A Deterministic Gated Lognormal Response Time Model to Identify Examinees with Item
Preknowledge

Murat Kasli, University of Miami

Cengiz Zopluoglu, University of Miami

Sarah Toton, Caveon

Please address all correspondence to:

Murat Kasli

Department of Educational and Psychological Studies
University of Miami
Max Orovitz Building

1570 Levante Ave, Coral Gables, FL 33146

Email: muratkasli@miami.edu

Acknowledgement: The authors would like to thank Dr. Andy Mroch and Dr. Mark Albanese for their support and feedback.

This study was proposed in Nov 2018 for NCBE's Joe E. Covington Award and was completed and supported as part of NCBE's Joe E. Covington Award for Research on Testing for Licensure.

Abstract

Response time (RT) information has recently attracted a significant amount of attention in the literature as it may provide meaningful information about item preknowledge. In this study, a Deterministic Gated Lognormal Response Time (DG-LNRT) model is proposed to identify examinees with potential item preknowledge using RT information. The proposed model is applied to a real experimental dataset provided by Toton and Maynes (2019) in which item preknowledge was manipulated, and its performance is demonstrated. Then, the performance of the DG-LNRT model is investigated through a simulation study. The model is estimated using the Bayesian framework via Stan. The results indicate that the proposed model is viable and has the potential to be useful in detecting cheating by using response time differences between compromised and uncompromised items.

Keywords: response time, lognormal response time model, item preknowledge, cheating

A Deterministic Gated Lognormal Response Time Model to Identify Examinees with Item Preknowledge

There has been growing interest in recent years in computer-based test delivery systems. Instead of paper-and-pencil based tests, many programs now administer tests via computers, which comes with many advantages both for test takers and test developers such as reduced costs of testing, increased test security, and immediate score reporting (Marianti, Fox, Avetisyan, Veldkamp & Tijmstra, 2014). While computer-based testing (CBT) may have many benefits, test security remains to be a significant issue. In particular, item preknowledge is a severe threat to test score validity when the items are overused or when there is a systematic effort to steal the item bank. Item preknowledge has become a concern for many large-scale testing programs (Shu, Henson & Luecht, 2013). For instance, most recently, there were speculations that the SAT questions administered in August 2018 had been leaked and sold on some online platforms before the actual test administration (Jesse & Murphy, 2018). Similarly, the ACT was not administered as scheduled in October 2018 due to a potential compromise in the testing process (Jaschik, 2018). Item preknowledge would provide an unfair advantage to the people who cheated and lead to serious consequences when it comes to high-stakes decisions. Hence, the interest in developing statistical methods to identify item preknowledge, a particular type of fraud in testing, has increased to ensure the validity of test scores.

With the increasing use of CBT, many testing programs are now able to collect response time (RT) information on items. Analyzing response time information may have utility for identifying unusual test-taking behavior, specifically when examinees had access to items before testing. Eckerly (2017) provides a detailed review of methods developed and proposed in the last two decades for identifying examinees with advanced knowledge of items. A significant number

of these methods either solely use response time information (e.g., Boughton, Smith, & Ren, 2017; Lee, 2018; Qian, Staniewska, Reckase, & Woo, 2016; van der Linden & Krimpen-Stoop, 2003; van der Linden & Guo, 2008) or use response time information along with item responses (e.g., Wang, Xu, Shang, & Kuncel, 2018).

In this study, a Deterministic Gated Lognormal Response Time Model (DG-LNRT), a hybrid version of the Deterministic Gated Item Response Theory Model (DGM; Shu, Henson & Luecht, 2013) and Lognormal Response Time Model (LNRT; van der Linden, 2007), is proposed to identify examinees with potential item preknowledge using response time information. DGM uses raw item responses and hypothesizes two latent ability parameters: true ability and cheating ability. DGM assumes that examinees respond to items using their true ability when they don't have prior knowledge, and they respond to items using their cheating ability when they have prior knowledge. The proposed DG-LNRT model is structurally similar to the DGM with a small nuance. It uses the latent speed parameter obtained from observed item response times instead of the latent ability parameter obtained from observed item responses. DG-LNRT hypothesizes two latent speed parameters for each individual, a latent speed parameter for responding to items when there is prior knowledge, and a second latent speed parameter for responding to items when there is no prior knowledge. In the next section, the DG-LNRT model is introduced, and an estimation framework using Stan (Stan Development Team, 2017) is provided. Then, the proposed model is applied to real experimental data in which item preknowledge was manipulated, and model performance is evaluated based on the known status of test-takers. Then, the results from a simulation study are summarized to evaluate the effectiveness of the model under certain conditions.

Model Description

DG-LNRT is a combination of LNRT (van der Linden, 2006) and DGM (Shu, Henson & Luecht, 2013). In this hybrid model, two latent speed parameters are hypothesized for each examinee based on whether or not they respond to items they had seen before. DG-LNRT can be specified as

$$\ln(y_{ij}) | \boldsymbol{\tau}_i, \boldsymbol{\tau}_c, \boldsymbol{\beta}, \boldsymbol{\alpha}, \mathbf{T}, \mathbf{C} \sim N(\mu_{ij}, \sigma_i), \quad (1)$$

where μ_{ij} and σ_i are defined as the following

$$\mu_{ij} = (\beta_i - \tau_{ij})^{1-T_j} \times \left[(1 - C_i) \times (\beta_i - \tau_{ij}) + C_i \times (\beta_i - \tau_{cj}) \right]^{T_j} \quad (2)$$

$$\sigma_i = \frac{1}{\alpha_i^2}. \quad (3)$$

In this model, $\ln(y_{ij})$ represents the natural logarithm of observed response time for the j^{th} examinee on the i^{th} item, $\boldsymbol{\beta} = \{\beta_1, \beta_2, \dots, \beta_I\}$ and $\boldsymbol{\alpha} = \{\alpha_1, \alpha_2, \dots, \alpha_I\}$ represent vectors of time-intensity and time-discrimination parameters for I items, respectively, $\boldsymbol{\tau}_i = \{\tau_{i1}, \tau_{i2}, \dots, \tau_{iJ}\}$ represents a vector of true latent speed parameters for J examinees when they respond to uncompromised items, $\boldsymbol{\tau}_c = \{\tau_{c1}, \tau_{c2}, \dots, \tau_{cJ}\}$ represents a vector of cheating latent speed parameters for J examinees when they respond to compromised items, $\mathbf{T} = \{T_1, T_2, \dots, T_J\}$ represents a vector of indicator variables for examinee status about whether or not the examinees have item preknowledge (0: honest examinee, 1: cheater examinee), and $\mathbf{C} = \{C_1, C_2, \dots, C_I\}$ represents an indicator variable for the item status about whether or not they are compromised (0: uncompromised item, 1: compromised item). The model can be broken down to the following four conditional models:

$$\ln(y_{ij}) | \tau_t, \tau_c, \beta, \alpha, \mathbf{T}, \mathbf{C} \sim \begin{cases} N(\beta_i - \tau_{tj}, \frac{1}{\alpha_i^2}), & \text{when } T_j = 0 \text{ and } C_i = 0 \\ N(\beta_i - \tau_{tj}, \frac{1}{\alpha_i^2}), & \text{when } T_j = 0 \text{ and } C_i = 1 \\ N(\beta_i - \tau_{tj}, \frac{1}{\alpha_i^2}), & \text{when } T_j = 1 \text{ and } C_i = 0 \\ N(\beta_i - \tau_{cj}, \frac{1}{\alpha_i^2}), & \text{when } T_j = 1 \text{ and } C_i = 1 \end{cases} \quad (4)$$

where τ_{tj} is the latent speed parameter for examinee j when responding to an uncompromised item, τ_{cj} is the latent speed parameter for examinee j when responding to a compromised item, α_i and β_i are parameters for the time discrimination and time intensity of the i^{th} item, respectively. C_i is an indicator variable describing the i^{th} item status where $C_i = 1$ indicates that the i^{th} item has been compromised and $C_i = 0$ indicates that the i^{th} item has not been compromised. T_j is an indicator variable for the j^{th} examinee where $T_j = 1$ indicates that the j^{th} examinee has item preknowledge if $\tau_{cj} > \tau_{tj}$ and $T_j = 0$ indicates that the j^{th} examinee does not have item preknowledge if $\tau_{tj} > \tau_{cj}$.

A Real Data Application

Data Description

The GRE Quantitative Reasoning Test data from a published study by Toton and Maynes (2019) were used to test the performance of the DG-LNRT model. The dataset consisted of 25 items from an out-of-circulation paper-and-pencil GRE Quantitative Reasoning Test. The test was computerized to collect response time information. There were 93 participants randomly assigned to one control and two experimental conditions. In the control condition (N=33), participants were given 40 minutes to complete the test and had no prior knowledge of the items. In the first experimental condition (N=30), participants arrived 20 minutes before the test and were allowed to study 12 of the 25 items without the key answers. In the second experimental condition (N=30),

participants similarly arrived 20 minutes before the test and were allowed to study the same 12 items with the key answers.

Analysis

The DG-LNRT model was fit to the data using the Stan (Stan Development Team, 2017) code developed for the proposed DG-LNRT model (see Appendix A for Stan model syntax). The parameters of DG-LNRT were estimated using the Hamiltonian Monte Carlo (HMC) algorithm as implemented in the rstan package (Stan Development Team, 2016) in R (R Core Team, 2018). The following priors were chosen for the model parameters:

- $\tau_t \sim N(0,1)$,
- $\tau_c \sim N(0,1)$,
- $\beta \sim N(\mu_1, \sigma_1)$ with the hyperparameters $\mu_1 \sim N(0,1)$ and $\sigma_1 \sim \exp(1)$.
- $\alpha \sim \text{Inv} - \gamma(5,5)$.

The current form of the model assumes that the set of compromised items is known and $\mathbf{C}$, a vector of zeros and ones indicating item status (compromised or uncompromised), must be given during the model fitting process as data. The primary interest in the model is estimating T_j , examinee status (honest or cheater). As T_j is an indicator variable estimated to be equal to one at each random draw when the posterior cheating latent speed parameter (τ_{cj}) is greater than the posterior true latent speed parameter (τ_{tj}), the average of the posterior samples of T_j is used as a probability estimate of having item preknowledge for examinee j . This set up of the model is very similar to how Shu, Henson & Luecht (2013) formulated DGM. Three Markov chains with random starting values were used, and each chain had 10,000 iterations, with the first 5,000 iterations in each chain as a warm-up, and the remaining 5,000 iterations for the model parameter

estimation. The convergence diagnostic measured $\hat{R}$, modified by Brooks and Gelman (1998), was used to assess the convergence of every parameter.

Since the examinees who had prior knowledge and who had no prior knowledge were known in the experimental study, this allowed us to examine three outcomes to evaluate the performance of the model by comparing the predicted probability estimate from the model for each individual (T_j) to the actual status of the individual (have item preknowledge or not). These outcomes were (a) the proportion of incorrectly identified individuals with no item preknowledge that were predicted to have preknowledge (false-positive rate), (b) the proportion of correctly identified individuals with item preknowledge that were predicted to have preknowledge (true-positive rate), and (c) the proportion of correctly identified individuals among all identified individuals (precision). These outcomes were computed based on the number of individuals detected by the model when the model predicted probability (T_j) is above a certain cut-off level. The following cut-off levels were used in the study: 0.50, 0.60, 0.70, 0.80, 0.90, 0.95, and 0.99.

Results

The distribution of $\hat{R}$ values for all model parameters are displayed in Figure 1. The $\hat{R}$ values ranged from 0.999 to 1.03 suggesting good convergence for all of the parameters estimated in the model. The mean, standard deviation, and 95% credible intervals with endpoints of 2.5 and 97.5 percentiles of posterior distributions were reported in Table 1 for the item parameters of α and β . The estimates of the α parameters ranged from 0.91 to 3.91, with an average of 1.3 and a standard deviation of 0.55. The estimates of the β parameters ranged from 3.17 to 4.78, with an average of 3.86 and a standard deviation of 0.39.

When the model predicted probability was compared to the actual binary status for each individual in the experimental data, the area under the receiver operating characteristic curve

(AUROC) was computed as 0.99 indicating that the DG-LNRT model was highly capable of distinguishing the individuals with item preknowledge from the examinees without item preknowledge. The outcome measures of false-positive rates, true-positive rates, and precision values for varying threshold levels are reported in Table 2. As seen in Table 2, precision values ranged from 0.96 to 1.00, with an average of 0.992. Similarly, true-positive rates ranged from 0.83 to 0.92 with an average of 0.89, and false-positive rates were between 0.00 and 0.061 with mean value of 0.013. The model performed very well in separating individuals in the control and experimental groups based on the probability estimate of having item preknowledge despite the small sample size and the imbalance between the numbers of individuals in the control and experimental group that potentially yielded more contamination of responses than expected in real settings.

Figure 1. Distribution of $\hat{R}$ values for all model parameters

Table 1. Item Parameter Estimates from Hamiltonian Monte Carlo through Stan

Item	$\hat{\alpha}$					$\hat{\beta}$				
	M	SD	95% CI	ESS	$\hat{R}$	M	SD	95% CI	ESS	$\hat{R}$
1	1.32	0.05	(1.22,1.43)	7292	1.000	3.72	0.10	(3.52,3.92)	459	1.005
2	1.28	0.05	(1.18,1.39)	8254	1.001	3.82	0.10	(3.63,4.02)	522	1.004
3	1.43	0.06	(1.32,1.55)	8982	1.000	4.35	0.10	(4.15,4.54)	407	1.006
4	1.17	0.05	(1.07,1.26)	7930	1.000	3.54	0.11	(3.33,3.75)	624	1.004
5	1.41	0.06	(1.30,1.53)	8399	1.000	4.31	0.10	(4.12,4.51)	407	1.006
6	1.14	0.05	(1.05,1.23)	9458	1.000	3.95	0.11	(3.73,4.17)	633	1.003
7	1.33	0.05	(1.23,1.44)	8671	1.000	3.92	0.10	(3.72,4.12)	441	1.006
8	1.09	0.04	(1.00,1.17)	9663	1.001	3.78	0.12	(3.55,4.01)	684	1.004
9	1.36	0.06	(1.26,1.47)	9929	1.000	4.50	0.10	(4.30,4.70)	428	1.005
10	1.02	0.04	(0.94,1.10)	11352	1.000	3.57	0.12	(3.32,3.81)	826	1.002
11	1.48	0.06	(1.36,1.61)	8843	1.000	4.78	0.10	(4.60,4.97)	406	1.006
12	1.23	0.05	(1.14,1.33)	9313	1.000	3.65	0.10	(3.45,3.85)	555	1.004
13	1.26	0.05	(1.16,1.35)	9747	1.001	4.03	0.11	(3.82,4.24)	479	1.004
14	0.91	0.03	(0.85,0.98)	10885	1.000	4.17	0.14	(3.89,4.45)	928	1.003
15	1.15	0.05	(1.07,1.24)	9388	1.000	4.23	0.11	(4.01,4.45)	547	1.005
16	1.07	0.04	(0.98,1.15)	9768	1.000	3.95	0.12	(3.71,4.18)	818	1.003
17	1.24	0.05	(1.15,1.35)	10338	1.000	3.38	0.11	(3.17,3.59)	500	1.005
18	1.32	0.06	(1.20,1.43)	7985	1.000	3.33	0.10	(3.13,3.52)	510	1.004
19	1.13	0.05	(1.04,1.22)	10229	1.000	4.11	0.12	(3.88,4.34)	565	1.004
20	1.28	0.06	(1.17,1.39)	6878	1.000	3.17	0.10	(2.97,3.37)	551	1.004
21	1.17	0.05	(1.08,1.27)	9805	1.000	3.53	0.11	(3.31,3.76)	558	1.004
22	1.12	0.05	(1.03,1.21)	9610	1.000	3.68	0.12	(3.45,3.91)	726	1.003
23	1.02	0.04	(0.94,1.10)	10826	1.000	3.60	0.13	(3.34,3.86)	761	1.003
24	0.93	0.04	(0.86,1.00)	11239	1.000	3.58	0.14	(3.29,3.86)	1083	1.002
25	0.94	0.04	(0.87,1.01)	9658	1.000	3.91	0.15	(3.62,4.20)	807	1.004

Note. M, SD, and 95% CI is the mean, standard deviation, and 95% credible interval of the posterior distribution. ESS is the effective sample size for the inferences based on the posterior distribution. ESS is the estimated number of independent draws after taking the autocorrelations into account. $\hat{R}$ is the potential scale reduction factor.

Table 2. Outcome Measures for the Varying Threshold Values

Cut-off Threshold	True Positive Rate ^a	False Positive Rate ^b	Precision
0.50	0.9167	0.0606	0.9649
0.60	0.9167	0.0303	0.9821
0.70	0.9167	0	1.0000
0.80	0.9167	0	1.0000
0.90	0.8833	0	1.0000
0.95	0.8667	0	1.0000
0.99	0.8333	0	1.0000

^a True positive rate is the proportion of correctly identified individuals out of 60 individuals in two experimental groups.

^b False positive rate is the proportion of incorrectly identified individuals out of 33 individuals in the control condition.

Simulation Study

A simulation study was also conducted to demonstrate the performance of the proposed DG-LNRT model and its accuracy to detect examinees with item preknowledge in different conditions. The simulation study manipulated four variables: test length, percentage of compromised items, sample size, and percentage of examinees with item preknowledge. The manipulated levels for each variable are shown in Table 3. A fully-crossed design yielded a total number of 60 (3x5x2x2) conditions, and we ran 50 replications within each condition.

Table 3. Simulation Design Conditions

Conditions	Levels
Test Length (N_i)	30,50 and 100
Percentage of compromised items (PCI)	10%,20%,30%,40% and 50%
Sample Size (N_p)	500 and 1000
Percentage of examinees with item preknowledge ($PEIP$)	5% and 10%

We used the following parameter specifications:

- α were normally distributed with a mean value of 4.02 and SD of 0.32,
- β were normally distributed with a mean value of 1.45 and SD of 0.10.
- latent speed parameters τ followed a multivariate normal distribution with a mean vector

$$\mu = [0 \quad 0.4] \text{ and covariance matrix } \Sigma = \begin{bmatrix} 1.0 & 0.8 \\ 0.8 & 1.0 \end{bmatrix}.$$
¹

This set of parameters assumed that the cheating latent speed parameter was, on average, 0.40 standard deviation above the true latent speeding parameter, and there was a positive correlation of 0.8 between the true and cheating latent speed parameters. In each replication, a random sample of items was selected for the compromised set of items based on the percentage of compromised items in the condition (e.g., five items were randomly selected as compromised items for the condition with 50-item test length and 10% compromised items), and a vector of zeros and ones for item status (**C**) was formed accordingly. Similarly, in each replication, a random sample of simulees was selected as having item preknowledge based on the percentage of examinees with item preknowledge and sample size in the condition (e.g., 25 simulees were randomly selected as having item preknowledge for the condition in which the sample size was 500 and percentage of examinees with item preknowledge was 5%), and a vector of zeros and ones for person status (**T**) was formed accordingly. After all the model parameters (α , β , τ) were generated based on the specifications described above, response time data for each individual on each item was generated based on Equations 1, 2, and 3. After data generation, the DG-LNRT model was fitted to each

¹ Readers may be curious about why the author(s) did not use the same parameter values obtained from real data demonstration. First, the author(s) were not aware of the existence of this real experimental data when the simulation study was designed and executed. Second, the experimental data included 60 individuals with item preknowledge and 33 individuals with no item preknowledge, so item parameters were probably highly biased because of the large amount of contamination due to item preknowledge. Item parameters used in the simulation study were obtained from another data made available by Cizek and Wollack (2017).

replication within each condition using Stan, as described above in the real data section. In the model fitting process, the chain length was set to 10,000, and the first 5,000 steps served as the warm-up period. After the model was fitted, the main parameter of interest was the probability of having item preknowledge estimated for each examinee. The average from the posterior distribution of T_j for each examinee was used as an estimate of the probability of having item preknowledge. The simulated examinee was flagged if T_j was above 0.95.

Similar to the real study, three outcomes were examined to evaluate the performance of the model by comparing the predicted probability outcome from the model for each simulated individual's status to the actual outcome of the simulated individual. These outcomes were (a) the proportion of incorrectly identified individuals with no item preknowledge that were predicted to have preknowledge (false-positive rate), (b) the proportion of correctly identified individuals with item preknowledge that were predicted to have preknowledge (true-positive rate), and (c) the proportion of correctly identified individuals among all identified individuals (precision). To examine the effect of simulation conditions on outcomes, the outcome measures were analyzed using a between-subjects factorial ANOVA after using an Arcsine transformation because all outcome measures were proportions and Arcsine transformation was used to stabilize variances and to normalize the outcome measures.

Results

The average false-positive rate, true-positive rate, and precision across 50 replications for each condition are reported in Table 4, Table 5, and Table 6, respectively. The eta-squared values from the ANOVA analyses of each outcome measure are reported in Table 7. The results for each outcome measure are separately discussed below.

Table 4. Summary Conditions – False-positive Rate ^{a,b}

PEIP	PCI	Np=500			Np=1000		
		Ni=30	Ni=50	Ni=100	Ni=30	Ni=50	Ni=100
0.05	0.10	0.0117	0.0085	0.0044	0.0096	0.0064	0.0055
0.05	0.20	0.0075	0.0035	0.0027	0.0071	0.0050	0.0029
0.05	0.30	0.0043	0.0035	0.0019	0.0067	0.0040	0.0017
0.05	0.40	0.0053	0.0025	0.0004	0.0045	0.0011	0.0026
0.05	0.50	0.0044	0.0043	0.0018	0.0042	0.0031	0.0024
0.10	0.10	0.0082	0.0091	0.0052	0.0101	0.0070	0.0060
0.10	0.20	0.0071	0.0044	0.0022	0.0076	0.0054	0.0037
0.10	0.30	0.0053	0.0024	0.0020	0.0046	0.0040	0.0025
0.10	0.40	0.0058	0.0022	0.0026	0.0054	0.0041	0.0024
0.10	0.50	0.0055	0.0049	0.0019	0.0048	0.0029	0.0025

^a PEIP, PCI, Np, and Ni represent the percentage of examinees with item preknowledge, percentage of compromised items, number of people, and number of items, respectively.

^b The numbers reported in this table are the average across 50 replications for each condition.

Table 5. Summary Conditions – True-positive Rate ^{a,b}

PEIP	PCI	Np=500			Np=1000		
		Ni=30	Ni=50	Ni=100	Ni=30	Ni=50	Ni=100
0.05	0.10	0.298	0.374	0.395	0.310	0.342	0.484
0.05	0.20	0.368	0.472	0.512	0.400	0.468	0.513
0.05	0.30	0.434	0.482	0.577	0.481	0.526	0.515
0.05	0.40	0.437	0.552	0.610	0.404	0.540	0.566
0.05	0.50	0.402	0.451	0.565	0.438	0.484	0.553
0.10	0.10	0.373	0.385	0.483	0.308	0.389	0.458
0.10	0.20	0.416	0.482	0.557	0.350	0.437	0.521
0.10	0.30	0.388	0.470	0.580	0.444	0.465	0.585
0.10	0.40	0.377	0.500	0.568	0.435	0.509	0.566
0.10	0.50	0.425	0.460	0.573	0.439	0.481	0.563

^a PEIP, PCI, Np, and Ni represent the percentage of examinees with item preknowledge, percentage of compromised items, number of people, and number of items, respectively.

^b The numbers reported in this table are the average across 50 replications for each condition.

Table 6. Summary Conditions – Precision ^{a, b}

PEIP	PCI	Np=500			Np=1000		
		Ni=30	Ni=50	Ni=100	Ni=30	Ni=50	Ni=100
0.05	0.10	0.603	0.698	0.821	0.637	0.741	0.821
0.05	0.20	0.727	0.878	0.912	0.749	0.831	0.902
0.05	0.30	0.838	0.885	0.946	0.790	0.870	0.943
0.05	0.40	0.812	0.935	0.989	0.819	0.964	0.921
0.05	0.50	0.825	0.851	0.941	0.850	0.893	0.924
0.10	0.10	0.830	0.830	0.911	0.771	0.856	0.894
0.10	0.20	0.872	0.925	0.968	0.837	0.899	0.940
0.10	0.30	0.894	0.958	0.969	0.914	0.928	0.963
0.10	0.40	0.887	0.962	0.962	0.900	0.934	0.964
0.10	0.50	0.899	0.914	0.972	0.909	0.949	0.961

^a PEIP, PCI, Np, and Ni represent the percentage of examinees with item preknowledge, percentage of compromised items, number of people, and number of items, respectively.

^b The numbers reported in this table are the average across 50 replications for each condition.

Table 7. Results of ANOVA

Predictor	η^2_{fpr}	η^2_{tpr}	η^2_p
Np	0.0148	0.0010	0.0078
Ni	0.1823	0.0742	0.3545
PEIP	0.0425	0.0390	0.1638
PCI	0.2626	0.1909	0.3600
Np x Ni	0.0233	0.0574	0.0023
Np x PEIP	0.0272	0.0207	0.0001
Ni x PEIP	0.0405	0.0255	0.0173
Np x PCI	0.0354	0.0248	0.0075
Ni x PCI	0.0744	0.0808	0.0238
PEIP x PCI	0.0495	0.1007	0.0181
Np x Ni x PEIP	0.0196	0.0464	0.0012
Np x Ni x PCI	0.0564	0.1331	0.0106
Np x PEIP x PCI	0.0433	0.0801	0.0041
Ni x PEIP x PCI	0.0607	0.0292	0.0116
Np x Ni x PEIP x PCI	0.0666	0.0918	0.0174

Note: PEIP, PCI, Np, and Ni represent the percentage of examinees with item preknowledge, percentage of compromised items, number of people, and number of items, respectively. η^2_{fpr} , η^2_{tpr} , and η^2_p represent the eta-squared values of false-positive rates, true-positive rates and precision, respectively.

Precision

Notice in Table 7, eta-squared values (η_p^2) indicated that the main effects of Ni and PCI had the largest effects on precision, each accounting for roughly 35-36% of the variance in precision. Similarly, the next largest main effect was for PEIP, which accounted for about 16% of the variance. Np and the interactions among variables in the model each accounted for less than 3% of the variance. Based on the largest eta-squared results, means of precision values were plotted for the corresponding main effects. As seen in Figure 2, the precision increased from 0.82 to 0.89 when the number of items increased from 30 to 50 and increased from 0.89 to 0.93 when the number of items increased from 50 to 100. Similarly, the precision increased from 0.84 to 0.91 when the percentage of examinees with item preknowledge increased from 5% to 10%. The precision values increased from 0.78 to 0.91 when the percentage of compromised items increased from 10% to 30%; however, it did not increase much for higher percentages of compromised items after leveling off.

True-positive Rate

According to the eta-squared values for the true-positive rates, the main effect of PCI had the largest effect, accounting for 19% of the variance. Also, the two-way interaction between PEIP and PCI and the three-way interaction between Np, Ni, and PCI were the next largest effect sizes, accounting for between about 10% and 13% of the variance in true-positive rates. As seen in the left panel of Figure 3, true-positive rates increased until PCI was increased up to 40% and then slightly decreased when PCI is equal to 50% for both levels of PEIP. Also, the interaction was disordinal. A higher level of PEIP (10%) yielded larger true-positive rates for lower levels of PCI while yielding lower true-positive rates for high levels of PCI. For the three-way interaction between Np, Ni, and PCI, as shown in the right panel of Figure 3, PCI had a relatively positive

effect on true-positive rates, and the changes in true-positive rates were varying based on the conditions. For the 30-item condition, true-positive rates were higher for the sample size of 500 ($N_p=500$) when the percentage of compromised items were between %10 and %20; however, the model was yielding higher true-positive rates for the sample size of 1000 when the percentage of compromised items were %30 or more.

Similarly, for the 50-item condition, the model had higher true-positive rates for the sample size of 500 when the percentages of compromised items were between %10 and %20. However, the model yielded higher true-positive rates when the sample size was equal to 1000 for the higher percentages of compromised items (e.g., %30, %40, and %50). For the 100-item condition, the pattern of relationships changed compared to the 30-item and 50-item test conditions. The model was yielding higher true-positive rates for the sample size of 1000 when the percentage of compromised items was between %10 and %20. However, the model had higher true-positive rates for the sample size of 1000 when the percentage of compromised items was %30 or more.

False-Positive Rate

Eta-squared values for the false-positive rate indicated that the main effects for N_i and PCI had the largest effect sizes, accounting for roughly 18% to 26% of the variance in false-positive rates, respectively. Moreover, as seen in Figure 4, false-positive rates decreased as N_i and PCI increased.

Figure 2. Main Effect Plots for Precision

Figure 3. Interaction Effect Plots for True-positive Rates

Figure 4. Main Effect Plots for False-positive Rates

Discussion

In this study, a model using response time information was proposed to identify examinees with item preknowledge. The model structure and estimation process of DG-LNRT using a Bayesian framework through Stan was presented. DG-LNRT incorporates a cheating latent speed ability in addition to true latent speed ability and classifies examinees with or without item preknowledge through an indicator variable T estimated from the difference in response time between the compromised and uncompromised items. The performance of the DG-LNRT model was illustrated by analyzing the GRE Quantitative Reasoning Test data of Toton and Maynes (2019). The findings from the real data analyses indicated that the model proposed here provided very high true-positive rates and precision while holding the false-positive rates under control.

To evaluate the effectiveness of the proposed model in a broader setting, a simulation study was also conducted. The results from the simulation study indicated that outcome measures changed based on the simulation conditions. The DG-LNRT model yielded better accuracy with larger sample sizes and larger numbers of items. For instance, when the sample size equals to 500, the false-positive rates ranged from 0.0019 to 0.0117 while the number of items is increasing. Similarly, for the sample size of 500, true-positive rates ranged from 0.298 to 0.573, and precision values ranged from 0.603 to 0.972. Additionally, for the sample size of 1000, ranges of the false-positive rate, true-positive rate, and precision were between 0.0011 and 0.0101, 0.308 and 0.585, and 0.637 and 0.964, respectively. The low levels of false-positive rates and particularly high levels of precision indicate that the model was able to distinguish the examinees with and without knowledge, and successfully recovered the examinees with item preknowledge while minimizing the number of incorrectly identified honest individuals.

Also, there were differences in outcome measures based on the percentages of examinees with item preknowledge (PEIP) and percentages of compromised items (PCI). The results from the ANOVA suggested that the PCI main effect accounted for the largest percentage of variance in each outcome measure. The main effects of PEIP and N_i also accounted for a relatively large proportion of variance in precision, and N_i also accounted for a relatively large proportion of the variance in the false-positive rate. The interaction effects of $PEIP \times PCI$ and $N_p \times N_i \times PCI$ explained relatively large percentages of variance in true-positive rates. Overall, the results from the simulation study demonstrate that the DG-LNRT model is viable and has the potential to be useful in detecting cheating by using response time differences between compromised and uncompromised items.

There are several limitations to this study. First, the simulation study assumed that the standardized effect size of the difference in latent speed parameter between examinees with or without item preknowledge was 0.4. This difference was estimated from an observational dataset provided by Cizek and Wollack (2017) in a separate setting. The dataset from Cizek and Wollack (2017) includes some examinees suspected to have item preknowledge; however, it is likely that this dataset has honest examinees incorrectly identified as cheaters or cheater examinees incorrectly identified as honest examinees. In addition, the true nature of testing fraud is not known for this dataset. Therefore, the difference in latent speed parameters between examinees with or without item preknowledge may be different in real practice, and different effect size values may yield different results. For instance, smaller effect size values would probably yield more inaccurate results, while a larger effect size would probably yield more accurate results. Future studies should attempt to estimate the difference in latent speed parameters between examinees

with or without item preknowledge using experimental datasets such as Toton and Maynes (2019), and these studies can provide us better insight for using more realistic simulations.

Further exploration of the performance of the DG-LNRT model manipulating the effect size of difference in response times would help address model performance under conditions different from those of the current study. Second, in this model, we assume that compromised items are known, and by using this assumption, we tried to detect examinees with item preknowledge. However, in practice, test companies may not be able to have the information about which set of items is compromised. Therefore, the current model proposed in this study can be expanded in future studies by relaxing the assumption that the item compromised status is known.

References

- Boughton, K. A., Smith, J., & Ren, H. (2017a). Using response time data to detect compromised items and/or people. *Handbook of quantitative methods for detecting cheating on tests*, 177-192.
- Brooks, S. P. and Gelman, A. (1998). General methods for monitoring convergence of iterative simulations. *Journal of Computational and Graphical Statistics*, 7(4), 434-455.
- Cizek, G. J., & Wollack, J. A. (2017). *Handbook of quantitative methods for detecting cheating on tests*. New York, NY: Routledge.
- Eckerly, C. (2017). Detecting preknowledge and item compromise: Understanding the status quo. In G. J. W. Cizek, J.A. (Ed.), *Handbook of quantitative methods for detecting cheating on tests* (pp. 101-123).
- Jaschik, S. (2018, October 1). ACT calls off Oct. 2 test. Retrieved from <https://www.insidehighered.com/admissions/article/2018/10/01/act-calls-oct-2-test-over-apparent-security-issue>
- Jesse, D. & Murphy, K. (2018, August 28). Reports of leaked SAT leave students wondering about integrity of test. Retrieved from <https://www.usatoday.com/story/news/2018/08/28/august-sat-test-reuse-cheating-college-board-october-2017/1120558002/>
- Lee, S. Y. (2018). *A Mixture Model Approach to Detect Examinees with Item Preknowledge*. The University of Wisconsin-Madison.
- Marianti, S., Fox, J.-P., Avetisyan, M., Veldkamp, B. P., & Tijmstra, J. (2014). Testing for aberrant behavior in response time modeling. *Journal of Educational and Behavioral Statistics*, 39, 426–451. <https://doi.org/10.3102/1076998614559412>
- Qian, H., Staniewska, D., Reckase, M., & Woo, A. (2016). Using Response Time to Detect Item Preknowledge in Computer-Based Licensure Examinations. *Educational Measurement: Issues and Practice*, 35(1), 38-47.
- R Core Team. (2018). R: A language and environment for statistical computing [Computer software]. Vienna, Austria: R Foundation for Statistical Computing. Available from <http://www.R-project.org>
- Stan Development Team (2017). *Stan Modeling Language Users Guide and Reference Manual, Version 2.16.0*. URL <http://mc-stan.org>.
- Stan Development Team (2016). “**rstan**: The R interface to Stan.” R package version 2.14.1, URL <http://mc-stan.org/>.

- Shu, Z., Henson, R., & Luecht, R. (2013). Using deterministic, gated item response theory model to detect test cheating due to item compromise. *Psychometrika*, 78, 481-497.
- Toton, S. L., & Maynes, D. D. (2019). Detecting Examinees with Pre-Knowledge in Experimental Data using Conditional Scaling of Response Times. In *Frontiers in Education* (Vol. 4, p. 49). Frontiers.
- van der Linden, W. J. (2006). A lognormal model for response times on test items. *Journal of Educational and Behavioral Statistics*, 31, 181–204.
- van der Linden, W. J., & Guo, F. (2008). Bayesian procedures for identifying aberrant response-time patterns in adaptive testing. *Psychometrika*, 73, 365–384.
- van der Linden, W. J. & van Krimpen&Stoop, E. M. L. A. (2003). Using response times to detect aberrant responses in computerized adaptive testing. *Psychometrika*, 68 (2), 251-265.
- Wang, C., Xu, G., Shang, Z., & Kuncel, N. (2018). Detecting Aberrant Behavior and Item Preknowledge: A Comparison of Mixture Modeling Method and Residual Method. *Journal of Educational and Behavioral Statistics*, 1076998618767123.

APPENDIX A

Stan Code for the DG-LNRT Model

```

code_rt<-'
data{
  int <lower=1> I;                // number of individuals
  int <lower=1> J;                // number of items
  int <lower=1> n_obs;            // number of observed responses
  int <lower=1> ind_person_obs[n_obs]; // person indicator of a response
  int <lower=1> ind_item_obs[n_obs];  // item indicator of a response
  int <lower=0,upper=1> i_status_obs[n_obs]; // item status indicator of
    a response
  real Y[n_obs];                // vector of response times
}
parameters {
  vector[J] beta;               // item specific time intensity parameter
  vector <lower=0> [J] alpha;    // item specific residual standard deviation,
  vector[I] tau_t;              // person specific latent working speed
  vector[I] tau_c;              // person specific latent working speed
  real mul;
  real <lower=0> signal;
}
transformed parameters {
  vector[I] T;
  for (i in 1:n_obs) {
    if(tau_t[ind_person_obs[i]]>tau_c[ind_person_obs[i]])
      T[ind_person_obs[i]] = 0;
    else
      T[ind_person_obs[i]] = 1; }}
model{
  tau_t ~ normal(0,1);
  tau_c ~ normal(0,1);
  mul ~ normal(0,1);
  signal ~ exponential(1);
  beta ~ normal(mul,signal);
  alpha ~ inv_gamma(5,5);

  for (i in 1:n_obs) {
    real p_t = beta[ind_item_obs[i]]-tau_t[ind_person_obs[i]];
    real p_c = beta[ind_item_obs[i]]-tau_c[ind_person_obs[i]];
    real p = (p_t^(1-T[ind_person_obs[i]]))*
      (((1-i_status_obs[i])*p_t +
        (i_status_obs[i])*p_c)^T[ind_person_obs[i]]);
    Y[i] ~ normal(p,1/(alpha[ind_item_obs[i]]^2));}}'

rt_stan <- stan(model_code = code_rt, data = data_rt,iter =10000, chains = 3,
control = list(max_treedepth = 15), cores = getOption("mc.cores", 4L))

```