimated parameter) and $\boldsymbol{\nu}$, we know the signals $\xi_{jt} = \vartheta_j + \nu_{jt}$
- By assumption $\pi_{j0}(\vartheta_j) = N(\mu_{j0}, \sigma_{j0}^2) = N(\bar{\vartheta}, \sigma_0^2)$
- Given $\boldsymbol{\nu}$ we can then infer the sequence of priors $\pi_0, \pi_1, \dots, \pi_{T-1}$:

$$\sigma_{j,t+1}^2 = \left(\frac{1}{\sigma_{jt}^2} + \frac{1}{\sigma_\nu^2}\right)^{-1}$$

$$\mu_{j,t+1} = \left(\frac{1}{\sigma_{jt}^2} + \frac{1}{\sigma_\nu^2}\right)^{-1}\left(\frac{1}{\sigma_{jt}^2}\mu_{jt} + \frac{1}{\sigma_\nu^2}\xi_{jt}\right)$$

Likelihood function

- $\theta \in \Theta$ is a vector of model parameters
- We know from the discussion above that $\Pr\{a_t = j|z_t, \pi_t; \theta\} = \Pr\{a_t = j|\boldsymbol{Q}_{t-1}, \boldsymbol{z}_t, \boldsymbol{\nu}; \theta\}$
- Define

$$l(\theta|\boldsymbol{Q}, \boldsymbol{z}) \equiv f(\boldsymbol{Q}|\boldsymbol{z}; \theta) = \int \left(\prod_{t=0}^{T} \Pr\{Q_t|\boldsymbol{Q}_{t-1}, \boldsymbol{z}_t, \boldsymbol{\nu}; \theta\}\right) f(\boldsymbol{\nu}; \theta) d\boldsymbol{\nu}$$

- High-dimensional integral
- Simulation estimator
  - Draw $\boldsymbol{\nu}^{(r)}$
  - Average over draws to simulate the likelihood
- $l(\theta|\boldsymbol{Q}, \boldsymbol{z}) = l_i(\theta|\boldsymbol{Q}_i, \boldsymbol{z}_i)$ is the likelihood contribution for one household i
  - Likelihood components conditionally independent across households given $\theta \Rightarrow$ joint likelihood is $l(\theta|\boldsymbol{Q}, \boldsymbol{z}) = \prod_i l_i(\theta|\boldsymbol{Q}_i, \boldsymbol{z}_i)$

Recap: Unobserved states and the likelihood function

- General approach
  1. Find a vector of random variables $\boldsymbol{\nu}$ with density $p(\boldsymbol{\nu};\theta)$ that allows to reconstruct the sequence of unobserved states
  2. Formulate the likelihood conditional on $\boldsymbol{\nu}$
  3. Integrate over $\boldsymbol{\nu}$ to calculate the likelihood conditional on data and θ only
- Note
  - Requires that numerical integration with respect to $\boldsymbol{\nu}$ or simulation from $p(\boldsymbol{\nu};\theta)$ is possible

Initial conditions

- Assumption so far: The first sample period is the first period when the consumers start buying the products
- What if the consumers made product choices before the sample period?
  - Would $\pi_{j0}(\theta) = N(\bar{\vartheta}, \sigma_0^2)$ be a valid assumption?
- Solution in Erdem and Keane (1996)
  - Split the sample periods into two parts, $t = 0, \ldots, T_0 - 1$ and $t = T_0, \ldots, T$
  - Make some assumption about the prior in period $t = 0$
  - Simulate the evolution of the priors conditional on the observed product choices and θ for periods $t \leq T_0$
  - Conditional on the simulated draw of $\pi_{j0}(\theta)$ formulate the likelihood using data from periods $t = T_0, \ldots, T$
- Note
  - The underlying assumption is that T_0 is large enough such that the effect of the initial condition in $t = 0$ vanishes

Unobserved heterogeneity

- The most common methods to incorporate unobserved heterogeneity in dynamic discrete choice models are based on a *finite mixture* or *latent class* approach
- Assume there is a finite number of M consumer/household types, each characterized by a parameter θ_m
- Let π_m be the fraction of consumers of type m in the population
- If there is no correlation between a consumer's type and the initial state, we can define the likelihood contribution for i as

$$l_i(Q_i; \theta, \pi) = \sum_{m=1}^{M} l_i(Q_i; \theta_m) \pi_m$$

  - Here, $\pi = (\pi_1, \dots, \pi_{M-1})$
- Obviously, the computational overhead increases in M

Unobserved heterogeneity

- Matters are more complicated if individuals are systematically in specific states depending on their type
  - For example, preference for a brand may be correlated with product experience or stockpiling of that brand
- In that case, $\pi_m(x_{i1}) = \Pr(i \text{ is of type } m | x_{i1}) \neq \pi_m$
- An easy solution (Heckman 1981) is to form an "auxiliary model" for this probability, i.e. make $\pi_m(x_{i1}; \tau)$ a parametrically specified function with τ to be estimated
- Example:

$$\pi_m(x_{i1}; \tau) = \frac{\exp(x_{i1}\tau_m)}{1 + \sum_{n=1}^{M-1} \exp(x_{i1}\tau_n)} \qquad \text{for } m = 1, \dots, M-1$$

Continuous types

- We now consider a more general form of consumer heterogeneity, which allows for a continuous distribution of types
- Consumer i's parameter vector ϑ_i is a function of a common component θ and some idiosyncratic component ω_i, which is drawn from a distribution with density $\phi(\cdot)$
  - ϕ is known, i.e., does not depend on any parameters that are estimated
- Example: $\theta = (\mu, L)$, and $\vartheta_i = \mu + L\omega_i$, $\omega_i \sim N(0, I)$

Continuous types

- Given (ω_i, θ), we can compute the choice specific value functions, calculate the CCP's $\Pr\{a_{it}|x_{it}; \omega_i, \theta\}$, and then obtain the likelihood contribution for individual i :

$$l_i(Q_i; \theta) = \int l_i(Q_i; \omega, \theta)\phi(\omega)d\omega$$

- If the integral is high-dimensional, we need to calculate it by simulation ($\omega^{(r)} \sim \phi$)

$$l_i(Q_i; \theta) \approx \frac{1}{R}\sum_{r=1}^{R} l_i(Q_i; \omega^{(r)}, \theta)$$

- The problem with this approach: instead of just re-solving for the value functions once when we change θ, we need to re-solve for the value function R times

Change of variables and importance sampling

- Ackerberg (2009) proposes a method based on a change of variables and importance sampling to overcome this computational challenge (see Hartmann 2006 for an application)
- Assume $\vartheta_i = \rho(\omega_i, \theta)$ (could additionally allow parameters to be a function of household characteristics)
- ϑ_i fully summarizes the behavior (choice-specific value functions) of household i
- Let $p(\vartheta_i|\theta)$ be the density of ϑ_i
  - The support of p must be the same for each θ
- In our example, $\vartheta_i \sim N(\mu, LL^T)$

Change of variables and importance sampling

- Let $g(\cdot)$ be a density that is positive on the support of p
  - For example, we could choose $g(\cdot) = p(\cdot|\theta_0)$ where θ_0 is some arbitrary parameter
- Then

$$
\begin{aligned}
l_i(Q_i;\theta) &= \int l_i(Q_i;\omega,\theta)\phi(\omega)d\omega \\
&= \int l_i(Q_i;\vartheta)p(\vartheta|\theta)d\vartheta \\
&= \int l_i(Q_i;\vartheta)\frac{p(\vartheta|\theta)}{g(\vartheta)}g(\vartheta)d\vartheta
\end{aligned}
$$

- In the second line we use a change of variables
- In the third line we use importance sampling

Change of variables and importance sampling

- We can now simulate the integral as follows:

$$l_i(Q_i;\theta) \approx \frac{1}{R}\sum_{r=1}^{R} l_i(Q_i;\vartheta^{(r)})\frac{p(\vartheta^{(r)}|\theta)}{g(\vartheta^{(r)})}$$

- Note that $\vartheta^{(r)}$ is drawn from g, a distribution that does not depend on the parameter vector θ
- Hence, as we change θ, only the weights $p(\vartheta^{(r)}|\theta)/g(\vartheta^{(r)})$ change, but not l_i!
  - We only need to calculate the value functions R times
- Intuition behind this approach: Change only the weights on different possible household types, not directly the households as we search for an optimal θ
- Based on the simulated likelihood function, we can define an SML (simulated maximum likelihood) estimator for θ

Understanding sequential, forward-looking learning

- Example: A firm launches a product, but is uncertain about its profitability
- Firm's profits from the product given by ϑ
- $\vartheta \in \{\vartheta_L, \vartheta_H\}$ could be either negative, $\vartheta_L < 0$, or positive, $\vartheta_H > 0$
- π_t: prior probability that profits are negative, $\pi_t = \Pr\{\vartheta = \vartheta_L\}$
- Firm decisions $a_t \in \{0,1\}$
  - $a_t = 0$ denotes that the firm scraps the product, and receives the payoff 0
  - $a_t = 1$ denotes that the firm stays in the market, and receives the payoff ϑ
- Assumption: If the firm stays in the market it observes profits and immediately learns the true value of ϑ

- Expected profit:

$$u_0(\pi_t) = 0$$
$$u_1(\pi_t) = \mathbb{E}(\vartheta|\pi_t) = \pi_t\vartheta_L + (1-\pi_t)\vartheta_H$$

- No latent payoff terms ϵ_{jt} in this model

- Suppose that the firm has information that the product is not profitable, that is

$$\mathbb{E}(\vartheta|\pi_t) = \pi_t\vartheta_L + (1-\pi_t)\vartheta_H < 0$$

- What action should the firm take—scrap or stay in the market?

Myopic vs. forward-looking decision making

- If the firm only cares about current profits, $\beta = 0$, then it should scrap the product:

$$u_1(\pi_t) < 0 = u_0(\pi_t)$$

- But what if $\beta > 0$?
- Let's approach this considering the impact of the current decision on future information, and the optimal decision that the firm can take based on future information

Choice-specific value functions

- If the firm stays in the market and sells the product, it learns the true level of profits
  - Hence, in the next period, $t + 1$, either $\pi_{t+1} = 0$ or $\pi_{t+1} = 1$
  - No more uncertainty about the profit level
- For the case of certainty we can easily calculate the value function:

$$v(1) = 0$$
$$v(0) = \frac{1}{1-\beta} \cdot \vartheta_H$$

- Choice-specific value functions for arbitrary π_t:

$$\begin{aligned} v_0(\pi_t) &= 0 \\ v_1(\pi_t) &= u_1(\pi_t) + \beta \mathbb{E}\left(v(\pi_{t+1}) | \pi_t, a_1 = 1\right) \\ &= (\pi_t \vartheta_L + (1-\pi_t)\vartheta_H) + \beta \left(\pi_t \cdot 0 + (1-\pi_t) \cdot \frac{1}{1-\beta} \cdot \vartheta_H\right) \\ &= \pi_t \vartheta_L + (1-\pi_t) \cdot \frac{1}{1-\beta} \cdot \vartheta_H \end{aligned}$$

Optimal forward-looking decision making under learning

- Keep the product in the market if and only if

$$v_1(\pi_t) = \pi_t \vartheta_L + (1-\pi_t) \cdot \frac{1}{1-\beta} \cdot \vartheta_H > 0 = v_0(\pi_t)$$

  - Reduces to myopic decision rule if $\beta = 0$
- Example
  - $\vartheta_L = -1$ and $\vartheta_H = 1$
  - Static decision making: Product will be scrapped iff $\pi_t \geq 0.5$
  - Suppose the firm is forward-looking, $\beta = 0.9$. Then the product will be scrapped iff $\pi_t \geq 10/11 \approx 0.909$

Overview

Overview

- Su and Judd (2012) propose an alternative approach to obtain the ML estimator for a dynamic decision process
  - Their method also works and has additional advantages for games
- They note that the nested fixed point approach is really only a special way of finding the solution of a more general constrained optimization problem
  - MPEC (mathematical programming with equilibrium constraints) approach
- There will often be better algorithms to solve an MPEC problem, which allows for faster and/or more robust computation of the ML estimate compared to the NFP approach

Implementation details

- Let's be precise on the computational steps we take to calculate the likelihood function
- We calculate the expected value function, $w(x;\theta)$, in order to derive the choice probabilities which allow us to "match" model predictions and observations
- On our computer, we will use some interpolation or approximation method to represent w
  - Representation will depend on a set of parameters, γ
  - γ represents the value of w at specific state points (interpolation), or coefficients on basis functions (approximation)
  - On our computer, γ completely defines w, $w \Leftrightarrow \gamma$

Implementation details

- The expected value function satisfies $w = \Gamma(w;\theta)$
- We can alternatively express this relationship as $\gamma = \Gamma(\gamma;\theta)$
  - Value function iteration on our computer proceeds by computing the sequence $\gamma^{(n+1)} = \Gamma(\gamma^{(n)};\theta)$, given some starting value $\gamma^{(0)}$
- For each parameter vector θ, there is a unique γ that satisfies $\gamma = \Gamma(\gamma;\theta)$
  - Denote this relationship as $\gamma = \psi(\theta)$

Implementation details

- Recall how we calculate the likelihood function:
  - Using the expected value function, $w \Leftrightarrow \gamma$, calculate the choice-specific values,

$$v_j(x;\theta) = u_j(x;\theta) + \beta \int w(x';\theta) f(x'|x,j;\theta)\, dx'$$

  - Then calculate the CCP's
- The equation above presumes that we use $w = \Gamma(w;\theta)$ to calculate the choice-specific values
- But we could alternatively use some arbitrary guess for w
  - Let γ be the parameter vector that summarizes some arbitrary w
  - Let $v_j(x;\theta,\gamma)$ be the corresponding choice-specific value function

MPEC approach

- We can use $v_j(x;\theta,\gamma)$ to calculate the CCP's and then the *augmented likelihood function*,

$$l(Q;\theta,\gamma)$$

- This expressions clearly shows what the likelihood function depends on:
  - Parameter values θ
  - The expected value function $w \Leftrightarrow \gamma$ describing how the decision maker behaves
- Our assumption of rational behavior imposes that the decision maker does not follow some arbitrary decision rule, but rather the rule corresponding to $\gamma = \Gamma(\gamma;\theta)$, denoted by $\gamma = \psi(\theta)$

MPEC approach

- We can thus express the likelihood estimation problem in its general form:

$$\max_{(\theta,\gamma)} \log(l(Q;\theta,\gamma))$$
$$\text{s.t. } \gamma - \Gamma(\gamma;\theta) = 0$$

  - Solve for the parameter vector (θ, γ)
  - The constraint is an equilibrium constraint, thus the term "MPEC"
- The nested fixed point algorithm is a special way of formulating this problem:

$$\max_{\theta} \log(l(Q;\theta)) \equiv \log(l(Q;\theta,\psi(\theta)))$$

- Note that both problem formulations define the same ML estimator

Discussion

- Benefits of MPEC estimation:
  - Avoid having to find an exact solution of the value function, $\gamma = \Gamma(\gamma;\theta)$, for each $\theta \Rightarrow$ speed advantage
  - The MPEC formulation allows for more robust convergence to the solution, because derivatives (of the objective and constraint) are easier to compute than in the NFP approach

Discussion

- Isn't a large state space ⇒ γ with many elements (say 100,000) an obvious obstacle to using the MPEC approach?
  - Good solvers, such as SNOPT or KNITRO can handle such problems
  - Will require a sparse Jacobian of the constraint, i.e. $\nabla \Gamma(\gamma; \theta)$ needs to be sparse
  - Will require that we are able to compute the Jacobian, $\nabla \Gamma(\gamma; \theta)$, which is an $L \times L$ matrix, where L is the number of elements in γ

Discussion

- How to calculate the Jacobian of the constraint?
  - Analytic derivatives → can be cumbersome and error-prone
  - Automatic differentiation (AD), available in C, C++, FORTRAN, and MATLAB → virtually no extra programming effort
    - Works fine for small-medium scale problems
    - *May* be difficult to implement (computer speed/memory requirements) given currently available AD software for large-scale problems

Overview

Overview

- Data are generated from the durable goods adoption model with two products discussed in Part I of the lecture
- We solve for w using bilinear interpolation
  - Exercise: Re-write the code using Chebyshev approximation, which is much more efficient in this example

Code documentation

- `Main.m`
  - Defines model parameters and price process parameters
  - Sets values for interpolator and initializes Gauss-Hermite quadrature information
  - If `create_data=1`, solves the decision process given parameters and simulates a new data set. Call to `simulate_price_process` to simulate prices and `simulate_adoptions` to simulate the corresponding adoption decisions. Plots aggregate adoption data based on the output of `calculate_sales`
  - Use script `display_DP_solution` to show choice probabilities and relative choice-specific value functions, $v_j(x) - v_0(x)$

Code documentation

- Main.m contains three estimation approaches:
  1. Use NFP estimator and MATLAB's built-in Nelder-Meade simplex search algorithm
  2. Use the TOMLAB package to find the NFP estimator. Numerical gradients are used
  3. Use TOMLAB and the MPEC approach
     - Allows for use of automatic differentiation using the TOMLAB/MAD module
- TOMLAB can be obtained at `http://tomopt.com/tomlab/` (ask for trial license)

Code documentation

- Solution of the optimal decision rule
  - Based on iteration on the integrated value function to solve for the expect value function
  - `Bellman_equation_rhs` updates the current guess of the expected value function
  - `Bellman_equation_rhs` uses `interpolate_2D` when taking the expectation of the future value

Code documentation

- NFP algorithm
  - Calls `log_likelihood`, which solves for the expected value function to calculate the CCP's
- MPEC algorithm
  - Calls `log_likelihood_augmented` with $\gamma \Leftrightarrow w$ supplied as parameters
  - `Bellman_equation_constraint` implements the equilibrium constraint, $\gamma - \Gamma(\gamma; \theta) = 0$

Overview

Omitted topics

- This lecture omits two recent estimation approaches:

1. Two-step estimators: Attempt to alleviate the computational burden inherent in NFP and MPEC estimation approaches (e.g. Pesendorfer and Schmidt-Dengler 2008)
2. Bayesian estimation using a new algorithm combining MCMC and value function iteration (Imai, Jain, and Ching 2009; Norets 2009)

Overview

Estimation and Identification

- Abbring, J. H. (2010): "Identification of Dynamic Discrete Choice Models," Annual Review of Economics, 2, 367-394
- Arcidiacono, P., and P. E. Ellickson (2011): "Practical Methods for Estimation of Dynamic Discrete Choice Models," Annual Review of Economics, 3, 363-394
- Ackerberg, D. A. (2009): "A new use of importance sampling to reduce computational burden in simulation estimation," Quantitative Marketing and Economics, 7(4), 343-376
- Aguirregabiria, V. and Mira, P. (2010): "Dynamic discrete choice structural models: A survey," Journal of Econometrics, 156, 38-67
- Dubé, J.-P., Hitsch, G. J., and Jindal, P. (2013): "The Joint Identification of Utility and Discount Functions from Stated Choice Data: An Application to Durable Goods Adoption," manuscript
- Fang, H., and Y. Wang (2013): Estimating Dynamic Discrete Choice Models with Hyperbolic Discounting, with an Application to Mammography Decisions," manuscript
- Heckman, J. (1981): "The Incidental Parameters Problem and the Problem of Initial Conditions in Estimating a Discrete Time-Discrete Data Stochastic Process," in C. Manski and D. L. McFadden (eds.): Structural Analysis of Discrete Data with Econometric Applications. MIT Press

Estimation and Identification (Cont.)

- Hotz, J. V., and R. A. Miller (1993): "Conditional Choice Probabilities and the Estimation of Dynamic Models," Review of Economic Studies, 60, 497-529
- Magnac, T. and Thesmar, D. (2002): "Identifying Dynamic Discrete Decision Processes," Econometrica, 70(2), 801-816
- Reiss, P. C., and F. A. Wolak (2007): "Structural Econometric Modeling: Rationales and Examples from Industrial Organization," in Handbook of Econometrics, Vol. 6A, ed. by J. J. Heckman and E. E. Leamer. Elsevier B. V.
- Rust, J. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," Econometrica, 55 (5), 999-1033
- Rust, J. (1994): "Structural Estimation of Markov Decision Processes," in Handbook of Econometrics, Vol. IV, ed. by R. F. Engle and D. L. McFadden. Amsterdam and New York: North-Holland
- Su, C.-L., and K. L. Judd (2012): "Constrained Optimization Approaches to Estimation of Structural Models," Econometrica, 80(5), 2213-2230
- Yao, S., C. F. Mela, J. Chiang, and Y. Chen (2012): "Determining Consumers' Discount Rates with Field Studies," Journal of Marketing Research, 49(6), 822-841

Other Estimation Approaches (Two-Step, Bayesian, ...)

- Aguirregabiria, V., and P. Mira (2002): "Swapping the Nested Fixed Point Algorithm: A Class of Estimators for Discrete Markov Decision Models," Econometrica, 70 (4), 1519-1543
- Aguirregabiria, V. and Mira, P. (2007): "Sequential Estimation of Dynamic Discrete Games," Econometrica, 75(1), 1-53
- Arcidiacono, P. and Miller, R. (2008): "CCP Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity," manuscript
- Bajari, P., L. Benkard, and J. Levin (2007): "Estimating Dynamic Models of Imperfect Competition," Econometrica, 75 (5), 1331–1370
- Imai, S., Jain, N., and Ching, A. (2009): "Bayesian Estimation of Dynamic Discrete Choice Models," Econometrica, 77(6), 1865-1899
- Norets, A. (2009): "Inference in Dynamic Discrete Choice Models With Serially Correlated Unobserved State Variables," Econometrica, 77(5), 1665-1682
- Pesendorfer, M. and Schmidt-Dengler, P. (2008): "Asymptotic Least Squares Estimators for Dynamic Games," Review of Economic Studies, 75, 901-928

Other Applications

- Erdem, T., and M. P. Keane (1996): "Decision-Making Under Uncertainty: Capturing Brand Choice Processes in Turbulent Consumer Goods Markets," Marketing Science, 15 (1), 1–20
- Hartmann, W. R. (2006): "Intertemporal Effects of Consumption and Their Implications for Demand Elasticity Estimates," Quantitative Marketing and Economics, 4, 325-349
- Hartmann, W. R. and Viard, V. B. (2008): "Do frequency reward programs create switching costs? A dynamic structural analysis of demand in a reward program," Quantitative Marketing and Economics, 6(2), 109-137
- Hitsch, G. J. (2006): "An Empirical Model of Optimal Dynamic Product Launch and Exit Under Demand Uncertainty," Marketing Science, 25(1), 25-50
- Misra, S. and Nair, H. (2011): "A Structural Model of Sales-Force Compensation Dynamics: Estimation and Field Implementation," Quantitative Marketing and Economics, 9, 211-257