

LSAC RESEARCH REPORT SERIES

- **Joint Modeling of Ability and Differential Speed Using Responses and Response Times**

Jean-Paul Fox

University of Twente, Enschede, the Netherlands

- **Law School Admission Council
Research Report 16-04
March 2016**

The Law School Admission Council (LSAC) is a nonprofit corporation that provides unique, state-of-the-art products and services to ease the admission process for law schools and their applicants worldwide. Currently, 222 law schools in the United States, Canada, and Australia are members of the Council and benefit from LSAC's services. All law schools approved by the American Bar Association are LSAC members. Canadian law schools recognized by a provincial or territorial law society or government agency are also members. Accredited law schools outside of the United States and Canada are eligible for membership at the discretion of the LSAC Board of Trustees; Melbourne Law School, the University of Melbourne is the first LSAC-member law school outside of North America. Many nonmember schools also take advantage of LSAC's services. For all users, LSAC strives to provide the highest quality of products, services, and customer service.

Founded in 1947, the Council is best known for administering the Law School Admission Test (LSAT[®]), with about 100,000 tests administered annually at testing centers worldwide. LSAC also processes academic credentials for an average of 60,000 law school applicants annually, provides essential software and information for admission offices and applicants, conducts educational conferences for law school professionals and prelaw advisors, sponsors and publishes research, funds diversity and other outreach grant programs, and publishes LSAT preparation books and law school guides, among many other services. LSAC electronic applications account for nearly all applications to ABA-approved law schools.

© 2016 by Law School Admission Council, Inc.

All rights reserved. No part of this work, including information, data, or other portions of the work published in electronic form, may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or by any information storage and retrieval system, without permission of the publisher. For information, write: Communications, Publishing, and Creative Services, Law School Admission Council, 662 Penn Street, PO Box 40, Newtown, PA, 18940-0040.

This study is published and distributed by LSAC. The opinions and conclusions contained in this report are those of the author(s) and do not necessarily reflect the position or policy of LSAC.

Table of Contents

Executive Summary	1
Introduction	1
The Variable Working Speed Model	3
The Lognormal Random Linear Variable Speed Model	5
The Lognormal Random Quadratic Variable Speed Model	6
Joint Model for Responses and Response Times.....	7
Identification	8
Parameter Estimation	9
Simulation Study	11
Modeling Variable Speed in the Amsterdam Chess Test Data	13
Discussion	18
References	20

Executive Summary

With computerized testing, it is possible to record both the responses of test takers to test questions (i.e., items) and the amount of time spent by a test taker in responding to each question. Various models have been proposed that take into account both test-taker ability and working speed, with many models assuming a constant working speed throughout the test. The constant working speed assumption may be inappropriate for various reasons. For instance, a test taker may need to adjust their pace due to time mismanagement, or a test taker who started out working too fast may reduce their working speed to improve their accuracy. A model is proposed here that allows for variable working speed. An illustration of the model using the Amsterdam Chess Test data is provided.

Introduction

Responses to items in a test reveal information about the accuracy of the responses (i.e., the degree of correctness), which is related to the ability being measured. With the introduction of computer-based testing, both responses and responses times (RTs), or the amount of time taken to respond to produce an answer, can be collected. RTs will reveal information about the working speed of the respondent. Traditionally, in psychological research, a speed–accuracy trade-off applies, with fast-working test takers often producing a greater number of incorrect responses compared to test takers who work slower. This is referred to as a *within-person relationship* between speed and ability, and in educational research this relationship is assumed (e.g., van der Linden, 2007). When speed and ability are assumed to be constant, however, no such relationship can be studied. The *between-person relationship* between ability and speed has also been studied, building on the information that differentiates one test taker from another. Various studies have reported a negative correlation, estimated at the population level, between working speed and ability of test takers. Empirical examples of Klein Entink, Fox, and van der Linden (2009) showed that higher-ability test takers tended to work at a slower speed than lower-ability test takers. Klein Entink, Kuhn, Hornke, and Fox (2009), Roberts and Stankov (1999), and van der Linden and Fox (2015) also have reported a negative correlation between ability and speed in their empirical examples.

In the common lognormal RT model of van der Linden (2006), it is assumed that the working speed of a test taker is constant throughout the test. The general item response theory (IRT) models are based on the principle that a test taker will use his or her cognitive knowledge to respond to the test items. Therefore, the relationship between ability and speed is assumed to be constant for each test taker, along with a constant working speed.

The assumption of a constant (latent) speed parameter corresponds to the assumption of a constant (latent) ability. However, it is reasonable to assume that test takers will vary their working speed during a test. Changes in time management could be required to finish the test in time, or test takers could decide to work slower to improve their level of accuracy. Working speed can also vary when test takers show aberrant response behavior, such as cheating or guessing (Marianti, Fox, Avetisyan, Veldkamp, & Tijmstra, 2014).

Evidence of variable working speed can also be found in psychological testing, where test takers are asked to do different performance tasks. Manipulating the

experimental conditions results in changes in test takers' response behavior. For example, test takers will work faster when the time pressure is increased, but the level of accuracy may or may not change; the speed–accuracy trade-off is different across different levels of the time-pressure condition. For example, Vandekerckhove, Tuerlinckx, and Lee (2011) defined a hierarchical diffusion model for two-choice RTs, where the parameters of the response process can vary across persons, items, and experimental conditions to model the underlying response process. Assink, van der Lubbe, and Fox (2015) used the hierarchical drift diffusion model to identify tunnel vision (i.e., tendency to focus exclusively on a limited view) due to time pressure. In an experiment, they found an interaction effect between the time-pressure condition and the RT, but not the response accuracy.

In educational testing, different joint models for ability and speed assume a constant speed parameter for persons. The hierarchical latent variable modeling of responses and RTs (Fox, Klein Entink, & van der Linden, 2007; Klein Entink et al., 2009; van der Linden, 2007; van der Linden & Glas, 2010) and the generalized linear IRT approach (Molenaar, Tuerlinckx, & van der Maas, 2015b) both assume a constant latent working speed parameter for each individual. The constant speed parameter is also assumed in the IRT approach of categorical RTs (e.g., De Boeck & Partchev, 2012; Partchev & De Boeck, 2012; Ranger & Kuhn, 2012) and the nonlinear regressions of IRT parameters on RTs (Ferrando & Lorenzo-Seva, 2007a, 2007b).

To model nonconstant working speed, a latent growth modeling approach is defined for the speed parameter. For each test taker, the within-person systematic differences in observed RTs conditional on the time intensities (i.e., the population-average time needed to complete each item) are modeled using latent variable modeling. An individual speed process is assumed, describing the changes in speed across items. So, test takers can work with different levels of speed during the test. Each individual speed process will be defined using random effects to model correlations between the RTs for each test taker. A linear (within-person) relationship is defined between individual RTs and random effects.

Furthermore, random effects are also used to define differences in speed process between test takers. This will generalize to the common lognormal speed model, where a random intercept is used to define differences in speed across test takers. The latent growth speed process will be a second level of the lognormal speed model.

In latent growth curve analysis, a time scale is needed to model the speed process and to define the individual variation in initial status and growth rate. In the present approach, the order in which the items are solved will define the underlying time scale describing the sequence of observed item RTs. Each item functions as a measurement occasion for speed, and each pattern of RTs is treated as longitudinal RT data with respect to the speed process. The measurement occasions (defining the time scale of the speed process) are defined on a scale from 0 to 1. The chosen time scale values are arbitrary and only represent the order in which the items are solved and reflect that the observations are made at equidistant time points. The time variable will be defined on this scale, where the first (last) measurement corresponds to the response to the first (last) item.

It will be shown that the latent growth model for working speed can be integrated with an IRT model for ability. Under the variable speed model, the ability parameter is influenced by the speed process parameters. This generalizes the univariate relationship between ability and a single speed variable within a test, since multiple

speed components are involved in this relationship. In this approach, ability will be influenced by a weighted average of the person-specific speed process parameters.

The Markov chain Monte Carlo (MCMC) method is used for parameter estimation, which enables joint estimation of all model parameters. The developed MCMC method is built on the estimation methods of Klein Entink et al. (2009) and Fox (2010), who developed MCMC schemes for joint models for responses and RTs assuming a constant working speed model.

Simulated and real data examples are given to illustrate the modeling framework. Data from the Amsterdam Chess Test (ACT; van der Maas & Wagenmakers, 2005) are used to model variable working speed using a linear and a quadratic speed component. A direct comparison is made with the hierarchical model of van der Linden (2007) and Klein Entink et al. (2009).

The Variable Working Speed Model

Van der Linden (2006, 2007) proposed a lognormal model for RTs using two parameters to describe item and individual variations in RTs. An item factor is defined that represents the time intensity of an item, and each time-intensity parameter represents the population-average time needed to complete the item. A person parameter is defined that represents the constant working speed as the systematic differences in RTs given the time intensities. For example, a test taker works more slowly (faster) than the average level in the population when the differences between RTs and time intensities are all positive (negative), since over items more (less) time is needed than the population-average time.

Let T_{ik} denote the RT of person i ($i = 1, \dots, N$) on item k ($k = 1, \dots, K$). A lognormal RT distribution is assumed, to account for the positively skewed characteristic of RT distributions, which leads to

$$\ln T_{ik} = \lambda_k - \zeta_i + \varepsilon_{ik}, \quad \varepsilon_{ik} \sim N(0, \sigma_{\varepsilon_k}^2). \quad (1)$$

The time-intensity parameter is represented by λ_k and the common speed parameter by ζ_i . The time-intensity parameter represents the population-average time needed to complete the item. The test takers are assumed to be randomly selected from a population. Therefore, the speed parameter is assumed to follow a normal population distribution

$$\zeta_i \sim N(\mu_{\zeta}, \sigma_{\zeta}^2). \quad (2)$$

In Fox et al. (2007) and Klein Entink et al. (2009), a time-discrimination parameter ϕ_k has been included as a slope parameter for speed. The time-discrimination parameter characterizes the sensitivity of the item to different speed levels of the test takers. This leads to the following specification of the lognormal speed model:

$$\ln T_{ik} = \lambda_k - \phi_k \zeta_i + \varepsilon_{ik}, \quad \varepsilon_{ik} \sim N(0, \sigma_{\varepsilon_k}^2). \quad (3)$$

From Equation (3), it follows that the time-discrimination parameter is also used to model the unexplained heterogeneity between time-pattern responses. This follows from the fact that the covariance between the response to item k and the response to item l of person i includes the time discriminations, which are given by

$$\begin{aligned}
\text{cov}(T_{ik}, T_{il}) &= \text{cov}(\lambda_k - \phi_k \zeta_i + \varepsilon_{ik}, \lambda_l - \phi_l \zeta_i + \varepsilon_{il}) \\
&= \text{cov}(\phi_k \zeta_i + \varepsilon_{ik}, \phi_l \zeta_i + \varepsilon_{il}) \\
&= \text{cov}(\phi_k \zeta_i, \phi_l \zeta_i) \\
&= \phi_k \text{var}(\zeta_i) \phi_l = \phi_k \sigma_\zeta^2 \phi_l.
\end{aligned} \tag{4}$$

Van der Linden (2015) defined the time-discrimination parameter to be a measurement error variance parameter such that $\sigma_\varepsilon^2 = 1/\phi_k$. In that case, the time discrimination (or error variance parameter) will not influence the covariance between RTs.

In latent growth modeling, a time scale is required for the observed responses. The times that the observations were made are represented on this scale. From this perspective, the items in a test can be viewed as measurement occasions for speed and ability. Therefore, a natural time scale can be defined from the collected RTs, since each RT also defines the time between two subsequent response observations. Subsequently, the RTs of each test taker, in the order in which the items were solved, also define the time between his or her measurement occasions. This time scale would be inappropriate if the test taker were permitted to take a break after finishing an item and before moving on to the next question.

When estimating all model parameters simultaneously, generating a unique timeline for each test taker increases the computational burden. Furthermore, the latent growth model is applied to a test situation that might take a few hours, such that the non-equidistant property of the time scale can hardly influence the results. Therefore, an equidistant timeline is defined, which leads to a common time scale for all test takers.

Let a time variable $\mathbf{X}_i = X_{i1}, X_{i2}, \dots, X_{iK}$, where $X_{i1} = 0$ represents the measurement occasions of test taker i . This time variable represents only the order of the items, which is used to model the speed process over time on an equidistant scale. The measurement of speed from the first item observation is defined as the intercept, and subsequent item observations can be used to model change in speed. Let $\mathbf{X}_{(i)} = X_{(i1)}, X_{(i2)}, \dots, X_{(iK)}$ denote the order in which the K items are made by person i . Then, a convenient time scale is defined by $X_{ik} = (X_{(ik)} - 1) / K$. The times are defined on scale from 0 to 1, where 1 is the upper bound representing an infinite number of items.

Note that the scale on which the latent variable working speed is measured is arbitrary. Therefore, the numerical values of the time scale for the speed process only need to address the order in which the items were solved and the assumed equidistant property of the measurements.

The Lognormal Random Linear Variable Speed Model

To introduce the latent growth model for speed, the lognormal RT model is extended with a linear growth term. This model will not be of particular interest in itself, since it is not realistic to assume that test takers will accelerate their speed of working in a linear way. However, the linear trend component can be used in combination with higher-order time components to model more complex processes of working speed.

The lognormal RT model with a linear trend for speed can be defined using the time variable X ; it follows that

$$\ln T_{ik} = \lambda_k - \phi_k (\zeta_{i0} + \zeta_{i1} X_{ik}) + \varepsilon_{ik}$$

$$\begin{pmatrix} \zeta_{i0} \\ \zeta_{i1} \end{pmatrix} \sim N \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_{\zeta_0}^2 & \rho_{\zeta_{01}} \\ \rho_{\zeta_{01}} & \sigma_{\zeta_1}^2 \end{pmatrix} \right). \quad (5)$$

The parameter ζ_{i0} represents the value of speed measured with the first item solved, also referred to as the *initial value of speed*. The parameter ζ_{i1} represents the random slope in speed, which means that test takers can differ in their growth rate of speed. Note that both random effects have a population mean of zero. This means that the average of time intensities defines the average time to complete the test. Furthermore, the population-average speed trajectory is constant and shows no changes in speed, since the means of the random effects are zero. So the population-average trajectory with zero values for the random speed variables represents a constant population-average level of speed throughout the test. Test takers can work faster than this average level, which corresponds to a positive initial speed value. Furthermore, test takers can show an increasing or decreasing trend in their speed rate, which is represented by a positive or negative growth rate, respectively.

Figure 1 represents this (random linear) variable speed model. The bottom scale represents the order in which the items are solved. The upper axis represents the real-time scale. For each item, an RT is observed to measure speed, and for each RT observation, an error term represents the measurement error that is involved in measuring speed. The latent speed measurements are modeled using a random intercept, referred to as I , and a random growth rate, referred to as S . The average level of speed I is measured by all item observations, where the growth rate is measured by all item observations excluding the first item. The variances of the growth model variables, I and S , define the between-person variability in initial speed value and growth rates. A covariance term is specified between the growth model variables. Test takers who worked too slowly at the start of the test might improve their speed to finish the test in time. Test takers who started working very fast might later decrease their speed (possibly improving their accuracy level), since by working fast initially they would have sufficient time to finish the test. This corresponds to a negative correlation between the growth model parameters.

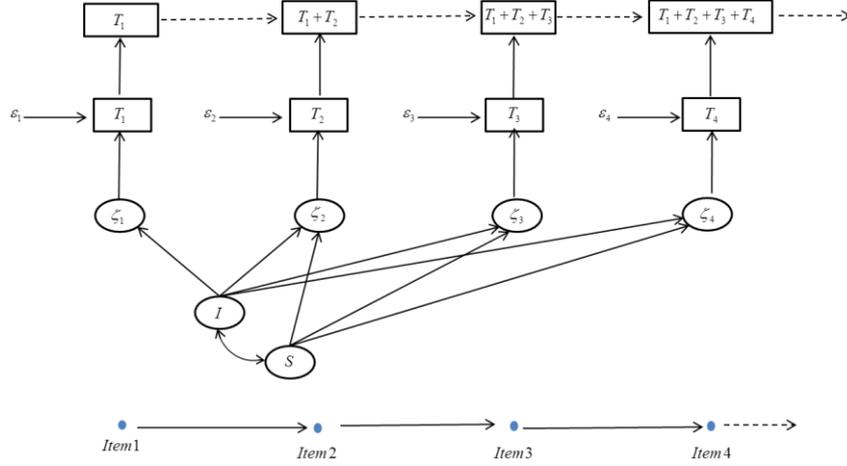


FIGURE 1: The lognormal random linear variable speed model

The Lognormal Random Quadratic Variable Speed Model

As stated above, to define a more complex speed process, the linear trend component is extended with a quadratic term. The linear trend can be used to model the speed processes of a test taker who starts to work faster and continues to work fast until the end of the test. However, a quadratic term can be used to decelerate or accelerate this linear trend. For example, a positive linear trend for speed can be decelerated by a negative quadratic term.

A random quadratic time component is included to define person-specific growth parameters. Then, each trajectory of working speed is modeled by an intercept, a linear trend, and a quadratic time component using individual parameters. The lognormal model with a random quadratic time variable is represented by

$$\ln T_{ik} = \lambda_k - \phi_k \left(\zeta_{i0} + \zeta_{i1} X_{ik} + \zeta_{i2} X_{ik}^2 \right) + \varepsilon_{ik}$$

$$\begin{pmatrix} \zeta_{i0} \\ \zeta_{i1} \\ \zeta_{i2} \end{pmatrix} \sim N \left(\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_{\zeta_0}^2 & \rho_{\zeta_{01}} & \rho_{\zeta_{02}} \\ \rho_{\zeta_{01}} & \sigma_{\zeta_1}^2 & \rho_{\zeta_{12}} \\ \rho_{\zeta_{02}} & \rho_{\zeta_{12}} & \sigma_{\zeta_2}^2 \end{pmatrix} \right). \quad (6)$$

In Figure 2, a graphical representation of the model is given. The bottom scale represents the order of responses to the items. The upper scale is the true time scale. The random intercept refers to the initial or average speed level, the linear trend is given by S , and a quadratic time component is given by Q . The random growth components are assumed to be correlated with common covariances across persons, according to the covariance matrix in Equation (6). In this model, the individual speed trajectories are modeled using three random effects, each with a mean of zero, such that the average time intensities define the average time to complete the test.

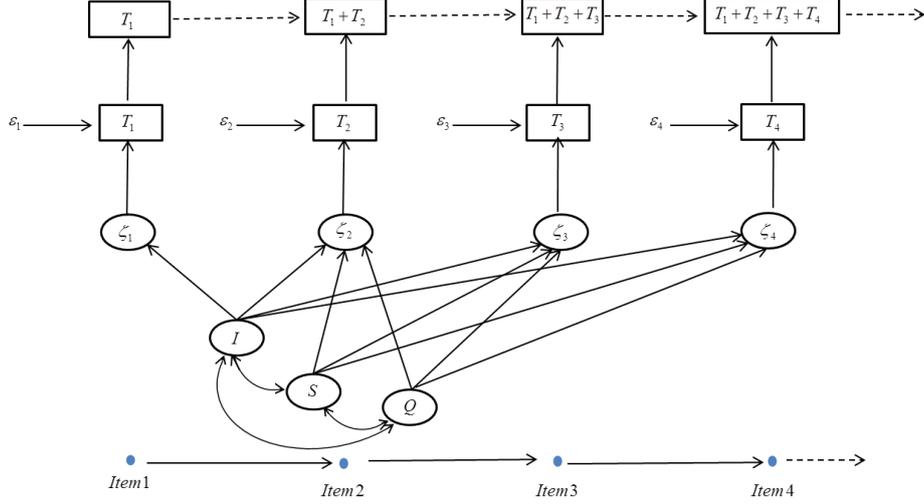


FIGURE 2: The lognormal random quadratic variable speed model

Joint Model for Responses and Response Times

Besides observing RTs, let Y_{ik} denote the response of person i ($i=1, \dots, N$) to item k ($k=1, \dots, K$). An IRT model is considered to model the item responses and to measure the ability of each test taker. When considering binary response data, we use a two-parameter normal ogive model with item discrimination parameter a_k and difficulty parameter b_k . Using the underlying latent response formulation, a latent response Z_{ik} is used, which is normally distributed with mean $a_k \theta_i - b_k$ and variance 1, and truncated from below (above) by zero when the response is correct (incorrect). So a correct response, $Y_{ik} = 1$, is the indicator that Z_{ik} is positive. The joint model for responses and RTs, allowing for variable speed is given by

$$\begin{aligned}
 Z_{ik} &= a_k \theta_i - b_k + \omega_{ik}, \quad \omega_{ik} \sim N(0, 1) \\
 \ln T_{ik} &= \lambda_k - \phi_k \left(\zeta_{i0} + \zeta_{i1} X_{ik} + \zeta_{i2} X_{ik}^2 \right) + \varepsilon_{ik}, \quad \varepsilon_{ik} \sim N\left(0, \sigma_{\varepsilon_k}^2\right) \quad (7) \\
 \begin{pmatrix} \theta_i \\ \zeta_{i0} \\ \zeta_{i1} \\ \zeta_{i2} \end{pmatrix} &\sim N \left(\begin{pmatrix} \mu_\theta \\ 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_\theta^2 & \rho_{\theta\zeta_0} & \rho_{\theta\zeta_1} & \rho_{\theta\zeta_2} \\ \rho_{\theta\zeta_0} & \sigma_{\zeta_0}^2 & \rho_{\zeta_01} & \rho_{\zeta_02} \\ \rho_{\theta\zeta_1} & \rho_{\zeta_01} & \sigma_{\zeta_1}^2 & \rho_{\zeta_12} \\ \rho_{\theta\zeta_2} & \rho_{\zeta_02} & \rho_{\zeta_12} & \sigma_{\zeta_2}^2 \end{pmatrix} \right)
 \end{aligned}$$

The prior distribution of the person parameters $(\theta_i, \zeta_i) = (\theta_i, \zeta_{0i}, \zeta_{1i}, \zeta_{2i})$ can be given as

$$\begin{pmatrix} \theta_i \\ \zeta_i \end{pmatrix} \sim N \left(\begin{pmatrix} \mu_\theta \\ \mathbf{0} \end{pmatrix}, \begin{pmatrix} \sigma_\theta^2 & \Sigma_{\theta\zeta} \\ \Sigma_{\zeta\theta} & \Sigma_\zeta \end{pmatrix} \right). \quad (8)$$

The relationship between speed and ability is defined by the covariance between ability and the speed components and is given by $\Sigma_{\theta\zeta}$. The ability parameter is influenced by the different speed components. This follows directly from the conditional distribution of ability given the speed variables. This distribution is given by

$$\theta_i | \zeta_i \sim N \left(\mu_\theta + \Sigma_{\theta\zeta} \Sigma_\zeta^{-1} (\zeta_i - \mu_\zeta), \sigma_\theta^2 - \Sigma_{\theta\zeta} \Sigma_\zeta^{-1} \Sigma_{\zeta\theta} \right). \quad (9)$$

Ability is influenced by the weighted average of the speed components, where the weights are defined by the covariance matrix $\Sigma_{\theta\zeta}$ times the inverse of the variance of speed components. When test takers do not vary their speed, only the first diagonal component of Σ_ζ will be larger than zero, showing the variability in constant speed values across test takers. The linear trend and quadratic change in speed will be around zero, which leads to a negligible influence of the remaining variable speed components on ability. When test takers vary their speed according to the quadratic variable speed model, the diagonal components of Σ_ζ will be larger than zero, and together with the covariance matrix $\Sigma_{\theta\zeta}$, will define the relation with ability. It follows that the constant speed model is generalized by allowing variable speed components to influence ability.

It will depend on the application whether changes in working speed will improve the accuracy of the responses. By measuring changes in working speed and modeling the relationship between speed and ability, it is possible to estimate the speed trajectories of test takers with different levels of ability. High-ability test takers may have different speed trajectories than low-ability test takers. The speed trajectories of test takers may also differ across tests. It will be possible to investigate the effects of time limits on test takers' speed changes, including those of proficient test takers. However, the benefits of estimating speed trajectories in relation to ability will depend on the application.

Identification

The latent scale of ability and speed needs to be identified. The mean and variance of the ability scale can be restricted to identify the scale. This can be accomplished by restricting the sum of item difficulties and product of discriminations, or by directly restricting the mean and variance of the ability parameter.

Next, for the variable speed model, the scale of the latent speed variable needs to be identified. This can also be accomplished by two restrictions. In the present description of the model, the mean of the speed parameter is set to zero to identify the mean of the speed scale. The average of the time-intensity parameters represents the population-average time needed to finish the test given an average working speed of zero. The variance of the speed scale is identified by fixing this

variance directly, or by restricting the product of time discriminations to one. For the joint model, the mean of each person parameter is restricted to zero, and the product of discriminations and time discriminations is restricted to one. These identification restrictions are also used by Klein Entink et al. (2009) and Fox (2010).

In the variable working speed model, an additional restriction is required, since the covariance between speed components is modeled by the time-intensity parameters and by the covariance matrix of the speed components, Σ_{ζ} . As mentioned above, the time-intensity parameters will influence the correlation between the RTs, which leads to an indeterminacy between the covariance parameters of speed and the time-intensity parameters. Therefore, as an additional constraint, the covariance matrix of the speed components is restricted to have zero nondiagonal terms, and the covariance between speed components is modeled by the time-intensity parameters. When the time-intensity parameters are all fixed to one, the covariance matrix of the speed components is a free matrix and no additional restriction is required. The residual errors are assumed to be independently distributed and do not influence the covariance modeling structure. When the ability and speed scale are identified, all higher-level model parameters will also be identified.

Parameter Estimation

The model parameters can be estimated using MCMC. The MCMC algorithms for the joint model with variable speed will follow the algorithms for the constant speed–ability joint models. In Fox (2010), the MCMC steps are fully explained for the so-called (constant speed) RTIRT model. The MCMC method was implemented in a modified version of the *cirt* R-program of Fox et al. (2007). The following sampling steps are required. At iteration $m = 1, \dots, M$,

1. For $k = 1, \dots, K$, sample item parameters from $p(\phi_k, \lambda_k, a_k, b_k | \mathbf{z}_k, \mathbf{t}_k, \boldsymbol{\theta}, \boldsymbol{\zeta}, \boldsymbol{\mu}_I, \boldsymbol{\Sigma}_I)$ using a multivariate normal prior with mean $\boldsymbol{\mu}_I$ and covariance matrix $\boldsymbol{\Sigma}_I$.
2. For $k = 1, \dots, K$, sample the residual variance in the lognormal model from $p(\sigma_{\varepsilon_k}^2 | \mathbf{t}_k, \boldsymbol{\zeta}, \lambda_k, \phi_k)$.
3. For $i = 1, \dots, N$, sample the ability parameter from $p(\theta_i | \boldsymbol{\zeta}_i, \boldsymbol{\mu}_\theta, \boldsymbol{\Sigma}_{\theta\zeta}, \mathbf{z}_i, \mathbf{t}_i)$.
4. Sample the hyperparameters $\boldsymbol{\mu}_I$ and $\boldsymbol{\Sigma}_I$ from $p(\boldsymbol{\mu}_I | \phi, \boldsymbol{\lambda}, \mathbf{a}, \mathbf{b}, \boldsymbol{\Sigma}_I)$ and $p(\boldsymbol{\Sigma}_I | \phi, \boldsymbol{\lambda}, \mathbf{a}, \mathbf{b})$.
5. For the constant speed model, sample the hyperparameter $\boldsymbol{\Sigma}_{\theta\zeta}$ from $p(\boldsymbol{\Sigma}_{\theta\zeta} | \boldsymbol{\theta}, \boldsymbol{\zeta}, \phi, \boldsymbol{\lambda}, \mathbf{a}, \mathbf{b})$.

For the variable speed model, several additional sampling steps are required. With an identification restriction on the covariance matrix of the person parameters, the sampling of the speed components $\boldsymbol{\zeta}$ and the free parameters of the covariance matrix requires a stepwise approach. The speed components are a priori independently and normally distributed. Each diagonal component of the covariance

matrix Σ_ζ is inverse-gamma distributed with an inverse gamma prior with parameters g_1 and g_2 . The conditional distribution of the variance parameter of speed component ζ_j ($j = 0, 1, 2$), is given by

$$\sigma_{\zeta_j}^2 \mid \boldsymbol{\mu}_\zeta, \boldsymbol{\zeta}_j, \mathbf{T} \sim IG\left(\frac{N}{2} + g_2, \sum_i (\zeta_{ji} - \mu_{\zeta_j})^2 / 2 + g_1\right), \quad (10)$$

and the three variance parameters define the diagonal of the covariance matrix Σ_ζ . The speed components are conditionally normally distributed, and it follows that

$$\zeta_i \mid \theta_i \sim N\left(\boldsymbol{\mu}_\zeta + \Sigma_{\zeta\theta} \sigma_\theta^{-2} (\theta_i - \mu_\theta), \Sigma_\zeta - \Sigma_{\zeta\theta} \sigma_\theta^{-2} \Sigma_{\zeta\theta}\right). \quad (11)$$

In this conditional distribution, the covariance $\Sigma_{\zeta\theta}$ in the mean term is considered to be a regression parameter. The conditional distribution of this parameter is normal with variance

$$\boldsymbol{\Omega} = E\left(\Sigma_{\zeta\theta}\right) = \left(\sigma_\theta^2 - \Sigma_{\theta\zeta} \Sigma_\zeta^{-1} \Sigma_{\zeta\theta}\right)^{-1} \left(\left(\Sigma_\zeta^{-1} (\boldsymbol{\zeta}_i - \boldsymbol{\mu}_\zeta)\right)^t \Sigma_\zeta^{-1} (\boldsymbol{\zeta}_i - \boldsymbol{\mu}_\zeta) \right) + \Sigma_\theta^{-1} \quad (12)$$

and mean

$$\Sigma_{\zeta\theta} = \boldsymbol{\Omega}^{-1} \left(\Sigma_\zeta^{-1} (\boldsymbol{\zeta}_i - \boldsymbol{\mu}_\zeta)^t (\theta_i - \mu_\theta) \left(\sigma_\theta^2 - \Sigma_{\theta\zeta} \Sigma_\zeta^{-1} \Sigma_{\zeta\theta}\right)^{-1} \right). \quad (13)$$

From the conditional distribution of θ_i given ζ_i , the distribution of the variance parameter can be derived. This variance parameter, $\sigma = \sigma_\theta^2 - \Sigma_{\theta\zeta} \Sigma_\zeta^{-1} \Sigma_{\zeta\theta}$, is inverse-gamma distributed with scale parameter

$$SS = \sum_i \left((\theta_i - \mu_\theta) - \Sigma_{\theta\zeta} \Sigma_\zeta^{-1} (\zeta_i - \boldsymbol{\mu}_\zeta) \right)^2 / 2 + g_1 \quad (14)$$

and shape parameter $N/2 + g_2$. Subsequently, from the sampled variance parameter, σ , a sampled value of the variance parameter σ_θ^2 can be obtained using the sampled value for Σ_ζ .

Without the identification restriction on the covariance matrix, Σ_ζ , the values of the complete covariance matrix of the person parameters,

$$\Sigma_P = \begin{pmatrix} \sigma_\theta^2 & \Sigma_{\theta\zeta} \\ \Sigma_{\zeta\theta} & \Sigma_\zeta \end{pmatrix}, \quad (15)$$

are sampled from an inverse-Wishart distribution with scale matrix

$$\sum_i \begin{pmatrix} \theta_i - \mu_\theta \\ \zeta_i - \mu_\zeta \end{pmatrix} \begin{pmatrix} \theta_i - \mu_\theta \\ \zeta_i - \mu_\zeta \end{pmatrix}^t + g_1 \quad (16)$$

and degrees of freedom $N + Q$, where Q is the number of random effects. The mean of the speed components, μ_ζ , is fixed to zero.

Simulation Study

In this simulation study, attention was focused on a variable speed process in the joint modeling of responses and RTs. The RTs were modeled according to a lognormal random quadratic variable speed model. This RT model included a random trend and a random quadratic time variable, which is represented by

$$\ln T_{ik} = \lambda_k - (\zeta_{i0} + \zeta_{i1} X_{ik} + \zeta_{i2} X_{ik}^2) + \varepsilon_{ik}, \quad \varepsilon_{ik} \sim N(0, \sigma_{\varepsilon_k}^2). \quad (17)$$

The random speed components had a mean of zero to identify the time intensities. The responses were modeled according to a two-parameter normal ogive model. The ability and speed parameters were assumed to be multivariate normally distributed, according to Equation (8).

The joint model for responses and RTs was used to generate the data, and a modified version of the *cirt* program of Fox et al. (2007) was used to estimate all model parameters. The item parameters were simulated from a multivariate normal distribution with the covariance matrix given in Table 1. The mean of the time discriminations and discrimination parameters was set to one, and the mean of the difficulty and time-intensity parameters was set to zero. The measurement error variance was set to .50 for each item. Furthermore, the model was identified by restricting the covariance between random speed components to zero, and restricting the product of time discrimination and discrimination parameters to one.

To evaluate the performance of the developed MCMC algorithm, a total of 50 datasets were simulated according to the joint model for 1,000 test takers responding to 20 and 40 items. A burn-in period of 1,000 iterations was used, and a total of 5,000 iterations were made to estimate all model parameters.

In Table 1, the true and final parameter estimates are given across the 50 simulated datasets. Under the header “LNIRTQ,” the parameter estimates are represented for replicated datasets of 20 and 40 items. The true value of the covariance matrix of the person parameters shows that the speed trajectories differ a lot across test takers. The random variation over test takers in the trend and quadratic components was around .50, given a variance of one across the average levels of speed. Furthermore, there was a positive covariance simulated between the random person components.

For the 20-item test-length condition, the estimated values of the person covariance matrix are close to the true values, despite the high level of variation in simulated speed trajectories. The variability in speed trajectories, which differed in their trend and quadratic components, was accurately estimated. Estimates of the

item covariance parameters also showed a good recovery of the true parameter values.

TABLE 1
Simulated and estimated parameter values (over 50 data replications) of the joint model with a lognormal random quadratic variable speed component for 1,000 test takers, and with 20 and 40 items

Variance Components		True Mean	LNIRTQ			
			K = 20		K = 40	
			Mean	SD	Mean	SD
Person Covariance Matrix						
Ability	σ_{θ}^2	1	1.031	0.059	1.009	0.052
	$\rho_{\theta\zeta_0}$	0.7	0.704	0.046	0.699	0.044
Speed	$\sigma_{\zeta_0}^2$	1	0.998	0.054	1.001	0.052
	$\rho_{\theta\zeta_1}$	0.2	0.194	0.052	0.187	0.039
	$\sigma_{\zeta_1}^2$	0.5	0.508	0.083	0.486	0.060
	$\rho_{\theta\zeta_2}$	0.1	0.111	0.048	0.107	0.037
	$\sigma_{\zeta_2}^2$	0.5	0.486	0.092	0.526	0.071
Item Covariance Matrix						
Discrimination	\sum_{11}	0.05	0.065	0.033	0.057	0.016
	\sum_{12}	0	-0.005	0.067	-0.002	0.043
	\sum_{13}	0	0	0.019	0.000	0.010
	\sum_{14}	0	0.001	0.073	0.002	0.042
Difficulty	\sum_{22}	1	1.054	0.365	1.040	0.244
	\sum_{23}	0	0	0.063	-0.005	0.040
	\sum_{24}	0	0.034	0.258	-0.017	0.163
Time Discrimination	\sum_{33}	0.05	0.068	0.027	0.059	0.014
	\sum_{34}	0	0.018	0.069	0.007	0.040
Time Intensity	\sum_{44}	1	1.122	0.388	0.981	0.230

Subsequently, the test length was increased from 20 to 40 items. So, twice as many RT observations were generated for each speed trajectory. In Table 1 under the label “K = 40,” the final estimates of the person and item covariance parameters for the 40-item test are given. It can be concluded that the estimated parameter values are close to true values. Furthermore, it follows directly that the estimated standard deviations are smaller than those based on the 20-item test data. Although not shown, the item parameters of the 20- and 40-item test and the measurement error variance were also accurately recovered across the 50 simulated datasets.

Modeling Variable Speed in the Amsterdam Chess Test Data

The ACT data of Van der Maas and Wagenmakers (2005) were used to identify the variable speed trajectories of 259 test takers, who responded to 40 chess tasks. The chess items were divided over three sections: tactical skill (20 items), positional skill (10 items), and end-game skill (10 items). Each item concerned a chess board situation, and the problem-solving task was to select the best possible move. The dichotomous response observations 1 (correct) and 0 (incorrect) as well as the RTs were stored. Fox (2010, p. 253), using the joint model of Klein Entink et al. (2009), analyzed the data using the RTIRT model to identify items not fitting the data.

The purpose of the present study was to investigate whether test takers worked with variable speed, and what type of speed trajectories could be identified. Furthermore, the complex between-person relationship between ability and random speed components was investigated. In a different approach, Molenaar et al. (2015a) considered a function of ability on speed in their generalized linear model for responses and RTs. They found a common curvilinear effect of ability on speed for the end-game items, where higher-ability test takers tended to use more time in contrast to lower-ability test takers, who started to answer faster at the end of the test. In their approach, higher-order interaction terms between ability and speed were used to obtain more insight into the relation between speed and ability, but the higher-order ability components were assumed to be fixed deterministic components (with a common effect across test takers), and it was assumed that speed did not have an influence on ability.

The RTIRT lognormal (constant) speed model and the lognormal (random quadratic) variable speed model (Equation (7)) were used to analyze the data. The MCMC algorithm was run for 10,000 iterations to estimate all model parameters, where a burn-in period of 1,000 iterations was used. In Figure 3, for the variable speed model, trace plots of four (variance and covariance) person parameters are given to show the fast convergence and the stable behavior of the MCMC chains. The MCMC algorithm converged rapidly without specifying informative starting values. The other trace plots showed similar behavior. The R-coda package was used to investigate the chains, and the commonly used convergence diagnostics (e.g., Geweke, Heidelberger, and Welch) did not show any issues.

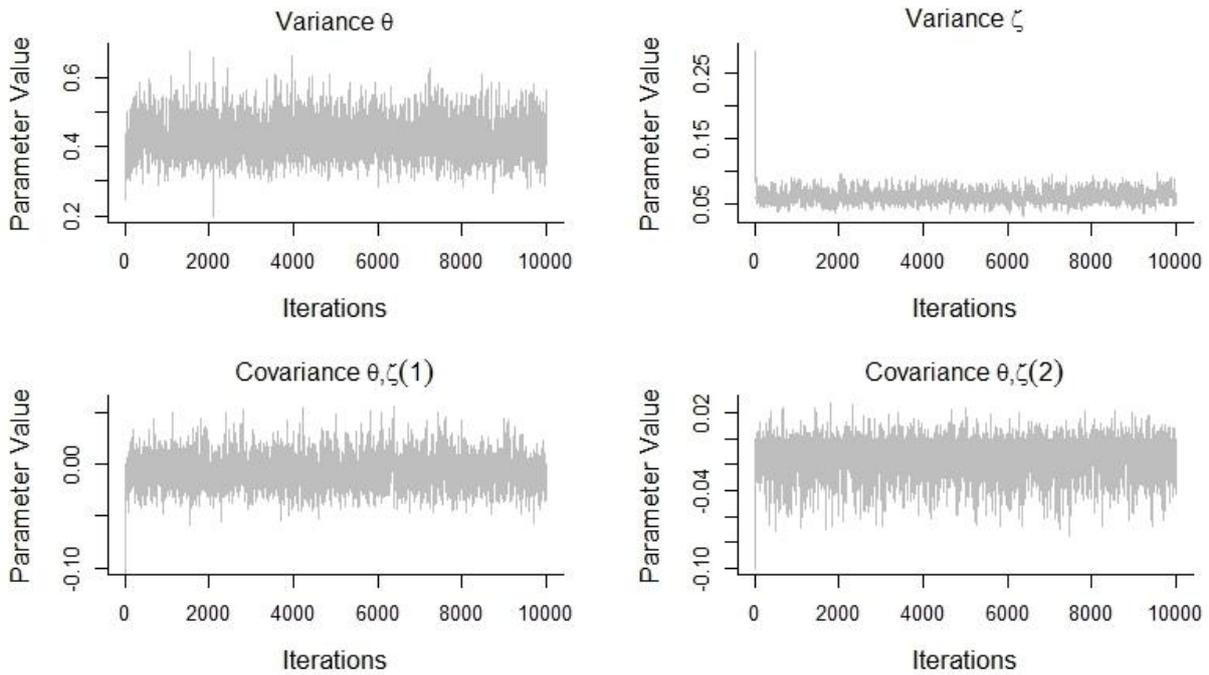


FIGURE 3: Trace plots of the ability and average speed population variance parameters, and the covariance between ability and the slope and quadratic speed components.

The RTIRT with a constant speed factor was fitted to the data. In Figure 4, the item parameter estimates of the 40-item test are given. It can be seen that there is sufficient variation in difficulty and item intensity to measure the ability and speed factors accurately at the different levels of the scale. The time discriminations are higher for the first 10 items (they define the tactical skill cluster), which means that responses to those items show more variation between slow- and fast-working test takers. The time discriminations for the end-game items were not as high, indicating less power to discriminate between test takers' working speed. Most items discriminated sufficiently between test takers' ability levels; only around five items had a low discriminating value of < 0.5 .

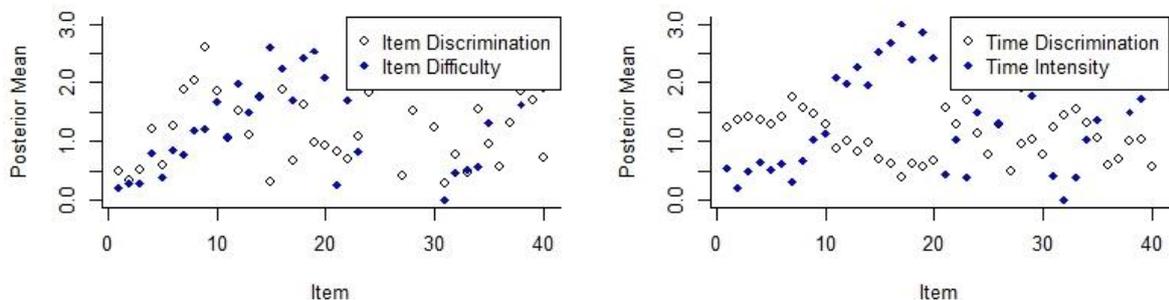


FIGURE 4: Item parameter estimates of the 40 chess items. The left-hand plot shows the discrimination (open diamond symbol) and difficulty (closed diamond symbol) estimates; the right-hand plot shows the time discrimination (open diamond symbol) and time intensity (closed diamond symbol) estimates.

The covariance estimates of the test taker's random factors and the item parameters of the joint model with a constant speed factor are given in Table 2. There is substantial between-item variation in difficulty and intensity but less variation in discrimination and time discrimination. The mean RT residual variation was around .25, and ranged from .15 to .50. The correlation between discrimination and difficulty was around .35, and between time discrimination and intensity around $-.92$. This strong negative correlation of $-.92$ showed that, for high time-intensive items, the speed factor did not explain much variation in RTs, whereas the speed factor did explain it for low time-intensive items. According to the model, for a time-intensive item, an increase in working speed does not have much effect on RT due to the low time-discrimination parameter. The influence on the RT due to a change in speed is much higher for low time-intensive items, which have high time discriminations.

The strong correlation of around .80 between item difficulty and item intensity was also apparent. The difficult items were clearly taking much more time to solve than the easy items.

TABLE 2
Amsterdam Chess Test: Covariance components and correlation estimates

Variance Components		Constant Speed			Variable Speed		
		Mean	SD	Cor.	Mean	SD	Cor.
Person Covariance Matrix							
Ability	σ_{θ}^2	0.317	0.040		0.425	0.139	
	$\rho_{\theta\zeta_0}$	0.105	0.013	0.640	0.115	0.016	0.720
Speed	$\sigma_{\zeta_0}^2$	0.085	0.008		0.06	0.009	
	$\rho_{\theta\zeta_1}$				-0.004	0.014	-0.018
	$\sigma_{\zeta_1}^2$				0.113	0.024	
	$\rho_{\theta\zeta_2}$				-0.014	0.015	-0.086
	$\sigma_{\zeta_2}^2$				0.062	0.015	
Item Covariance Matrix							
Discrimination	Σ_{11}	0.479	0.175		0.533	0.188	
	Σ_{12}	0.252	0.138	0.352	0.289	0.150	0.371
	Σ_{13}	0.030	0.052	0.108	0.019	0.064	0.053
	Σ_{14}	0.067	0.151	0.079	0.147	0.07	0.408
Difficulty	Σ_{22}	1.073	0.259		1.141	0.271	
	Σ_{23}	-0.268	0.085	-0.645	-0.317	0.105	-0.603
	Σ_{24}	1.008	0.272	0.794	0.414	0.111	0.786
Time Discrimination	Σ_{33}	0.161	0.047		0.242	0.076	
	Σ_{34}	-0.450	0.123	-0.915	-0.130	0.047	-0.536
Time Intensity	Σ_{44}	1.503	0.375		0.243	0.058	

For the person parameters, there was not much variation in speed levels (around .085) or ability levels (around .32). Under the constant working speed assumption, the correlation between ability and speed was around .65, which showed that high-ability test takers were also completing the items faster. They were able to identify the solution to the chess problem faster than the low-ability test takers.

This covariance structure holds under the assumption that test takers were working with a constant speed. To investigate variable speed trajectories of test takers, the joint model with the random quadratic variable speed model was also fitted. In Table 2, the covariance estimates are given under the header “Variable Speed.” For the covariance between item parameters, it can be seen that the strong correlation between time discrimination and time intensity diminished to $-.54$. Apparently, the additional speed components explained the greater variation in RTs, reducing the strong correlation between time discrimination and intensity. The correlation between the item discrimination and time intensity increased to $.41$. The chess items that were highly discriminating in ability were also the time-intensive items. This relates to the positive correlation between ability and speed. It is likely that the test takers showed different speed behavior in responding to well-discriminating items depending on their ability.

The correlation between the average speed level and ability was around $.72$, which was slightly higher than it was under the constant speed model. The corresponding 95% highest posterior density (HPD) interval was $[0.637, 0.785]$. A slightly negative correlation of $-.02$ was estimated between ability and the random slope speed component (95% HPD interval equaled $[-0.157, 0.099]$). This means that high-ability test takers were more likely to decrease their speed in a linear way. The correlation between ability and the quadratic speed component was around $-.09$ (95% HPD interval equaled $[-0.266, 0.035]$), which means that high-ability test takers were more likely to show an acceleration in the negative trend in speed. However, both estimated correlation parameters were not significantly different from zero, since zero was included in the 95% HPD intervals.

From the trace plots of the covariance parameters (see Figure 3), it can also be seen that the drawn covariance values are not significantly different from zero. Since the correlations with ability were not significantly different from zero, the characteristics of the speed trajectories could not be explained by differences in ability.

In Figure 5, the ability estimates are plotted against the random components of speed. It can be seen that there is a strong positive relation between ability and average speed, where the relation between ability and the slope and quadratic speed components is not significantly different from zero. The estimated average speed component was conditionally estimated on the two other random speed components, which accounted for nonconstant speed behavior. The positive correlation between the linear and quadratic speed component showed that a more negative (positive) trend in speed was accelerated, leading to an even slower (higher) working speed.

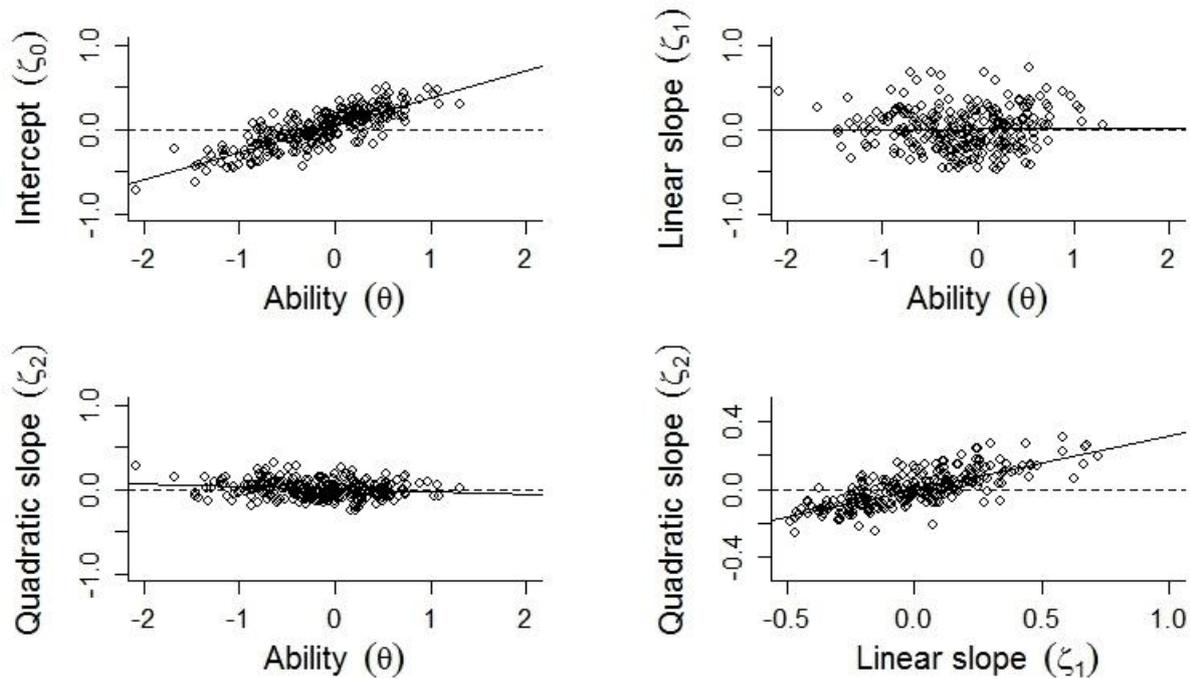


FIGURE 5. Random person parameter estimates; the average speed (ζ_0), slope speed (ζ_1), and the quadratic speed (ζ_2) component plotted against ability (θ). The slope of speed is plotted against the quadratic speed component.

In Figure 6, from the total sample of $N = 259$ test takers, the fitted item-specific working-speed measurements of 20 high- and low-ability test takers were plotted. The test takers started working at different speeds: The high-ability group started to work faster than the low-ability group. Some high-ability test takers increased their working speed toward the end of the test, but others decreased their level of working speed around halfway through the test. The low-ability test takers showed the opposite behavior. Most of the low-ability test takers increased their working speed halfway through the test; only a few showed a constant decrease in working speed. It is possible that high-ability test takers (the 10% highest scoring test takers) were more focused and eager to ensure that all items were correct, whereas low-ability test takers (the 10% lowest-scoring test takers) might have been less motivated halfway through the test, proceeding more quickly through the latter half of the test. Molenaar et al. (2015a) reported on this pattern, based on a higher-order interaction effect between ability and speed. With the quadratic variable speed model, each test taker's trajectory of working speed was estimated, showing the patterns while controlling for the correlation with ability. This made it possible to estimate the variable working speed behavior of each test taker.

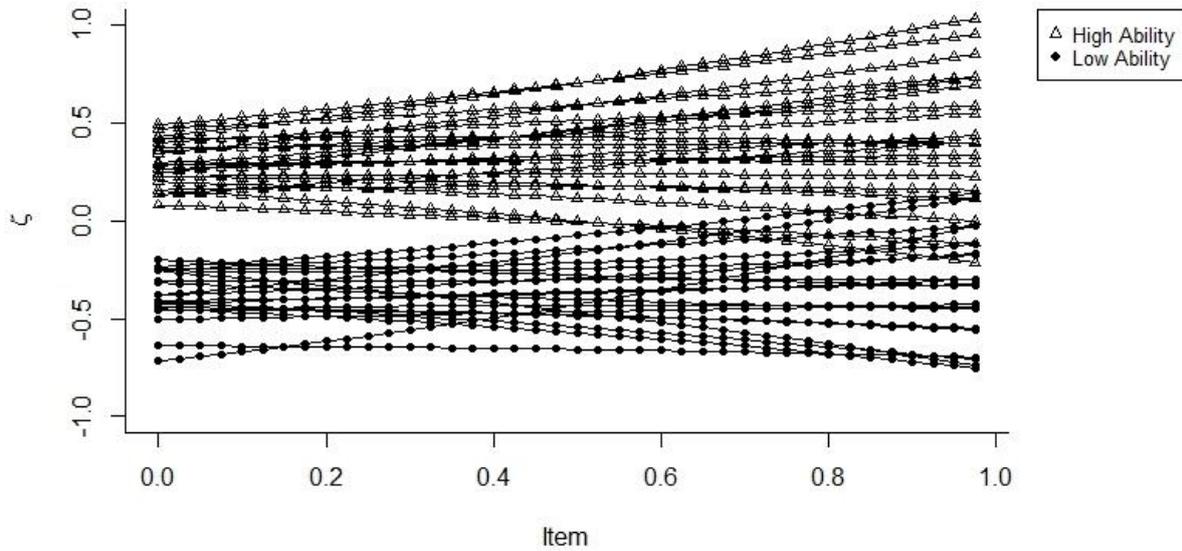


FIGURE 6. *Fitted latent speed trajectories over items of high- and low-ability test takers*

Discussion

Computer-based testing makes it possible to collect RT information as well as response information, by simply recording the total time spent on each item and the response to each item, respectively. RTs can be used to make more accurate inferences about test takers' ability (e.g., van der Linden et al., 2010) and item characteristics. RTs can also reveal new information about test characteristics, test takers' response behavior, and test takers' ability that would not be identified when using response information alone.

The latent growth model for working speed can be used to measure variable working speed according to a time scale defined by the order in which the items were answered. In the present model, the random slope of speed and random quadratic speed components were added to model deviations from a constant speed model. The extension to higher-order random effect components can be made. However, this will require a sufficient number of item observations to estimate all model parameters. The higher-order terms can also be included for groups of test takers using discrete latent random effects. When the trajectory of speed includes higher-order components, the relation between working speed and ability can be defined as the weighted correlation between ability and all the speed components. In that case, the additional random speed components are used to control for nonconstant speed to improve the estimation of the relation between speed and ability.

Note that the order in which the items are answered does not influence the ability estimate. The estimation of the latent speed trajectories depends on both the RT information and the order in which the items are answered. In the estimation of each speed trajectory, the information both for each item-specific RT observation and for the relationships between RTs is used. So, if the order of observed RTs were changed, a different trajectory would be estimated, since the relationships between RTs would be different.

This model, which is a generalization of the constant speed model proposed by van der Linden (2007), can be used to measure a more complex relationship

between ability and speed. From a model-building perspective, it is recommended that one first evaluate the fit of the hierarchical model for responses and RTs before fitting a differential speed component. More research is needed to develop information criteria (e.g., Bayesian information criterion [BIC], deviance information criterion [DIC]) that are able to identify the best joint model for responses and RTs among a set of competing models. For example, a straightforward implementation of the DIC is not going to produce reliable results, since the estimation of the number of effective parameters, which is required to compute the DIC, is very complex when the model contains many random effects, multiple outcomes of different types (categorical and continuous), and multiple link functions (linear and nonlinear). Future research will focus on the procedure of Klein Entink et al. (2009) and Fox (2010, p. 241–242), who considered a DIC based on the integrated likelihood (e.g., see Berger, Brunero, & Wolpert, 1999), where ability and speed were integrated out. This simplifies the computation of the penalty term, since it is no longer based on the random person parameters, and leads to a more accurate estimate of the number of model parameters.

The so-called cross-relation between speed and accuracy was also modeled by Molenaar et al. (2015a), who considered different functions of higher-order ability components on speed. They introduced two person factors: ability and speed. In the proposed model, several random person variables were introduced to better describe this relationship by assuming a variable speed model. The MCMC algorithm developed in this report (see section titled Parameter Estimation) can handle numerous random effects, since it is a simulation-based estimation procedure.

The situation where test takers are limited in their responding due to time constraints is referred to as test speededness. However, when speed is not of interest it should also not interfere with the measurement of ability. Speededness is considered to be a threat to the validity of the test scores; it inadvertently interferes with the performance level of the test takers. Research has focused on detecting test speededness as a threat to test validity. Goegebeur, De Boeck, Wollack, and Cohen (2008) and Chang, Tsai, and Hsu (2014) considered a mixture modeling approach and defined a speeded class to identify test takers whose performance is affected by the time limit. The general idea is that the time limit influences the probability of an item being answered, without considering RT information. Shao and Cheng (2015) considered a change-point model to identify speeded test takers, and considered removing the speeded responses to improve ability estimation. Given the RT and response information, the joint model for ability and speed, using a latent growth model for speed, can provide insight into test speededness. Test takers' fitted speed trajectories can be used to identify (strategic) speed behavior while accounting for differences in ability. An increase in speed at the end of the test would indicate the influence of a time limit on the test performance, whereas a decrease in speed would show the opposite.

Test speededness has a negative influence on test-taker performance. Test fraud, in contrast, is intended to have an opposite effect on performance. Test takers may try to positively distort their responses to improve their achievement scores, which overestimates the test takers' true achievement level. Nowadays, there is an increased interest in test fraud detection (e.g., Kingston & Clark, 2014), where attention is focused on test-taking effort. Test takers may show solution behavior or rapid guessing behavior, where the guessed responses provide no information about the true achievement level of the test takers. Schnipke and Scrams (1997) considered the use of RTs to identify rapid guessing behavior for a speeded test.

Wise and Kong (2005) also considered RTs to measure RT effort, which addresses the proportion of true solution behavior in contrast to rapid guessing behavior. Wang and Xu (2015) developed a hierarchical mixture model to determine whether test takers' response strategies could be identified as rapid guessing behavior or solution behavior. Although these approaches consider two different strategies that a test taker might use, they do not consider the actual speed trajectory that underlies the observed RTs. Under the joint model, the speed trajectory in relation to ability will give a more accurate description of the test engagement of the test taker. Therefore, extreme speed trajectories, indicating that responses are given without evaluating the meaning of the question, can indicate rapid guessing behavior, where the speed components are related to ability. Furthermore, a test taker's quadratic time-component effect will show whether a test taker accelerates or decelerates rapidly during the test. This information can be used to address inconsistencies in response patterns. More generally, potential threats to the validity of the test scores (e.g., guessing, cheating) can be evaluated by exploring the speed trajectories of test takers in relation to ability. Statistical tests similar to the tests for aberrant speed behavior reported by Mariani et al. (2014) could be developed to identify extreme changes in speed behavior.

Several model extensions could be considered to make this model suitable for multiple group or multiple latent group settings, or for polytomous or nominal response data. The multiple group modeling approach of Azevedo, Andrade, and Fox (2012) might be used to extend the joint modeling of responses and RTs to a multiple group setting. Another interesting extension would be to consider a latent growth model for ability. This would lead to a multivariate latent growth modeling framework for ability and speed, to model changes in the factor variables (speed and ability) over time. Then, changes in speed and ability could be jointly modeled to investigate, for example, changes in the accuracy–speed trade-off over time.

References

- Assink, N., van der Lubbe, R. H., & Fox, J. P. (2015). Does time pressure induce tunnel vision? An examination with the Eriksen Flanker Task by applying the hierarchical drift diffusion model. In Y. Wang, P. Borne, & I. Rudas (Eds.), *New developments in computational intelligence and computer science* (pp. 30–40). Proceedings of the International Conference on Applied Physics, Simulation and Computers, Vienna, Austria.
- Azevedo, C. L., Andrade, D. F., & Fox, J.-P. (2012). A Bayesian generalized multiple group IRT model with model-fit assessment tools. *Computational Statistics and Data Analysis*, *56*(12), 4399–4412.
- Berger, J. O., Brunero, L., & Wolpert, R. L. (1999). Integrated likelihood methods for eliminating nuisance parameters. *Statistical Science*, *14*, 1–28.
- Chang, Y. W., Tsai, R. C., & Hsu, N. J. (2014). A speeded item response model: Leave the harder till later. *Psychometrika*, *79*(2), 255–274.
- De Boeck, P., & Partchev, I. (2012). IRTrees: Tree-based item response models of the GLMM family. *Journal of Statistical Software*, *48*(1), 1–28.

- Ferrando, P. J., & Lorenzo-Seva, U. (2007a). An item response theory model for incorporating response time data in binary personality items. *Applied Psychological Measurement, 31*(6), 525–543.
- Ferrando, P. J., & Lorenzo-Seva, U. (2007b). A measurement model for Likert responses that incorporates response time. *Multivariate Behavioral Research, 42*(4), 675–706.
- Fox, J.-P. (2010). *Bayesian item response modeling: Theory and applications*. New York: Springer.
- Fox, J.-P., Klein Entink, R., & van der Linden, W.J. (2007). Modeling of responses and response times with the package CIRT. *Journal of Statistical Software, 20*(7), 1–14.
- Goegebeur, Y., De Boeck, P., Wollack, J. A., & Cohen, A. S. (2008). A speeded item response model with gradual process change. *Psychometrika, 73*(1), 65–87.
- Kingston, N., & Clark, A. (2014). *Test fraud: Statistical detection and methodology*. London: Taylor & Francis.
- Klein Entink, R. H., Fox, J.-P., & van der Linden, W.J. (2009). A multivariate multilevel approach to the modeling of accuracy and speed of test takers. *Psychometrika, 74*(1), 21–48.
- Klein Entink, R. H., Kuhn, J. -T., Hornke, L. F., & Fox, J. -P. (2009). Evaluating cognitive theory: A joint modeling approach using responses and response times. *Psychological Methods, 14*, 54–75.
- 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*(6), 426–451.
- Molenaar, D., Tuerlinckx, F., & van der Maas, H. L. (2015a). A generalized linear factor model approach to the hierarchical framework for responses and response times. *British Journal of Mathematical and Statistical Psychology, 68*(2), 197–219.
- Molenaar, D., Tuerlinckx, F., & van der Maas, H. L. (2015b). A bivariate generalized linear item response theory modeling framework to the analysis of responses and response times. *Multivariate Behavioral Research, 50*(1), 56–74.
- Partchev, I., & De Boeck, P. (2012). Can fast and slow intelligence be differentiated? *Intelligence, 40*(1), 23–32.
- Ranger, J., & Kuhn, J. T. (2012). A flexible latent trait model for response times in tests. *Psychometrika, 77*(1), 31–47.
- Roberts, R. D., & Stankov, L. (1999). Individual differences in speed of mental processing and human cognitive abilities: Toward a taxonomic model. *Learning and Individual Differences, 11*, 1–120.

- Schnipke, D. L., & Scrams, D. J. (1997). Modeling item response times with a two-state mixture model: A new method of measuring speededness. *Journal of Education Measurement*, 34(3), 213–232.
- Shao, C., Li, J., & Cheng, Y. (2015). Detection of test speededness using change-point analysis. *Psychometrika*, 1–24. doi: 10.1007/s11336-015-9476-7
- Vandekerckhove, J., Tuerlinckx, F., & Lee, M. D. (2011). Hierarchical diffusion models for two-choice response times. *Psychological Methods*, 16(1), 44.
- van der Linden, W. J. (2006). A lognormal model for response times on test items. *Journal of Educational and Behavioral Statistics*, 31(2), 181–204.
- van der Linden, W. J. (2007). A hierarchical framework for modeling speed and accuracy on test items. *Psychometrika*, 72(3), 287–308.
- van der Linden, W. J. (2015). Lognormal response-time model. In W.J. van der Linden (Ed.), *Handbook of item response theory: Vol 1. Models*. Boca Raton, FL: Chapman & Hall/CRC.
- van der Linden, W. J., & Fox, J. -P. (2015). Joint hierarchical modeling of responses and response times. In W.J. van der Linden (Ed.), *Handbook of item response theory: Vol 1. Models*. Boca Raton: FL: Chapman & Hall/CRC.
- van der Linden, W. J., & Glas, C. A. W. (2010). Statistical tests of conditional independence between responses and/or response times on test items. *Psychometrika*, 75(1), 120–139.
- van der Linden, W. J., Klein Entink, R. H., & Fox, J. -P. (2010). IRT parameter estimation with response times as collateral information. *Applied Psychological Measurement*, 34, 327–347.
- van der Maas, H. L., & Wagenmakers, E. J. (2005). A psychometric analysis of chess expertise. *The American Journal of Psychology*, 118, 29–60.
- Wang, C., & Xu, G. (2015). A mixture hierarchical model for response times and response accuracy. *British Journal of Mathematical and Statistical Psychology*, 68, 456–477.
- Wise, S. L., & Kong, X. (2005). Response time effort: A new measure of examinee motivation in computer-based tests. *Applied Measurement in Education*, 18, 163–183.