Biometrics 65, 505–513
June 2009
DOI: 10.1111/j.1541-0420.2008.01113.x

Nested Markov Compliance Class Model in the Presence of Time-Varying Noncompliance

Julia Y. Lin,[1,*] Thomas R. Ten Have,[2] and Michael R. Elliott[3]

[1]Center for Multicultural Mental Health Research, Cambridge Health Alliance-Harvard Medical School, Somerville, Massachusetts 02143, U.S.A.
[2]Department of Biostatistics and Epidemiology, University of Pennsylvania, Philadelphia, Pennsylvania 19104, U.S.A.
[3]Department of Biostatistics and Institute of Social Research, University of Michigan, Ann Arbor, Michigan 48109, U.S.A.
[*]*email:* jlin@charesearch.org

Summary. We consider a Markov structure for partially unobserved time-varying compliance classes in the Imbens–Rubin (1997, *The Annals of Statistics* **25,** 305–327) compliance model framework. The context is a longitudinal randomized intervention study where subjects are randomized once at baseline, outcomes and patient adherence are measured at multiple follow-ups, and patient adherence to their randomized treatment could vary over time. We propose a nested latent compliance class model where we use time-invariant subject-specific compliance principal strata to summarize longitudinal trends of subject-specific time-varying compliance patterns. The principal strata are formed using Markov models that relate current compliance behavior to compliance history. Treatment effects are estimated as intent-to-treat effects within the compliance principal strata.

Key words: Geriatric depression; Hidden Markov model; Latent class; Longitudinal compliance class model; Noncompliance; Principal stratification.

1. Introduction

In randomized intervention studies where interventions are administered repeatedly, subject adherence to the randomized treatment may vary over time. In addition, the effect of the treatment from previous time points on the outcome may be nontransient. We propose a longitudinal compliance class model with decay parameters for treatment effects that uses a nested principal stratification structure to characterize longitudinal compliance patterns over time within which intent-to-treat (ITT) effects are estimated. We consider a Markov structure for the time-varying subject adherence to randomized treatment. We illustrate the model with analysis of the "Prevention of Suicide in Primary Care Elderly: Collaborative Trial" (PROSPECT; Bruce et al., 2004).

The PROSPECT study was a randomized intervention study targeted at elderly patients with depression in primary care practices. There were two treatment groups: usual care and the intervention. In the usual care group, patients received standard care. In the intervention group, patients were assigned to meet with health specialists who educated patients, their families, and physicians about depression, treatment, and monitored adherence to treatment. Primary care practices were randomized to the treatments rather than individual patients to prevent contamination of treatments among patients within the same practice and for practicality. Patients were followed for 2 years from the initial randomization. Clinical depression outcome and adherence to randomized treatment were measured at 4, 8, 12, 18, and 24 months. There were 598 patients in the study. The clinical outcome of interest was the severity of depression measured by the Hamilton depression score (HAMD). We consider an all-or-none treatment adherence measured by whether patients met with the health specialists at least once since the previous follow-up period. We are interested in investigating the effect of the intervention on depression severity accounting for treatment adherence over time.

When subjects do not adhere to the treatment to which they are randomized, subject noncompliance could confound the relationship between the treatment and the outcome. Therefore, it is important to account for subject noncompliance when estimating the effect of the treatment. One way to do that is by using principal stratification strategies (Frangakis and Rubin, 1999, 2002). Angrist, Imbens, and Rubin (1996) and Imbens and Rubin (1997) proposed to use compliance classes to describe subject compliance behaviors within which ITT contrasts are made to estimate the causal effect of the treatment on the outcome.

Cross-sectional studies with two treatment arms, experimental treatment and control treatment, have four possible compliance classes: compliers, always-takers, never-takers, and defiers. Compliers are those who would adhere to the treatment to which they are assigned; always-takers are those who would seek the experimental treatment regardless of their treatment assignment; never-takers are those who would opt

for the control treatment regardless of their treatment assignment; and defiers are those who would refuse the treatment to which they are assigned and choose to receive the other treatment.

In studies, such as the PROSPECT, where those assigned to the control treatment have no access to the experimental treatment, there are only compliers and never-takers. Always-takers and defiers cannot exist because those randomized to the control treatment cannot receive the experimental treatment. The compliance classes for those assigned to the experimental treatment in this study design are observed. Subjects assigned to and receiving the experimental treatment are compliers; subjects assigned to the experimental treatment but receiving the control treatment are never-takers. The compliance classes for those assigned to the control treatment are unobserved.

We propose an extension of the cross-sectional model in Imbens and Rubin (1997) to longitudinal settings. Yau and Little (2001) proposed an extension where outcome was measured repeatedly over time; however, adherence to intervention was only recorded once and did not vary. Our proposed model allows treatment adherence to vary over time. In Frangakis et al. (2004), outcome was repeatedly measured over time, and subject compliance could vary over time. This model differs from our proposed model in two ways: (i) we restrict our method to study designs where randomization status do not change over time; and (ii) we propose a nested model structure that uses subject-specific time-invariant principal strata to summarize subject-specific time-varying compliance behavior. The subject-level time-invariant strata allows us to classify subjects based on their longitudinal compliance, and relate longitudinal compliance to outcomes.

In the presence of time-varying compliance behaviors, it may be useful to consider patterns of longitudinal compliance behavior when examining longitudinal outcomes. Subjects with different compliance trajectories may differ in treatment outcomes. We may make inferences on different longitudinal compliance patterns and the longitudinal outcomes associated with those patterns. In a study like the PROSPECT where there are two possible compliance classes and five follow-up visits, we have 32 (2^5) possible compliance patterns. It may be impractical and not clinically meaningful to look at the longitudinal outcomes in all of the 32 patterns. Hence, it may be more helpful to have summary measures of the longitudinal compliance patterns in the data, and look at longitudinal outcomes within broader latent classes.

We use the nested latent class model framework proposed by Lin, Ten Have, and Elliott (2008) to accommodate time-varying latent compliance classes by specifying broader principal strata that summarize the compliance classes. The nested latent class model involves two levels of compliance class models. The first level uses subject-specific time-varying compliance classes to describe the time-varying treatment adherence; the second level uses subject-specific time-invariant compliance "superclasses" to summarize the longitudinal patterns of compliance classes. The superclass defined here is a principal stratum in the sense that the superclass is a function of compliance classes, in which the compliance classes describe the relationship between treatment received and treatment randomization, and that the function itself is not affected by the actual treatment randomization. Treatment received is a function of the compliance classes and the treatment randomization. It is consistent with the definition of principal stratum in Frangakis and Rubin (2002), and similar to the principal stratum in Frangakis et al. (2004). The superclass is a "coarser" principal stratum. The ITT effect of the intervention stratified on compliance superclass, or the principal effect (Frangakis and Rubin, 2002), is estimated to control for longitudinal subject treatment noncompliance.

Lin et al. (2008) makes the conditional independence (CI) assumption that compliance classes at each time point within an individual are independent from each other given the individual's compliance superclass and baseline covariates. In other words, knowing the compliance superclass and subject baseline characteristics, the history of compliance behaviors does not provide any more information on the current compliance behavior. This may be a strong assumption, which we now propose to assess with a Markov model for the time-varying compliance classes. We fit a latent transitional model (Collins and Wugalter, 1992) incorporating covariates in estimating transitional probabilities (Reboussin, Liang, and Reboussin, 1999). We assume a first-order Markov structure for the compliance classes given superclass and baseline covariates where compliance behaviors are assumed to depend on the compliance class in the previous time point. Modeling the Markov structure of the time-varying compliance classes will allow us to: (i) utilize information from history of compliance to predict compliance behaviors; and (ii) examine how history of compliance relates to compliance behavior.

As another extension of Lin et al. (2008), this article considers the nontransient effect of treatment over time. In the PROSPECT we may consider the decay of the ITT effect of the treatment on the outcome. It is conceivable that information ascertained in meetings with health specialists may have lasting effects on the subjects and their treatment outcomes.

We will define notation, discuss assumptions, principal effects, the parametric model, parameter estimation, the handling of missing outcomes, and assessment of model fit in Section 2. Then we will proceed to discuss the analysis results in Section 3, and make concluding remarks in Section 4.

2. Nested Compliance Class Model

2.1 *Notation*

Let Z_i denote the randomization status for subject i, where $i = (1, \ldots, N)$, and $Z_i \in (0, 1)$ for usual care and the intervention, respectively. Similarly, let D_{ij} denote the time-varying treatment received for subject i at time j, where $j = (1, 2, 3, 4, 5)$ for 4, 8, 12, 18, and 24 months, respectively, and $D_{ij} \in (0, 1)$ for usual care and intervention, respectively. Note that Z_i does not have the subscript j because we are restricting to designs where randomization does not change over time. Let Y_{ij} denote the observed outcome for subject i at time j. We use $\mathbf{Z}, \mathbf{D}, \mathbf{Y}$ to denote vectors of Z_i, D_{ij}, and Y_{ij}.

Following Little and Rubin (2000), we use $Y_{ij}(Z)$ to denote the partially latent potential outcome, an outcome that would have been observed, for subject i at time j if randomized to treatment Z. Let C_{ij} denote membership of the partially latent compliance classes for subject i at time j. In

the PROSPECT, because those randomized to the usual care group have no access to the intervention, there are only two possible compliance classes: compliers and never-takers; therefore, $C_{ij} \in (c, n)$. We use $\mathbf{C}$ to denote the vector of C_{ij}. The proposed principal stratification strategy uses compliance "superclasses" to summarize the longitudinal compliance patterns in the data within which we can stratify on and compare potential outcomes. It precludes the confounding when stratifying on observed postrandomization compliance patterns. Let U_i denote membership of the latent superclass for subject i, where $U_i = (1, \ldots, K)$ for assumed K numbers of latent superclasses. We use $\mathbf{U}$ to denote the vector of U_i.

Subject-level baseline covariates $\mathbf{A}_i$ and $\mathbf{Q}_i$ are used in modeling the outcome and compliance probabilities, respectively. We use $\mathbf{A}$ and $\mathbf{Q}$ to denote vectors of $\mathbf{A}_i$ and $\mathbf{Q}_i$.

We use uppercase letter to denote random variables or indices of potential outcomes (e.g., $Y_{ij}(Z)$), and lowercase letter to denote realized or observed values of random variables or indices (e.g., $Z_i = z$).

2.2 *Assumptions*

We make the randomization (Rubin, 1978), stable unit-treatment value (SUTVA; Rubin, 1986), and model assumptions to identify causal model parameters. We assume that potential outcomes, latent compliance classes, and latent compliance superclasses (which are assumed to be baseline characteristics) are independent of the randomization assignment status. We make the no interference assumption of the SUTVA and assume that the potential outcomes of an individual are not influenced by the treatment assignment of another individual. We also make the consistency assumption of the SUTVA which assumes that the potential outcome of a certain treatment will be the same regardless of the treatment assignment mechanism. It implies that the observed outcome is a function of the potential outcomes and treatment assignment: $Y_{ij} = Z_i * Y_{ij}(1) + (1 - Z_i) * Y_{ij}(0)$. The SUTVA assumption is violated when there is interference among subjects or when there are versions of treatments not represented by the treatment indicator variable.

2.3 *Principal Effects*

We utilize the compliance superclasses to summarize the longitudinal compliance patterns and estimate the ITT effects stratified on these superclasses. A compliance superclass is a latent subject-level principal stratum that is time invariant, and is considered to be a prerandomization characteristic that allows us to model potential outcomes conditional on prospective postrandomization behavior.

Our effect of interest is the principal effect of treatment assignment on the outcome within a compliance superclass at time j:

$$E[Y_{ij}(Z = 1) \,|\, U_i = k] - E[Y_{ij}(Z = 0) \,|\, U_i = k]. \tag{1}$$

It is an ITT contrast stratified on the compliance superclass. Because the superclasses defined here create baseline principal strata summarizing longitudinal compliance behaviors and do not represent specific longitudinal compliance patterns, the principal effect may sacrifice straightforward causal interpretation. The interpretation of the principal effects relies on the interpretation of the superclasses. Nonetheless, it allows us to consider the effect of treatment randomization controlling for longitudinal compliance.

The principal effect can be defined by observed outcomes under the randomization and the SUTVA consistency assumption:

$$\begin{aligned} &E[Y_{ij}(Z = 1) \,|\, U_i = k] - E[Y_{ij}(Z = 0) \,|\, U_i = k] \\ &\quad = E[Y_{ij}(Z = 1) \,|\, Z_i = 1, U_i = k] \\ &\qquad - E[Y_{ij}(Z = 0) \,|\, Z_i = 0, U_i = k] \\ &\quad = E[Y_{ij} \,|\, Z_i = 1, U_i = k] - E[Y_{ij} \,|\, Z_i = 0, U_i = k]. \end{aligned} \tag{2}$$

The first equal sign follows from the randomization assumption, which says that randomization is independent of baseline characteristics (e.g., potential outcomes) conditional on baseline covariates (e.g., compliance superclass). The second equal sign follows from the SUTVA consistency assumption, which implies that the observed outcome given treatment assignment z is the potential outcome for treatment assignment $Z = z$.

2.4 *Parametric Model*

The CI model proposed in Lin et al. (2008) assumes that longitudinal compliance classes within an individual are independent given compliance superclass and baseline covariates. Under the current proposed model we relax the CI assumption. We assume compliance classes are dependent on the compliance classes at one or more previous time points, the compliance superclass, and baseline covariates. As one reviewer pointed out, this model is a hidden Markov model similar to those used in "mover-stayer" applications (Langeheine and Van de Pol, 2002).

Following the CI model, we assume outcomes within individuals are independent given randomization, time-varying compliance class, baseline covariates, and subject-level random effect.

$$\begin{aligned} &\big(Y_{ij} \,|\, C_{i1}, \ldots, C_{ij}, Z_i = z, \mathbf{A}_i, \mathbf{W}_i, \\ &\boldsymbol{\lambda}, \zeta(t, j), \boldsymbol{\gamma}, \boldsymbol{\varphi}_i, \sigma^2\big) \overset{\text{ind}}{\sim} N(\mu_{ijz}, \sigma^2), \\ &\mu_{ijz} = \sum_{t=1}^{j} \left[\sum_{\eta'} I(C_{it} = \eta', Z_i = z) \lambda_{t\eta' z} \zeta(t, j) \right] \\ &\qquad + \mathbf{A}_i^T \boldsymbol{\gamma} + \mathbf{W}_i^T \boldsymbol{\varphi}_i. \end{aligned} \tag{3}$$

The conditional mean of the outcome has three components: compliance class-specific effect of randomization, the effect of baseline covariates, and the subject-specific random effects to account for within-subject correlation in the outcomes. The compliance class-specific effect of randomization on the outcome is represented by $\sum_{t=1}^{j} [\sum_{\eta'} I(C_{it} = \eta', Z_i = z) \lambda_{t\eta' z} \zeta(t, j)]$, where $\lambda_{t\eta' z}$ for $t \le j$ describes the compliance-class specific ITT effect of the treatment on the outcome, $\boldsymbol{\lambda}$ denotes the vector of $\lambda_{t\eta' z}$, and $\zeta(t, j)$ modifies that ITT effect at time t on the outcome at time j. The effect of the baseline covariates on the outcome is represented by $\mathbf{A}_i^T \boldsymbol{\gamma}$, where $\mathbf{A}_i$ denotes the vector of baseline covariates of subject i, and the column vector $\boldsymbol{\gamma}$ denotes the corresponding coefficients. The random effects $\boldsymbol{\varphi}_i$ are used to account for within-subject correlation in the outcomes, where $\mathbf{W}_i$ denotes the random effect design matrix for subject i. In our preliminary analysis

we found small within-practice correlation (0.075); hence, clustering by primary care practice was ignored, as in Bruce et al. (2004) and Small et al. (2006). We consider a random subject-level intercept model.

To model the nontransient effect of the treatment on subsequent outcomes, we use the parameter $\zeta(t, j)$ to modify the impact of the ITT effect at time t on the outcome at time j. We can assume a transient relationship where the outcome at time j is not dependent on the ITT effect at time t (i.e., $\zeta(t, j) = I(t = j)$); assume a nontransient relationship where the outcome at time j is dependent on the cumulative ITT effect of current and all prior time periods (i.e., $\zeta(t, j) = I(t \leq j)$); or assume a decaying relationship where the outcome at time j is dependent on the cumulative ITT effect of current and all prior time periods, but the influence of past treatment effects diminish as time lag increases (i.e., $\zeta(t, j) = e^{-\tau(j-t)}$ where $\tau > 0$). Preliminary analysis of the data using a decay model suggested $\tau \to \infty$, or a transient relationship. Hence, we consider the transient model:

$$\mu_{ijz} = \sum_{\eta'} [I(C_{ij} = \eta', Z_i = z)\lambda_{j\eta' z}] + \mathbf{A}_i^T \boldsymbol{\gamma} + \mathbf{W}_i^T \boldsymbol{\varphi}_i. \quad (4)$$

To relax the CI assumption of the time-varying compliance classes of the CI model, we propose a Markov compliance class (MCC) model where the compliance classes are dependent on past compliance behavior. Similar to the CI model, we assume that compliance superclass is an underlying factor that drives subject compliance over time. We model the compliance class at the first time point conditional on the compliance superclass and baseline covariates $\mathbf{Q}_i$ using logit models: $P(C_{i1} = \eta \mid U_i = k, \mathbf{Q}_i) = \omega_{k\eta}(\mathbf{Q}_i)$ and $\omega_{k\eta}(\mathbf{Q}_i) = \exp(\alpha_{0k\eta} + \boldsymbol{\alpha}_{1\eta}\mathbf{Q}_i)/[\sum_{\eta'} \exp(\alpha_{0k\eta'} + \boldsymbol{\alpha}_{1\eta'}\mathbf{Q}_i)]$ where $\sum_\eta \omega_{k\eta}(\mathbf{Q}_i) = 1 \,\forall\, k$. We constrain $\alpha_{0k\eta}$ and $\boldsymbol{\alpha}_{1\eta}$ for one of the compliance class η to be 0 for identifiability. In the presence of more than two compliance classes, we can use multinomial logit models instead of logistic models to model the compliance probabilities.

We assume subject compliance superclass $(U_i = k) \sim$ Multinomial$(1, p_k)$, where $\sum_k p_k = 1$. Compliance superclass between subjects are assumed to be independent: $f(\mathbf{U}) = \prod_{i=1}^N f(U_i = k)$ for $k = 1, \ldots, K$, where $f(.)$ denotes the distribution function.

We utilize latent transition models (Collins and Wugalter, 1992) to characterize the Markov process of compliance classes across time. In this article, we consider a nonstationary first-order MCC model. The number of model parameters in multiple-order Markov models increases exponentially without additional constraints such as stationarity. Because of the lack of good predictors of compliance transitions, we assume that there are no associated covariates influencing the transitional probabilities. Covariates can be incorporated using logit models as in Reboussin et al. (1999). We assume the compliance class transitions $(C_{ij} = \eta \mid C_{i,j-1} = \eta', U_i = k) \sim$ Multinomial$(1, \pi_{kj\eta'\eta})$, where $\sum_\eta \pi_{kj\eta\eta} = 1 \,\forall\, k, j, \eta'$. The joint distribution of the compliance classes given compliance superclass then becomes:

$$\begin{aligned} &P(C_{i1}, \ldots, C_{i5} \,|\, U_i, \mathbf{Q}_i) \\ &\quad = P(C_{i1} \,|\, U_i, \mathbf{Q}_i) P(C_{i2} \,|\, C_{i1}, U_i) \cdots P(C_{i5} \,|\, C_{i4}, U_i). \end{aligned} \quad (5)$$

If compliance class and compliance superclass memberships, and missing outcomes are known, the joint distribution of the complete data for subject i given the model specifications is as follows:

$$\begin{aligned} &f(Y_{i1}, \ldots, Y_{i5}, \boldsymbol{\varphi}_i, C_{i1}, \ldots, C_{i5}, U_i \,|\, Z_i, \mathbf{A}_i, \mathbf{Q}_i, \mathbf{W}_i, \boldsymbol{\theta}) \\ &\quad = f(Y_{i1}, \ldots, Y_{i5} \,|\, \boldsymbol{\varphi}_i, C_{i1}, \ldots, C_{i5}, U_i, Z_i, \mathbf{A}_i, \mathbf{Q}_i, \mathbf{W}_i, \boldsymbol{\theta}) \\ &\qquad \times f(\boldsymbol{\varphi}_i \,|\, C_{i1}, \ldots, C_{i5}, U_i, Z_i, \mathbf{A}_i, \mathbf{Q}_i, \mathbf{W}_i, \boldsymbol{\theta}) \\ &\qquad \times f(C_{i1}, \ldots, C_{i5} \,|\, U_i, Z_i, \mathbf{A}_i, \mathbf{Q}_i, \mathbf{W}_i, \boldsymbol{\theta}) \\ &\qquad \times f(U_i \,|\, Z_i, \mathbf{A}_i, \mathbf{Q}_i, \mathbf{W}_i, \boldsymbol{\theta}) \\ &\quad = f(Y_{i1}, \ldots, Y_{i5} \,|\, C_{i1}, \ldots, C_{i5}, Z_i, \mathbf{A}_i, \mathbf{W}_i, \lambda, \boldsymbol{\gamma}, \boldsymbol{\varphi}_i, \sigma^2) \\ &\qquad \times f(\boldsymbol{\varphi}_i \,|\, \Sigma_\varphi) f(C_{i1}, \ldots, C_{i5} \,|\, U_i, \mathbf{Q}_i) f(U_i), \end{aligned} \quad (6)$$

where $\boldsymbol{\theta} = (\boldsymbol{\lambda}, \boldsymbol{\gamma}, \sigma^2, \Sigma_\varphi)$.

Knowing the time-varying compliance classes, the superclass does not provide additional information on the longitudinal compliance behavior. Therefore, we assume that the potential outcomes are conditionally independent of the superclasses given compliance classes. However, because superclasses are functions of the compliance classes, we can use estimated effects associated with the compliance classes to estimate effects associated with the superclasses.

Under these model specifications, the principal ITT effect of the intervention on the outcome stratified on compliance superclass defined in equation (1) becomes

$$\begin{aligned} &E[Y_{ij}(Z = 1) \,|\, U_i = k] - E[Y_{ij}(Z = 0) \,|\, U_i = k] \\ &\quad = \sum_{\eta'} (\lambda_{j\eta' 1} - \lambda_{j\eta' 0})\, P(C_{ij} = \eta' \,|\, U_i = k). \end{aligned} \quad (7)$$

2.5 *Estimation*

We use Bayesian Markov chain Monte Carlo (MCMC) methods to estimate model parameters. For details of the priors and the conditional draws of the Gibbs sampler, please refer to the Web Appendix.

2.6 *Missing Outcome Imputation*

To deal with missing outcomes we assume a latent ignorable missing data mechanism (Peng, Little, and Raghunathan, 2004), which assumes missing at random given latent compliance class and covariates. At each iteration of the MCMC procedure, we impute the missing outcomes conditional on compliance classes, treatment randomization, baseline covariates, and subject-level random effects. We draw the missing outcome Y_{ij}^{mis} for subject i at time j from its predictive distribution given current values of parameters $C_{ij}, \lambda_{j\eta z}, \boldsymbol{\gamma}, \boldsymbol{\varphi}_i, \sigma^2$, and vector of observed outcomes $\mathbf{Y}^{\text{obs}}$.

$$\begin{aligned} &\left(Y_{ij}^{\text{mis}} \,\middle|\, \mathbf{Y}^{\text{obs}}, C_{ij}, Z_i = z, \mathbf{A}_i, \mathbf{W}_i, \lambda_{j\eta z}, \boldsymbol{\gamma}, \boldsymbol{\varphi}_i, \sigma^2\right) \sim N\left(\mu_{ijz}^*, \sigma^2\right) \\ &\qquad \mu_{ijz}^* = \sum_{\eta'} [I(C_{ij} = \eta', Z_i = z)\lambda_{j\eta' z}] + \mathbf{A}_i^T \boldsymbol{\gamma} + \mathbf{W}_i^T \boldsymbol{\varphi}_i. \end{aligned} \quad (8)$$

2.7 *Model Fit Assessment*

We compare the fits of the MCC model and the CI model by comparing the posterior predictive distributions (PPD; Gelman et al., 2004) of the time-varying compliance classes. Let G_m denote the number of individuals in the m^{th} of the 32

possible longitudinal compliance patterns and let κ_m be the estimated probability of exhibiting the m^{th} longitudinal compliance pattern. We consider the χ^2-type statistics:

$$S^{\text{obs}} = \sum_m \frac{\left(G_m^{\text{obs}} - N\kappa_m\right)^2}{N\kappa_m\left(1-\kappa_m\right)} \quad \text{and} \quad S^{\text{rep}} = \sum_m \frac{\left(G_m^{\text{rep}} - N\kappa_m\right)^2}{N\kappa_m\left(1-\kappa_m\right)}, \tag{9}$$

where G_m^{obs} is the observed statistics and G_m^{rep} is the repeated statistic obtained from draws of the parameters generated by the Gibbs sampler. The PPD p-value is then given by: $\sum_l I[(S^{\text{obs}})^l < (S^{\text{rep}})^l]/\sum_l 1$, where $(S^{\text{obs}})^l$ and $(S^{\text{rep}})^l$ denote the S^{obs} and S^{rep} from the l^{th} Gibbs draw. A PPD p-value close to 0.50 indicates a good fit of the model to the data.

3. Results

We demonstrate the MCC model with analysis of the PROSPECT data and compare the results to the analysis under the CI model. In the PROSPECT, those randomized to the usual care group do not have access to the intervention; therefore, there are only two compliance classes: compliers and never-takers. Goodman (1974) suggests that we can only identify at most three latent compliance superclasses given five dichotomous compliance classes; hence, we consider a maximum of three superclasses.

Unrecorded treatment received (D_{ij}) is assumed to be 0, indicating no visits with health specialists. In this analysis we let $\mathbf{A}_i$ be the baseline HAMD score and baseline suicidal ideation. We adjust for the baseline HAMD because we are interested in the change in HAMD scores from baseline. Treatment randomization failed to balance the proportion of subjects with suicidal ideation at baseline between the treatment groups; therefore, we adjust for it in modeling the outcome. We let $\mathbf{Q}_i$ be the baseline HAMD score in estimating the compliance probabilities in the CI model and in estimating the initial compliance probabilities in the MCC model.

We use relatively flat priors in the Bayesian MCMC estimation of the model parameters because we do not have strong prior inclinations. Following Garrett and Zeger (2000) and Ten Have et al. (2004), we assume $\boldsymbol{\alpha} \sim MVN(\mathbf{0}, \Sigma_\alpha = diag(50, 4))$. The difference in variance component in the priors reflect the different scaling of the covariates. A larger variance is used for binary covariates (i.e., intercept) and a smaller variance is used for continuous covariates (i.e., baseline HAMD score). The identifiability of the $\boldsymbol{\alpha}$ parameter is checked by comparing the prior and the posterior distributions (Garrett and Zeger, 2000). We assume the prior $(\pi_{kj\eta'c}, \pi_{kj\eta'n}) \sim$ Dirichlet(0.01, 0.01) $\forall k$, j, η' for the transitional probabilities. This is equivalent to adding 0.01 subject to each of the $(C_{i,j-1} = \eta', C_{ij} = \eta \mid U_i = k)$ groups. Let $\boldsymbol{\beta} = [\lambda_{1c0}, \ldots, \lambda_{5n1}, \boldsymbol{\gamma}]$, and we assume $\boldsymbol{\beta} \sim MVN(\mu_\beta = \mathbf{0}, \Sigma_\beta = 1000 \times \mathbf{I})$ and $\sigma^2 \sim \text{Inv} - \chi^2(\nu_\sigma = 1, \psi = 1/10)$. For the random effect variance parameter we assume $\Sigma_\varphi \sim \text{Inv} - \chi^2(\nu_\varphi = 1, \Gamma = 1/10)$. We assume the prior $(p_1, \ldots, p_K) \sim$ Dirichlet$(1, \ldots, 1)$, assigning a priori one subject to each of the K superclasses.

To assess the convergence of the MCMC chains we used the Gelman–Rubin $\hat{R}$ statistic (Gelman et al., 2004, p. 296–297), and $\hat{R} < 1.1$ is accepted as evidence of convergence. We ran three chains of the CI model for 10,000 iterations each with the first 1000 iterations discarded as burn-in, and ran three chains of the MCC model for 150,000 iterations each with the first 75,000 iterations discarded as burn-in. The maximum $\hat{R}$ was 1.05 and 1.08 for the CI and the MCC models, respectively.

We present the results under the CI model as specified in Lin et al. (2008), then the results under the MCC model, followed by comparison of the two models. We can assess the conditional independence assumption made under the CI model by comparing the fit of the CI model to the fit of the MCC model to the data.

3.1 *Conditional Independence Model*

In Lin et al. (2008) we found that the three-class CI model has a better fit to the data than the two-class CI model. Hence, we compare the three-superclass CI model to the MCC model. Table 1 shows the time- and superclass-varying compliance probabilities under the CI model assuming the average baseline HAMD of 18.1, and Table 2 shows the ITT effect of randomization on the outcome within each compliance superclass adjusting for the baseline HAMD and baseline suicidal ideation.

Table 1 shows that the first superclass under the CI model consists of subjects who are noncompliant at the 4-month follow-up and become even more noncompliant for the remainder of the study (low compliers). The second superclass consists of subjects who are highly compliant for the first 8 months and become increasingly noncompliant (decreasing

Table 1

Posterior means and 95% *credible intervals* (*in parentheses*) *for the time- and compliance superclass-varying compliance probabilities assuming the average baseline HAMD of* 18.1 *and superclass probabilities under the CI model*

Time	Low compliers	Decreasing compliers	High compliers
4 months	0.43 (0.33, 0.53)	0.99 (0.96, 1.00)	1.00 (0.98, 1.00)
8 months	0.01 (0.00, 0.07)	0.99 (0.94, 1.00)	1.00 (0.99, 1.00)
12 months	0.01 (0.00, 0.04)	0.51 (0.36, 0.66)	1.00 (0.98, 1.00)
18 months	0.06 (0.02, 0.12)	0.11 (0.00, 0.28)	0.99 (0.98, 1.00)
24 months	0.04 (0.01, 0.09)	0.01 (0.00, 0.07)	0.83 (0.77, 0.90)
$P(U_i)$	0.28 (0.23, 0.33)	0.16 (0.12, 0.22)	0.56 (0.50, 0.62)

Table 2
Posterior means and 95% *credible intervals* (*in parentheses*) *for the ITT contrasts of the outcome within compliance superclasses under the CI model*

Time	Low compliers	Decreasing compliers	High compliers
4 months	−7.54 (−10.05, −2.00)	−1.35 (−3.23, 0.10)	−1.32 (−3.20, 0.09)
8 months	−3.39 (−7.24, 0.81)	−0.93 (−2.78, 0.83)	−0.92 (−2.78, 0.86)
12 months	0.84 (−2.21, 3.95)	−0.61 (−2.11, 1.05)	−2.03 (−3.86, −0.14)
18 months	1.44 (−1.40, 4.07)	1.28 (−1.35, 3.85)	−1.34 (−3.33, 0.64)
24 months	0.04 (−2.58, 2.69)	0.10 (−2.61, 2.85)	−1.50 (−3.72, 0.63)

Table 3
Posterior means and 95% *credible intervals* (*in parentheses*) *for the time- and compliance superclass-varying compliance probabilities assuming the average baseline HAMD of* 18.1 *and superclass probabilities under the MCC model*

Time	Increasing noncompliers	Erratic compliers	High compliers
4 months	0.66 (0.53, 0.80)	0.38 (0.00, 1.00)	0.99 (0.88, 1.00)
8 months	0.38 (0.20, 0.56)	0.83 (0.07, 1.00)	0.98 (0.86, 1.00)
12 months	0.19 (0.00, 0.40)	0.32 (0.00, 1.00)	0.99 (0.86, 1.00)
18 months	0.10 (0.02, 0.31)	0.93 (0.12, 1.00)	0.96 (0.76, 1.00)
24 months	0.02 (0.00, 0.07)	0.66 (0.00, 1.00)	0.88 (0.65, 1.00)
$P(U_i)$	0.42 (0.25, 0.56)	0.04 (0.00, 0.15)	0.54 (0.42, 0.72)

compliers). The third superclass consists of subjects who are highly compliant but become less compliant at the last follow-up visit (high compliers). More than half of the subjects are high compliers and about a quarter of subjects are low compliers, leaving decreasing compliers as the smallest superclass.

The log odds of compliance for every unit increase in the baseline HAMD and its 95% credible interval is 0.003(−0.04, 0.05) suggesting those with more severe depression at baseline (higher baseline HAMD) may be slightly more likely to comply with treatment assignment than those with less severe depression at baseline.

The within-superclass ITT contrasts of equation (7) are shown in Table 2. The contrasts suggest strong direct effect of randomization at the 4-month follow-up in the low-complier superclass, which consists of largely never-takers unlikely to meet with health specialists regardless of the treatment assigned. After the first year, only the high compliers randomized to the intervention group, who are still highly likely to meet with their health specialists, showed greater reduction in the HAMD than high compliers in usual care. None of the superclasses show strong ITT effects on depression after 2 years.

3.2 *Markov Compliance Class Model*

The MCC model relaxes the CI assumption of the time-varying compliance classes given compliance superclass and baseline covariates, and instead, assumes a first-order Markov structure for the time-varying compliance classes given compliance superclass. We present results under the three-compliance superclass model.

The log odds of compliance at 4 months adjusting for baseline HAMD are −0.52(−1.87, 0.81), −3.61(−15.56, 4.37), and 4.99(1.11, 13.69) for the first, second, and third superclass, respectively. This suggests that those in the first and second superclasses are less likely to comply with their treatment assignment whereas those in the third superclass are more likely to comply with their treatment assignment. Our model assumes that the association between the baseline HAMD and compliance at 4 months is the same across all three superclasses. The log odds of 4-month compliance for a unit increase in the baseline HAMD is 0.07(0.01, 0.13) suggesting that those with more severe depression are more likely to comply with treatment assignment.

Table 3 shows the time-varying compliance probabilities when we assume the average baseline HAMD score of 18.1. The first superclass consists of subjects who are likely to comply with assigned treatment at 4 months then compliance decreases over time (increasing noncompliers). The second superclass consists of subjects who exhibit erratic compliance behavior with abrupt increases and decreases in compliance probabilities (erratic compliers). The third superclass consists of subjects who are highly compliant then compliance decreased slightly during the last 6 months (high compliers). More than half of the subjects are high compliers, less than half are increasing noncompliers, and only a small portion are erratic compliers.

The transitional probabilities of the time-varying compliance classes within each superclass in Table 4 shows that increasing noncompliers and high compliers are more likely to stay in the complier class if they are in the complier class in the previous time point than if they are in the never-taker class then switch to the complier class. Subjects in the high-complier superclass are more likely to transition to the complier class than subjects in the increasing noncomplier superclass. We do not see any clear patterns in the transitional probabilities of the erratic compliers.

Table 4
Posterior means and 95% *credible intervals* (*in parentheses*) *of the transitional probabilities under the MCC model*

Superclass	j	$P(C_{i,j} = c \mid C_{i,j-1} = c, U_i)$	$P(C_{i,j} = c \mid C_{i,j-1} = n, U_i)$
Increasing noncomplier	2	0.57 (0.34, 0.77)	0.01 (0.00, 0.06)
	3	0.45 (0.00, 0.77)	0.01 (0.00, 0.03)
	4	0.27 (0.00, 1.00)	0.06 (0.02, 0.12)
	5	0.10 (0.00, 0.51)	0.02 (0.00, 0.05)
Erratic complier	2	0.67 (0.00, 1.00)	0.56 (0.00, 1.00)
	3	0.31 (0.00, 1.00)	0.48 (0.00, 1.00)
	4	0.64 (0.00, 1.00)	0.78 (0.00, 1.00)
	5	0.68 (0.00, 1.00)	0.54 (0.00, 1.00)
High complier	2	1.00 (0.99, 1.00)	0.15 (0.00, 1.00)
	3	1.00 (1.00, 1.00)	0.44 (0.00, 1.00)
	4	0.97 (0.84, 1.00)	0.54 (0.00, 1.00)
	5	0.91 (0.76, 1.00)	0.46 (0.00, 1.00)

Table 5
Posterior means and 95% *credible intervals* (*in parentheses*) *for the ITT contrasts of the outcome within compliance superclasses under the MCC model*

Time	Increasing noncompliers	Erratic compliers	High compliers
4 months	−5.19 (−7.33, −3.04)	−8.32 (−15.33, −0.76)	−1.46 (−3.05, −0.04)
8 months	−2.70 (−5.21, −0.34)	−1.39 (−4.71, 0.58)	−0.89 (−2.57, 0.77)
12 months	0.52 (−1.92, 3.13)	−0.01 (−3.41, 3.75)	−2.10 (−3.81, −0.37)
18 months	1.55 (−1.05, 4.23)	−1.28 (−3.29, 1.48)	−1.38 (−3.23, 0.50)
24 months	0.48 (−2.12, 2.95)	−1.31 (−4.57, 2.35)	−2.02 (−4.53, 0.11)

The posterior means and credible intervals of the within-compliance superclass ITT contrasts in equation (7) are displayed in Table 5 showing strong ITT effect at 4 months in the erratic compliers, which consists of mostly never-takers unlikely to meet with health specialists, suggesting direct effect of randomization. This direct effect seems to dissipate over time. We also see an ITT effect at 4 months in the high compliers, which consists of almost entirely compliers who are likely to meet with health specialists if assigned to the intervention, suggesting an effect of the intervention. Consistent with the results under the CI model, at the end of the first year we see greater decrease in HAMD in the high compliers assigned to the intervention than high compliers assigned to the usual care. It suggests that meeting with health specialists help improve depression, although none of the 95% credible intervals exclude 0 at the end of 2 years.

3.3 *Model Comparison*

Under the CI and the MCC compliance class structures we identified a superclass of high compliers, who are highly compliant with slight decrease in compliance at the last follow-up. We also identified a superclass with decreasing compliance, although the compliance probability under the CI model starts out much higher at 4 months and decreases at a faster rate over subsequent follow-ups than under the MCC model. Under the CI model we identified a superclass of subjects who are noncompliant, with no clear compliance trajectory. Under the MCC model we identified a superclass of subjects exhibiting erratic compliance behavior with fluctuating compliance probabilities and no clear trend in their compliance class transitions.

We saw similar within-compliance superclass ITT effects under both the CI and the MCC models. The ITT effects were larger in noncompliant subjects than compliant subjects at the 4-month follow-up, suggesting a direct effect of randomization early on. This is most evident in the low compliers under the CI model and the erratic compliers under the MCC model, both of which consist of mostly never-takers at 4 months. However, this direct effect seems to dissipate over time. At the end of 2 years we see the largest ITT effect in the high compliers under both the CI and the MCC models, which consist of mostly compliers.

Assessment of the fits of the PPD to the data using the χ^2-type statistics in equation (9) yields a PPD p-value of 0.0057 under the three-superclass CI model and 0.1549 under the MCC model, suggesting a better fit of the MCC model. The three-class MCC model also has a better fit than the two-class MCC model (PPD p-value =0.0089).

4. Discussion

Lin et al. (2008) proposed a CI model of the time-varying compliance classes that assumes the compliance classes within an individual are independent given compliance superclass and baseline covariates. In this article, we proposed a Markov model that assumes the compliance classes at each time point are dependent on the previous compliance behaviors, compliance superclass, and baseline covariates. The model also

accommodates possible nontransient ITT effects of previous treatment on the outcome using a decay parameter.

Under the MCC model we found those who are more depressed at baseline are more likely to comply with their assigned treatment at 4 months. The same trend was also found under the CI model. More depressed patients may be more eager to treat their depression and more likely to adhere to their prescribed treatment. Physicians may also monitor more depressed patients more closely, thus increasing treatment compliance.

The proposed MCC model provides information on how history of compliance relates to compliance behavior that was not considered in the CI model. People are creatures of habit—those who complied with the assigned treatment in the previous follow-up period were more likely to comply again than those who were noncompliant in the previous follow-up period.

We saw evidence of direct effect of randomization during the first 4 months; though in the long run, compliant subjects who were meeting with health specialists showed greater improvement in their depression than noncompliant subjects. The presence or availability of the health specialists may have had a positive impact on the patients' depression outcome initially regardless of whether they actually met, but to benefit from the intervention longitudinally, the patients had to have met the health specialists.

In our model, we assumed the potential outcomes are conditionally independent of the superclasses given compliance classes. The reviewers pointed out that a more parsimonious alternative would assume that the potential outcomes are conditionally independent of the compliance classes given the superclasses. However, from an interpretive point of view, it is easier to interpret compliance class-specific ITT estimates than to interpret superclass-specific ITT estimates. Additionally, the ITT effects within each of the compliance classes correspond to better estimators than do the ITT effects within the broader superclasses given that compliance classes at each time point provide more information than superclass alone.

Comparing the PPD to the data showed that the MCC model has a better fit than the CI model. In our future research, we plan to explore covariates that relate to compliance superclasses and time-varying compliance classes to further improve the fit of the MCC model.

Although the outcome model helps to identify the ITT effects within compliance classes under the normality and constant variance assumptions, if we have (i) only compliers and never-takers, and (ii) good pretreatment predictors of compliance, then a parametric outcome model is not necessary for identifiability of the ITT effects. In our application, we satisfy the first condition, but only weakly satisfy the second condition; hence, our results may be sensitive to the normality assumption. See Rubin (2006) for more discussion on identifiability of principal strata with parametric assumptions and covariates.

In a simulation study in Gallop et al. (under review) we found that results are sensitive to the violation of the homogeneous variance assumption when the sample size is small. Additional assumptions, such as the exclusion restriction (ER) assumption, may be needed to relax the homogeneous variance assumption. However, making the ER assumption may be unreasonable in the PROSPECT given we found possible direct effect of randomization. In our future work, we would like to explore alternative models to relax the homogeneous variance assumption.

Cheng and Small (2006) proposed a principal stratification method for a cross-sectional three-treatment arm trial. Following their strategy, with possible additional assumptions, such as the ER and the monotonicity assumptions, we can extend our proposed method to accommodate studies with more than two treatment arms. The number of possible compliance patterns increases exponentially with increasing numbers of active treatment arms and time points. Utilizing the superclasses may provide even greater benefit under these types of settings.

5. Supplementary Materials

The Web Appendix referenced in Section 2.5 is available under the Paper Information link at the *Biometrics* website `http://www.biometrics.tibs.org`.

Acknowledgements

JYL was supported by grant T32-MH065218, TRTH was supported by grant R01-MH61892, and MRE was supported by grant P30-MH066270. We thank the associate editor and the two reviewers for their insightful comments that helped to improve this manuscript.

References

Angrist, J. D., Imbens, G. W., and Rubin, D. B. (1996). Identification of causal effects using instrumental variables. *Journal of the American Statistical Association* **91,** 444–455.

Bruce, M. L., Ten Have, T. R., Reynolds, C. F., Katz, I. I., Schulberg, H. C., Mulsant, B. H., Brown, G. K., McAvay, G. J., Pearson, J. L., and Alexopoulos, G. S. (2004). Reducing suicidal ideation and depressive symptoms in depressed older primary care patients. *Journal of the American Medical Association* **291,** 1081–1091.

Cheng, J. and Small, D. S. (2006). Bounds on causal effects in three-arm trials with non-compliance. *Journal of the Royal Statistical Society, Series B* **68,** 815–836.

Collins, L. M. and Wugalter, S. E. (1992). Latent class model for stage-sequential dynamic latent variables. *Multivariate Behavioral Research* **27,** 131–157.

Frangakis, C. E. and Rubin, D. B. (1999). Addressing complications of intention-to-treat analysis in the combined presence of all-or-none treatment-noncompliance and subsequent missing outcomes. *Biometrika* **86,** 365–379.

Frangakis, C. E. and Rubin, D. B. (2002). Principal stratification in causal inference. *Biometrics* **58,** 21–29.

Frangakis, C. E., Brookmeyer, R. S., Varadhan, R., Safaeian, M., Vlahov, D., and Strathdee, S. A. (2004). Methodology for evaluating a partially controlled longitudinal treatment using principal stratification, with application to a needle exchange program. *Journal of the American Statistical Association* **99,** 239–249.

Gallop, R., Small, D., Lin, J. Y., Elliott, M. R., Joffe, M., and Ten Have, T. R. (under review). Mediation analysis with principal stratification.

Garrett, E. S. and Zeger, S. L. (2000). Latent class model diagnosis. *Biometrics* **56,** 1055–1067.

Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B. (2004). *Bayesian Data Analysis*, 2nd edition. New York: Chapman and Hall.

Goodman, L. A. (1974). Explanatory latent structure analysis using both identifiable and unidentifiable models. *Biometrika* **61,** 215–231.

Imbens, G. W. and Rubin, D. B. (1997). Bayesian inference for causal effects in randomized experiments with noncompliance. *The Annals of Statistics* **25,** 305–327.

Langeheine, R. and Van de Pol, F. (2002). Latent Markov chains. In *Applied Latent Class Analysis*, J. A. Hagenaars and A. L. McCutcheon (eds), 304–341. Cambridge, U.K.: Cambridge University Press.

Lin, J. Y., Ten Have, T. R., and Elliott, M. R. (2008). Longitudinal compliance class model in the presence of time-varying noncompliance. *Journal of the American Statistical Association* **103,** 462–473.

Little, R. J. and Rubin, D. B. (2000). Causal effects in clinical and epidemiological studies via potential outcomes: Concepts and analytical approaches. *Annual Review of Public Health* **21,** 121–145.

Peng, Y., Little, R. J. A., and Raghunathan, T. E. (2004). An extended general location model for causal inferences from data subject to noncompliance and missing values. *Biometrics* **60,** 598–607.

Reboussin, B. A., Liang, K. Y., and Reboussin, D. M. (1999). Estimating equations for a latent transition model with multiple discrete indicators. *Biometrics* **55,** 839–845.

Rubin, D. B. (1978). Bayesian inference for causal effects. *The Annals of Statistics* **6,** 34–58.

Rubin, D. B. (1986). Statistics and causal inference. Comment: Which ifs have causal answers. *Journal of the American Statistical Association* **81,** 961–962.

Rubin, D. B. (2006). Causal inference through potential outcomes and principal stratification: Application to studies with "censoring" due to death. *Statistical Science* **21,** 299–309.

Small, D. S., Ten Have, T. R., Joffe, M. M., and Cheng, J. (2006). Random effects logistic models for analysing efficacy of a longitudinal randomized treatment with non-adherence. *Statistics in Medicine* **25,** 1981–2007.

Ten Have, T. R., Elliott, M. R., Joffe, M., Zanutto, E., and Datto, C. (2004). Causal models for randomized physician encouragement trials in treating primary care depression. *Journal of the American Statistical Association* **99,** 16–25.

Yau, L. H. Y. and Little, R. J. (2001). Inference for the complier-average causal effect from longitudinal data subject to noncompliance and missing data, with application to a job training assessment for the unemployed. *Journal of the American Statistical Association* **96,** 1232–1244.

Received December 2006. *Revised May* 2008.
Accepted May 2008.