Dynamic Paired Comparison Models with Stochastic Variances

Mark E. Glickman*
Boston University

Abstract

In paired comparison experiments, the worth or merit of a unit is measured through comparisons against other units. When paired comparison outcomes are collected over time and the merits of the units may be changing, it is often convenient to assume the data follow a non-linear state-space model. Typical paired comparison state-space models that assume a fixed (unknown) autoregressive variance do not account for the possibility of sudden changes in the merits. This is a particular concern in modeling, for example, cognitive ability in human development; cognitive ability not only changes over time, but can change abruptly. We explore a particular extension of conventional state-space models for paired comparison data that allows the state variance to vary stochastically. Models of this type have recently been developed and applied to modeling financial data, but can be seen to have applicability in modeling paired comparison data. The model is applied to a data set involving chess game outcomes. A filtering algorithm is also derived that can be used in place of likelihood-based computations when the number of objects being compared is large.

Keywords: Approximate Bayesian estimation, Bradley-Terry model, Chess, State-space models, Stochastic volatility, Thurstone-Mosteller model.

*Address for correspondence: Department of Mathematics and Statistics, Boston University, 111 Cummington Street, Boston, MA 02215, USA. E-mail: mg@math.bu.edu

1 Introduction

Paired comparison data arises when objects are compared to elicit a preference or a degree of preference. The literature on paired comparison modeling is vast, spanning fields such as statistics, marketing, psychology, and decision sciences. Background on fundamental issues in paired comparison modeling along with examples can be found in David (1988) and Bradley (1984). A common situation is to observe paired comparison data over time where the underlying value or worth of the objects are changing. This might occur, for example, in comparing preferences towards the value of marketed products or services, or in the outcomes of games played between competitors whose abilities may be changing over time. Recent work (Glickman, 1993; Fahrmeir and Tutz, 1994; Glickman, 1999) has adopted a state-space approach to modeling such data where the underlying merit parameters change as a Gaussian stochastic process. While this approach can often appropriately describe the change in merits, it can be too restrictive if the merits can undergo sudden shifts, or if interventions occur that change the merit of an object quickly. For example, in a marketing context, if a product is reported to be defective or dangerous, it will likely lose merit quickly. When comparing human cognitive skill through games such as chess, younger players may undergo quick increases in ability (Simonton, 1997) which the simpler state-space models cannot capture.

This paper describes an extension of the usual state-space models for paired comparison data that allows for sudden movement in the underlying merit parameters. The extension is closely related to the stochastic volatility model (Jacquier et al., 1994; Capobianco, 1996; Uhlig, 1997) developed in the context of modeling financial time series data. Our model involves letting not only the merit of objects change stochastically, but letting the variance of the state process change stochastically as well. In this extension, sudden shifts in merit are reflected through the variance of the change in merits becoming large. In Section 2, the model allowing the variance of the state process to change stochastically is developed.

This is followed in Section 3 by an application of the model to a data set on chess game outcomes. An algorithm is then presented in Section 4, extending an algorithm developed in Glickman (1999), that approximates likelihood-based computations using nearly closed form calculations. A situation in which one might want to use such an algorithm is when many objects are being compared, so that exact likelihood-based methods become computationally intractable. We provide a summary of the model, and consider directions of extensions in Section 5.

2 A stochastic variance paired comparison model

Suppose $K_{ij}^{(t)}$ comparisons are to take place between objects i and j at time t. Time is assumed to be discretized into periods of equal duration, so that t, which indexes a time interval, takes on integer values. We treat data observed within a time period as occurring at the start of the period. Let $Y_{ijk}^{(t)}$ be 1 if i is preferred to j in the k-th comparison between the two objects at time t, and 0 if j is preferred. Let $\gamma_i^{(t)}$ and $\gamma_j^{(t)}$ be the merits of the objects at time t. We assume for fixed t that the probability i is preferred to j during the k-th comparison is given by

$$\mathrm{P}(Y_{ijk}^{(t)} = 1) = F(\gamma_i^{(t)} - \gamma_j^{(t)},\ \theta,\ \boldsymbol{x}_i^{(t)},\ \boldsymbol{x}_j^{(t)}), \tag{1}$$

where F is a specified probability function monotonically increasing in $\gamma_i^{(t)} - \gamma_j^{(t)}$, θ is a vector of other model parameters, and $\boldsymbol{x}_i^{(t)}$ and $\boldsymbol{x}_j^{(t)}$ are covariate information for objects i and j at time t. This model, the linear paired comparison model, assumes that preference probabilities are functions of the merit parameters only through their difference. When there are no other model parameters or covariates, two common special cases of this model include the Bradley-Terry model (Bradley and Terry, 1952) when F is a standard logistic distribution function, and the Thurstone-Mosteller model (Thurstone, 1927; Mosteller, 1951) when F is a Gaussian distribution function. In practice, the choice among paired comparison models can usually be assessed only with a large amount of data (Stern, 1992).

The model specification in (1) is sufficiently general to include many extensions of basic linear paired comparison models. For example, the inclusion of order effects can be modeled by letting the preference probability be a function of $\gamma_i^{(t)} - \gamma_j^{(t)} + w_{ijk}^{(t)}\delta$, where δ is the order effect, and $w_{ijk}^{(t)}$ is 1 if object i is presented first, and 0 otherwise on the k-th comparison. For example, with a logistic distribution function, this parameterization corresponds to the model by Davidson and Beaver (1977).

The parameters $\gamma_i^{(t)}$ are assumed to change over time through a stochastic process. Our model assumes for each object i,

$$\gamma_i^{(t+1)} \mid \gamma_i^{(t)}, \sigma_i^{2(t+1)} \sim \mathrm{N}(\gamma_i^{(t)},\ \sigma_i^{2(t+1)}), \tag{2}$$

so that the innovations in merit follow a normal distribution centered at 0, and with a variance that depends on time. We further assume a model for the change in variance,

$$\log \sigma_i^{2(t+1)} \mid \sigma_i^{2(t)},\ \tau^2 \sim \mathrm{N}(\log \sigma_i^{2(t)},\ \tau^2). \tag{3}$$

This aspect of the model allows for the innovations in the process for $\gamma_i^{(t)}$ to have a variance that may be stochastically varying.

Because adding the same constant to all the $\gamma_i^{(t)}$ results in an equivalent model specification, an additional assumption is necessary to ensure identifiability. This can be accomplished by assuming

$$\gamma_i^{(0)} | \omega^2 \sim \mathrm{N}(0,\ \omega^2), \tag{4}$$

so that the $\gamma_i^{(t)}$ at $t = 0$ may be viewed as drawn from a common distribution centered at 0.

Assuming I objects and T time periods, the likelihood for this model can be written as

$$\mathrm{L}(\boldsymbol{\gamma}, \boldsymbol{\sigma}^2, \theta, \tau^2, \omega^2 \mid \boldsymbol{y}, \boldsymbol{x}) =$$
$$\left(\prod_{i=1}^{I} \mathrm{N}(\gamma_i^{(0)} \mid 0,\ \omega^2) \right) \times$$

$$\prod_{t=1}^{T}\left(\prod_{i<j}\prod_{k=1}^{K_{ij}^{(t)}} F(\gamma_i^{(t)}-\gamma_j^{(t)},\theta,\boldsymbol{x}_i^{(t)},\boldsymbol{x}_j^{(t)})^{y_{ijk}^{(t)}}(1-F(\gamma_i^{(t)}-\gamma_j^{(t)},\theta,\boldsymbol{x}_i^{(t)},\boldsymbol{x}_j^{(t)}))^{1-y_{ijk}^{(t)}}\right)\times$$
$$\prod_{t=0}^{T-1}\prod_{i=1}^{I}\mathrm{N}(\gamma_i^{(t+1)}\mid\gamma_i^{(t)},\sigma_i^{2(t)})\mathrm{N}(\log\sigma_i^{2(t+1)}\mid\log\sigma_i^{2(t)},\tau^2), \tag{5}$$

where $\boldsymbol{\gamma}$ and $\boldsymbol{\sigma}^2$ are the arrays of the $\gamma_i^{(t)}$ and $\sigma_i^{2(t)}$, and $\mathrm{N}(\cdot|\cdot,\cdot)$ is a normal density of the first argument with the given mean and variance.

Together with (1), equations (2) and (3) form a state-space model with a stochastic variance for paired comparison data. This model can be viewed as an extension of the more usual constant variance model, where $\sigma_i^{2(t)}=\sigma^2$ is assumed for all i and t, in which case (3) is no longer a component of the model. The constant variance model has been analyzed by Glickman (1993), Fahrmeir and Tutz (1994) and Glickman (1999). An important limitation of the constant variance model is that it does not account for the possibility of sudden shifts, innovations, or periods of uncertainty in the process for the $\gamma_i^{(t)}$. For example, in modeling the development of human expertise, one might expect bursts of cognitive development that would not be predicted by the constant variance model. In a marketing context, the worth of a product may change suddenly relative to its competitors, and this change may be poorly described by the constant variance model.

The model for the time-varying variance has close connections with stochastic volatility models from finance (for example, Jacquier et al., 1994). Recent work in modeling of financial time series data has explored the applicability of stochastic volatility models in which the variance of a portfolio index or stock price is assumed to be changing according to (3). A major difference between our model and more conventional stochastic volatility models is that stochastic volatility assumes the variance of observations is changing, whereas in our model the variance of the process governing the merits is changing. Because the observations in our model are the results of preferences, and are therefore binomial, the variance of observations is determined from the mean. Thus, in the paired comparison situation, it would not be meaningful to assume a stochastic process on the observation variance except as it is translated through the process on the merit parameters. It does make sense in our situation to assume that the underlying process on the merits can undergo sudden shifts, and this can be captured through a process assumed on the variance of the merit process.

Model fitting can be accomplished in a Bayesian framework through Markov chain Monte Carlo (MCMC) simulation from the posterior distribution via the Gibbs sampler. A choice of a prior distribution, convenient for model fitting, would assume a product of independent inverse-Gamma densities for τ^2 and ω^2 with low degrees of freedom to reflect initial uncertainty, and a non-informative density (e.g., normal with large variance, or uniform) on θ. Note that the $\gamma_i^{(t)}$ have distributions that are already specified conditionally. Assuming the functional form of F in (1) is tractable (e.g., the Bradley-Terry model, the Thurstone-Mosteller model, or various extensions), then, given τ^2, ω^2, and the $\sigma_i^{2(t)}$ for all i and t, the conditional posterior distribution of the remaining parameters has a form that is common to non-linear state-space models. Recognizing that the Bradley-Terry and the Thurstone-Mosteller models are particular examples of generalized linear models (Critchlow and Fligner, 1991), implementation of the Gibbs sampling steps to simulate the $\gamma_i^{(t)}$ and θ can follow Zeger and Karim (1991), Karim and Zeger (1992), Glickman (1993), and Oh (1997). Sampling the $\sigma_i^{2(t)}$ conditional on the $\gamma_i^{(t)}$, τ^2, and ω^2, a necessary step in MCMC algorithms for fitting stochastic volatility models, is straightforward and can be carried out as in Jacquier et al. (1994). The conditional distributions of τ^2 and ω^2 given the remaining parameters are inverse-Gamma, so that sampling for this step can be performed directly.

3 Example: Best chess players of all time

The model of Section 2 can be applied to a data set consisting of all known results of chess games from 1857 to 1991 played among 88 of the world's all-time best chess players. For chess data, the merit (or strength), $\gamma_i^{(t)}$, of a player can be inferred through game outcomes, which are the results of paired comparisons. The outcome of the k-th comparison between competitors i and j at time t, $y_{ijk}^{(t)}$, is 1 if player i defeats j, and 0 if player j defeats i. The data set, which consists of 15664 outcomes of games played among 1367 pairs of players was compiled by Professor Nathan Divinsky. Not all $\binom{88}{2} = 3828$ pairs of players competed against each other due to non-overlapping chess careers. A detailed account of the data appears in Keene and Divinsky (1989). Several models of chess playing strength have been fit by this data, including models of Elo (1978), Joe (1990), Henery (1992) and Glickman (1999).

For our analysis, we group game outcomes into periods of one year. Thus we act as if all of the games were played simultaneously at the beginning of each year, with innovations in merit and in variance changing over the remainder of the year. The data consist of 135 periods, though some years (e.g., 1859, 1874, and 1875) contain no game outcomes. For years in which no games were recorded, there is no likelihood contribution from data in (5), though the terms for the change in $\gamma_i^{(t)}$ and $\sigma_i^{2(t)}$ still appear.

One aspect of chess outcome data that the model must account for is the existence of ties. Not only does this third possible paired comparison outcome occur in chess, but it occurs frequently. Several extensions to common paired comparisons models have addressed ties as a third outcome, including the extensions by Davidson (1970) and Rao and Kupper (1967) to the Bradley-Terry model, and by Greenberg (1965) for the Thurstone-Mosteller model. Instead of treating a tie as a third outcome to the model, we adopt an approach which acts as if ties are not really observed, but that they are viewed as half the contribution of a win and a loss. In other words, we assume two ties contains the same information about players' strengths as a win followed by a loss (or vice versa). Thus if p_{ij} is the probability that i defeats j, the contribution to the likelihood of a tie would be $\sqrt{p_{ij}(1 - p_{ij})}$. This approach to ties in paired comparisons can also be found in Glickman (1999).

We assume that outcome probability for the k-th game played between players i and j in period t follows the Bradley-Terry model,

$$P(Y_{ijk}^{(t)} = 1) = \frac{\exp(\gamma_i^{(t)})}{\exp(\gamma_i^{(t)}) + \exp(\gamma_j^{(t)})}. \tag{6}$$

Information about the games, such as the player with the first move, time control for the game, or other covariate information was not recorded. Player's ages were not incorporated into the model, because determining an appropriate functional relationship between age and playing strength is beyond the scope of this paper. Rather than fitting a model in which each player has strength parameters for all 135 years, we restrict the model to years from the first to last appearances in the data set. The median career length for the 88 players in this data set is 29.5 years (the distribution is left-skewed because recent players in the data set may still be early in their careers).

In addition to fitting the stochastic variance model of (2) and (3), we also fit a constant variance model assuming

$$\gamma_i^{(t+1)} \mid \gamma_i^{(t)}, \sigma^2 \sim \mathrm{N}(\gamma_i^{(t)},\ \sigma^2), \tag{7}$$

where σ^2 is a single variance governing the change in $\gamma_i^{(t)}$ over time. This particular model was examined by Glickman (1993), Fahrmeir and Tutz (1994) and Glickman (1999). A vague but proper prior distribution is assumed for σ^2 in our analysis.

Both models were fit by MCMC simulation, with burn-in periods of 100,000 iterations at which point the model was diagnosed to have reached stationarity through trace plots and informal diagnostics (e.g., Geweke, 1992). The large number of iterations of MCMC is required because successive parameter draws are highly dependent. Model summaries were computed based on the empirical distribution of simulated parameter values for every 250-th draw for 150,000 iterations beyond the 100,000-th iteration, resulting in 600 draws per parameter. Because of the massive number of parameters per model (2629 for the $\gamma_i^{(t)}$ alone), only 600 parameter values for each parameter were saved due to limitations in disk space.

The two models result in comparable inferences. In general, the marginal posterior distribution of the $\gamma_i^{(t)}$ have large overlap across the two models. The typical trend for an individual's $\gamma_i^{(t)}$ over time is low early in the player's career, a peak in the middle, and then a slow decline towards the end. This finding is consistent with previous analyses of this data set (Joe 1990, Glickman 1999).

It is interesting to examine the comparison between the variance parameters in the two models. A 95% central posterior interval for σ, the standard deviation in the change in $\gamma_i^{(t)}$, in the constant variance model is (0.1223, 0.1541) with a posterior median of 0.1379. By comparison, among the collection of 2629 posterior medians of the $\sigma_i^{(t)}$ in the stochastic variance model, the median value is only 0.1103. Also, the posterior variability of the $\sigma_i^{(t)}$ are, not surprisingly, much larger in the stochastic variance model; the distribution of 2629 posterior medians of the $\sigma_i^{(t)}$ ranges between 0.003505 and 0.5736, and the median width of the corresponding 95% central posterior intervals is 0.37. The large variability in the distribution of posterior medians of the $\sigma_i^{(t)}$ indicates, informally, that the stochastic variance model is accounting for the differing erraticism in strength changes for different players. More formal model selection would involve methods such as in Carlin and Chib (1995).

The contrast between the constant variance and stochastic variance models can be seen by examining model summaries for two of the chess players in the data set. Figure 1 displays summary information of the playing strength for Efim Bogoljubov, former world-champion contender from Russia whose career spanned the 1910s through the 1950s. The top two plots show the posterior mean along with pointwise 95% central posterior intervals over time for Bogoljubov's strength parameters. The first plot shows the summaries from the constant variance model, and the second from the stochastic variance model. Inference about Bogoljubov's strength over time is more precise in the stochastic variance model than

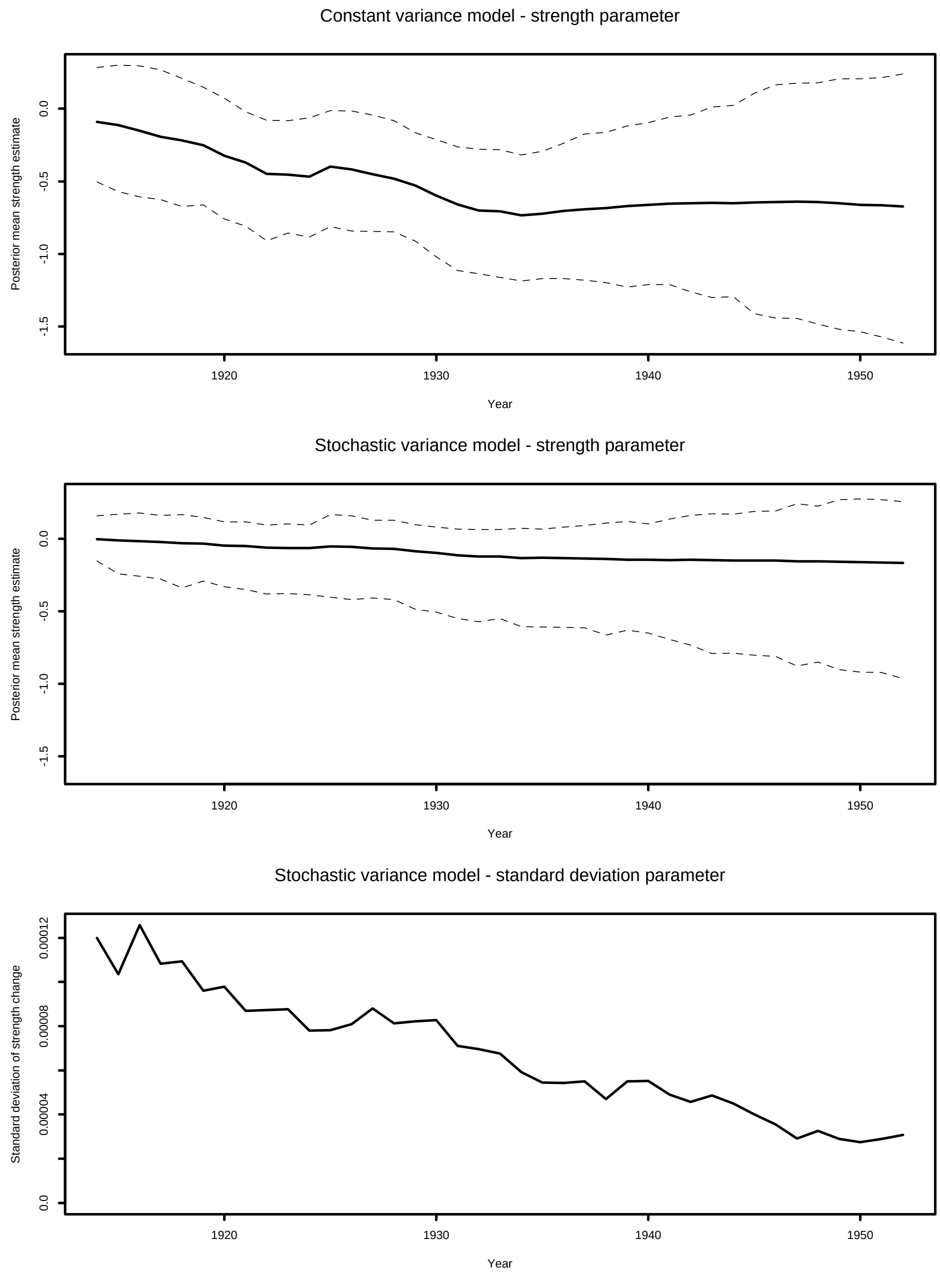

Figure 1: Model summaries for Efim Bogoljubov. *Top*: Posterior means for the $\gamma_i^{(t)}$ in the constant variance model, with pointwise approximate 95% central posterior intervals. *Middle*: Posterior means for the $\gamma_i^{(t)}$ in the stochastic variance model, with pointwise approximate 95% central posterior intervals. *Bottom*: Posterior means for the $\sigma_i^{2(t)}$.

in the constant variance model as evidenced by the narrower 95% posterior intervals. In both cases, the intervals are wider later on in Bogoljubov's career because fewer games were played. The trend in posterior means for the stochastic variance model is smoother than in the constant variance model. As shown in the bottom plot, which displays the posterior means of the variance parameters, the variation in Bogoljubov's strength from year to year was inferred to be quite small (posterior means of the $\sigma_i^{(t)}$ were always less than 0.00013), which accounts both for the smoothness of the strength parameter and the greater precision in the stochastic variance model.

The model summaries of Bogoljubov can be compared with those of José Capablanca, former world champion from Cuba who competed from around 1910 through the late 1930s. Model summaries for Capablanca can be found in Figure 2. In contrast to Bogoljubov, inferences about Capablanca's strength parameters are more precise in the constant variance model than in the stochastic variance model. Capablanca's posterior mean strength is seen to be more erratic in the stochastic variance model, and more smoothly varying in the constant variance model. The opposite relationship to Bogoljubov can be explained by the large inferred values of the $\sigma_i^{(t)}$, as seen in the bottom plot. The posterior means of the $\sigma_i^{(t)}$ for Capablanca range from 0.15 to 0.25, which is the large relative to other players in the data set.

4 Analysis for comparing many objects

When many objects are being compared over time, an exact likelihood-based approach (e.g., maximum likelihood, Bayesian analysis) may become computationally intractable. For example, when rating populations of chess players, or competitors in online gaming systems which can attract tens of thousands of players, a likelihood-based analysis would not be possible to perform in real time. Instead, a simple forward filtering algorithm may be prefer-

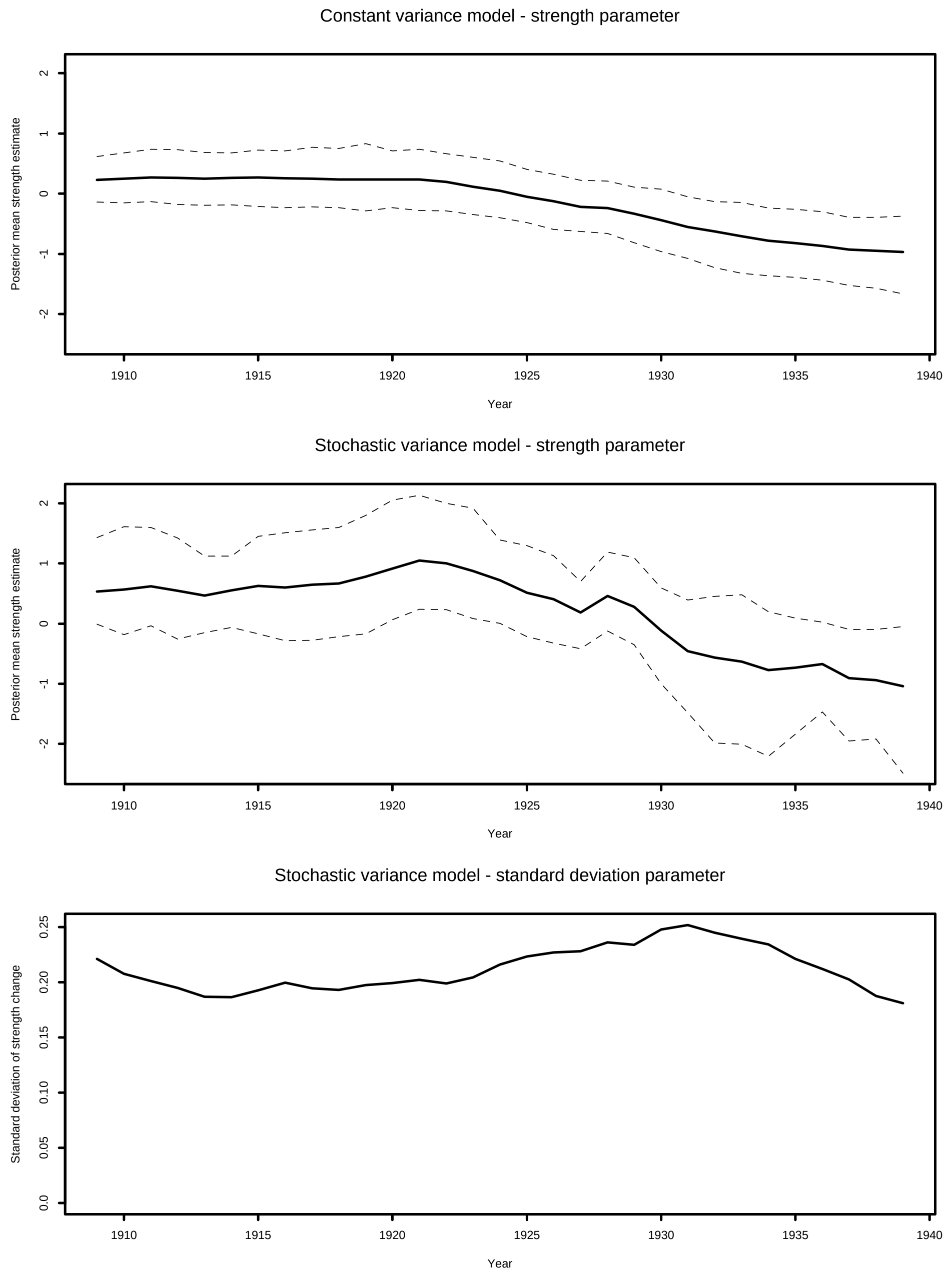

Figure 2: Model summaries for José Capablanca. *Top*: Posterior means for the $\gamma_i^{(t)}$ in the constant variance model, with pointwise approximate 95% central posterior intervals. *Middle*: Posterior means for the $\gamma_i^{(t)}$ in the stochastic variance model, with pointwise approximate 95% central posterior intervals. *Bottom*: Posterior means for the $\sigma_i^{2(t)}$.

able. Glickman (1999) develops an approximate Bayesian analysis for the constant variance model in which the $\gamma_i^{(t)}$ are updated sequentially with the acquisition of new data. The result of this analysis is an approximation to the marginal posterior distribution for the most recent merit parameter for any object. Rather than perform an analysis that jointly infers all parameters simultaneously, the approach taken in Glickman (1999) involves updating the merit parameter of an object by integrating out information about other objects through the prior distribution rather than the posterior distribution. While this approach results in a loss of efficiency, there are computational advantages that allow for the derivation of a simple algorithm. The procedure in Glickman (1999) can be extended to be consistent with the stochastic variance model.

The algorithm to update objects' merit parameters proceeds in the following manner. It is assumed in the algorithm that τ^2 (the variance of the change in $\log \sigma_i^{2(t)}$ over time) and ω^2 (the variance of the $\gamma_i^{(0)}$) have been estimated and fixed in advance. This can be accomplished by fitting the model using an exact likelihood procedure to a set of data of manageable size, and determine the values of τ^2 and ω^2 that maximize the marginal posterior distribution of these two parameters. For the remainder of the development, these parameters are assumed fixed.

1. At the end of the time period $t-1$, each object's merit parameter, $\gamma_i^{(t-1)}$, has an approximating normal marginal posterior distribution with known mean and variance. Also, each object has a current (and known) variance parameter, $\sigma_i^{2(t-1)}$, describing the variability of the time change in merit for that object.

2. Collect all comparisons during time period t (time periods are assumed to be equally spaced).

3. For each object individually, perform appropriate calculations (described below) to determine the updating of $\sigma_i^{2(t-1)}$ to $\sigma_i^{2(t)}$, and then the updating of the distribution of

$\gamma_i^{(t-1)}$ to that of $\gamma_i^{(t)}$.

These calculations are repeated for each time period as data is observed.

To understand the calculations that result in an estimate of $\sigma_i^{2(t)}$, suppose an object in question has a merit $\gamma^{(t-1)}$ that can be summarized at time $t-1$ by the marginal posterior distribution

$$\gamma^{(t-1)} \sim \mathrm{N}(\mu^{(t-1)},\ \phi^{2(t-1)}). \tag{8}$$

Assume that, during period t, this object is compared to others indexed by $j = 1, \ldots, m$, with marginal posterior distributions

$$\gamma_j^{(t-1)} \sim \mathrm{N}(\mu_j^{(t-1)},\ \phi_j^{2(t-1)}). \tag{9}$$

The stochastic variance model assumes

$$\gamma^{(t)} \sim \mathrm{N}(\gamma^{(t-1)},\ \sigma^{2(t)}) \tag{10}$$

$$\log \sigma^{2(t)} \sim \mathrm{N}(\log \sigma^{2(t-1)},\ \tau^2) \tag{11}$$

Further, let $\hat{\gamma}^{(t)}$ be the maximum value of $\gamma^{(t)}$ in the likelihood for period t, integrated over the *prior* distribution of the other objects' $\gamma_j^{(t)}$, and let ν^2 be the associated asymptotic variance of $\hat{\gamma}^{(t)}$ from the marginalized likelihood. When integrating over the distribution of the other objects' merits, we assume $\log \sigma_j^{2(t)} = \log \sigma_j^{2(t-1)}$ for object j, even though this is only true in expectation. Both $\hat{\gamma}^{(t)}$ and ν^2 can be computed using iterative numerical procedures (though an approximation is used in the algorithm that follows). The distribution of $\hat{\gamma}^{(t)}$ can be approximated by

$$\hat{\gamma}^{(t)} \sim \mathrm{N}(\gamma^{(t)},\ \nu^2). \tag{12}$$

Combining (8), (10) and (12), integrating out $\gamma^{(t-1)}$ and $\gamma^{(t)}$, and for notational convenience letting $\alpha^{(t-1)} = \log \sigma^{2(t-1)}$ and $\alpha^{(t)} = \log \sigma^{2(t)}$, we have

$$\hat{\gamma}^{(t)} \sim \mathrm{N}(\mu^{(t-1)},\ \phi^{2(t-1)} + \exp(\alpha^{(t)}) + \nu^2). \tag{13}$$

From (11), we have

$$\alpha^{(t)} \sim \mathrm{N}(\alpha^{(t-1)},\ \tau^2) \tag{14}$$

Noting that all other parameters are known, the approximate marginal posterior density of $\alpha^{(t)}$ is the product of the densities in (13) and (14), so that the marginal log-posterior, up to an additive constant, is given by

$$\begin{aligned} \log p(\alpha^{(t)} \mid \boldsymbol{y}^{(t)}) &= -\frac{1}{2}\frac{(\alpha^{(t)}-\alpha^{(t-1)})^2}{\tau^2} \\ &\quad -\frac{1}{2}\log(\phi^{2(t-1)}+\exp(\alpha^{(t)})+\nu^2) - \frac{1}{2}\frac{(\hat{\gamma}^{(t)}-\mu^{(t-1)})^2}{\phi^{2(t-1)}+\exp(\alpha^{(t)})+\nu^2} \end{aligned} \tag{15}$$

where $\boldsymbol{y}^{(t)}$ denotes the collection of comparison outcomes during time period t. Rather than compute $\hat{\gamma}^{(t)}$ and ν^2 numerically, we use approximations derived in Glickman (1999). In particular, we approximate $(\hat{\gamma}^{(t)} - \mu^{(t-1)})$ by a Taylor series expansion through the linear term, and approximate ν^2 by the curvature around $\mu^{(t-1)}$ rather than $\hat{\gamma}^{(t)}$. This yields

$$\hat{\gamma}^{(t)} - \mu^{(t-1)} \approx \nu^2 \sum_{j=1}^{m}\sum_{k=1}^{n_j} g(\phi_j^{2(t-1)})\{y_{jk}^{(t)} - \mathrm{E}(y|\mu^{(t-1)}, \mu_j^{(t-1)}, \phi_j^{2(t-1)})\} \tag{16}$$

with

$$\nu^2 \approx \left[\sum_{j=1}^{m} n_j g(\phi_j^{2(t-1)})^2 \mathrm{E}(y|\mu^{(t-1)}, \mu_j^{(t-1)}, \phi_j^{2(t-1)})\{1 - \mathrm{E}(y|\mu^{(t-1)}, \mu_j^{(t-1)}, \phi_j^{2(t-1)})\}\right]^{-1} \tag{17}$$

where $y_{jk}^{(t)}$ is the result of the k-th comparison of the object with object j during period t, n_j is the number of times the object is compared to object j, and

$$\begin{aligned} g(\phi^2) &= \frac{1}{\sqrt{1+3\phi^2/\pi^2}}, \\ \mathrm{E}(y|\mu,\mu_j,\phi_j^2) &= \frac{1}{1+\exp(-g(\phi_j^2)(\mu-\mu_j))}. \end{aligned}$$

The algorithm proceeds by estimating $\alpha^{(t)}$ (and therefore $\sigma^{2(t)}$) by maximizing over (15). This can be accomplished using a numerical algorithm. For example, because the first term in (15) is maximized by $\alpha^{(t)} = \alpha^{(t-1)}$, and the second two terms in (15) are maximized by setting

$$(\hat{\gamma}^{(t)} - \mu^{(t-1)})^2 = \phi^{2(t-1)} + \exp(\alpha^{(t)}) + \nu^2$$

so that $\alpha^{(t)} = \log((\hat{\gamma}^{(t)} - \mu^{(t-1)})^2 - \phi^{2(t-1)})$ assuming

$$(\hat{\gamma}^{(t)} - \mu^{(t-1)})^2 - \phi^{2(t-1)} - \nu^2 > 0, \tag{18}$$

the Newton-Raphson algorithm will converge to the maximum quickly if the initial value of $\alpha^{(t)}$ is selected to be between $\alpha^{(t-1)}$ and $\log((\hat{\gamma}^{(t)} - \mu^{(t-1)})^2 - \phi^{2(t-1)})$. If (18) is false, then the second two terms reach their supremum when $\exp(\alpha^{(t)})$ is set to 0. In this situation, because these latter two terms are bounded above as $\alpha^{(t)} \to -\infty$, the first term dominates, and convergence of the Newton-Raphson algorithm is quick when choosing an initial value of $\alpha^{(t)}$ less than $\alpha^{(t-1)}$.

Once $\sigma^{2(t)}$ is estimated, we set $\phi'^{2(t)} = \phi^{2(t-1)} + \sigma^{2(t)}$, which is the prior variance for $\gamma^{(t)}$ accounting for the passage of time from period $t-1$ to t. Now the algorithm of Glickman (1999) may be applied directly to obtain $\mu^{(t)}$ and $\phi^{2(t)}$, the marginal posterior mean and variance of $\gamma^{(t)}$, the merit parameter for the object in question. These are given by

$$\begin{aligned}
\mu^{(t)} &= \mu^{(t-1)} + \frac{1}{1/\phi'^{2(t)} + 1/\nu^2} \sum_{j=1}^{m} \sum_{k=1}^{n_j} g(\phi_j^{2(t-1)}) \{ y_{jk}^{(t)} - \mathrm{E}(y | \mu^{(t-1)}, \mu_j^{(t-1)}, \phi_j^{2(t-1)}) \} && (19) \\
\phi^{2(t)} &= \left(\frac{1}{\phi'^{2(t)}} + \frac{1}{\nu^2} \right)^{-1}. && (20)
\end{aligned}$$

Details of the derivations appear in Glickman (1999).

To assess the accuracy of the approximation algorithm, simulated data under varying parameter values were generated, and nominal coverage was compared with the results of simulations. We performed a total of 32 simulation sets which are summarized in Table 1. We considered two values (0, 0.5) for the prior mean, μ, four different numbers of comparisons, m, (10, 50, 200, 1000), and four different values of the standard deviation for the change in $\log \sigma^{2(t)}$, τ (0, 0.3, 0.7, 1.2). These values were chosen to span plausible range of values that might be expected in practice. For each fixed combination of μ, m and τ, data were generated in the following manner: Values of $\phi^{(t-1)}$ and $\sigma^{(t-1)}$ were fixed at 0.173 and 0.0576, respectively, so that a 95% prior interval around $\gamma^{(t-1)}$ had length 2/3, and that the standard

Prior Mean	Number of Comparisons	$\tau = 0$	$\tau = 0.3$	$\tau = 0.7$	$\tau = 1.2$
$\mu = 0$	$m = 10$	(.5210, .9460)	(.4928, .9476)	(.4892, .9450)	(.4886, .9414)
	$m = 50$	(.4942, .9498)	(.5026, .9562)	(.5042, .9504)	(.4794, .9414)
	$m = 200$	(.5074, .9500)	(.4890, .9514)	(.4822, .9450)	(.4840, .9466)
	$m = 1000$	(.4896, .9484)	(.4922, .9482)	(.4894, .9402)	(.4908, .9448)
$\mu = 0.5$	$m = 10$	(.5048, .9524)	(.4912, .9504)	(.4968, .9440)	(.4932, .9384)
	$m = 50$	(.5082, .9486)	(.5044, .9460)	(.5060, .9484)	(.4924, .9398)
	$m = 200$	(.4984, .9466)	(.5034, .9540)	(.4924, .9472)	(.4930, .9436)
	$m = 1000$	(.4988, .9498)	(.4918, .9472)	(.4914, .9446)	(.4868, .9460)

Table 1: Results of the approximating algorithm on simulated data. For each analysis, 5000 replications were simulated and nominal 50% and 95% central posterior intervals were constructed. The pairs of values in the parentheses for each analysis consist of the proportion of 5000 replications in which the simulated value of γ was contained in the nominal 50% and 95% intervals, respectively.

deviation of the $\gamma^{(t-1)}$ was three times larger than the standard deviation in the change of the $\gamma^{(t-1)}$ over time. A value of $\sigma^{2(t)}$ was simulated conditional on $\sigma^{2(t-1)}$ and τ, and then a value of $\gamma^{(t)}$ was simulated given $\mu^{(t-1)}$, $\phi^{(t-1)}$ and $\sigma^{2(t)}$. For the m other objects involved in the comparisons, the collection of $\mu_j^{(t-1)}$ were simulated from a normal distribution with mean 0 and standard deviation 0.173. The $\phi_j^{(t-1)}$ and $\sigma_j^{(t)}$ were generated from scaled χ^2 distributions on 20 degrees of freedom with means of 0.173 and 0.0576, respectively. Values of $\gamma_j^{(t)}$ were generated by the same process as $\gamma^{(t)}$. Finally, the outcome of comparisons were generated from the Bradley-Terry model given the generated values of the $\gamma^{(t)}$ and the $\gamma_j^{(t)}$. The algorithm was then applied (ignoring parameter values at time t) to determine the parameters $\mu^{(t)}$ and $\phi^{(t)}$ of the approximating normal posterior to $\gamma^{(t)}$. Approximate 50% and 95% normal central posterior interval for $\gamma^{(t)}$ were calculated as $\mu^{(t)} \pm z\phi^{(t)}$ with $z = 0.6745$ or 1.96. It was noted whether the generated value of $\gamma^{(t)}$ was contained in this interval. This process was repeated 5000 times, and the fraction of times in which the true parameter value was contained in the intervals is summarized in Table 1.

The results of the simulations demonstrate that the algorithm produces close to nominal coverage under varying conditions. Table 1 reveals that the nominal 50% and 95% posterior intervals contain roughly 0.5 and 0.95 of the generating value of $\gamma^{(t)}$. The accuracy of the approximation algorithm does not seem to change by varying μ. However, there appears to be a small loss of efficiency when m becomes larger, and when τ is at its largest. In these cases, the actual coverage is consistently smaller than nominal coverage, indicating that the posterior intervals are not wide enough. The discrepancy does not seem large enough to be of great practical concern.

5 Discussion

The dynamic paired comparison model presented in this paper extends previous work by allowing the variance of the state process to change stochastically. Considering such models increases the flexibility to describe phenomena where the underlying characteristics may undergo sudden shifts or changes in paradigm. Fitting the stochastic variance model for paired comparison data can be carried out using standard Bayesian computational machinery. For situations where many objects are being compared over time, in which case a full likelihood analysis may be too computationally intensive, this paper demonstrates an approximating algorithm that can be carried out with far less of a computational burden.

Aside from the increased flexibility, the stochastic variance model has an important benefit over the constant variance model when, for example, it is suspected that interventions may affect the process for the merits. Under the constant variance model, the variance in an object's merit prior to observing data at time t must decrease after the data is observed. This is a direct consequence of (20) where the posterior variance must be less than the prior variance. However, under the stochastic variance model, the variance may *increase* after data is observed. This reflects the increased uncertainty about the merit parameter after observing unusual data.

A variety of extensions can be incorporated into the stochastic variance model. One extension is to incorporate covariate information in the change in $\sigma_i^{2(t)}$ over time. For example, in the context of human cognitive development, older people may stabilize in merit so that the change in $\sigma_i^{2(t)}$ may be assumed to be negatively related with age. Another extension involves using a different distributions other than normal for describing the change in $\gamma_i^{(t)}$, and log-normal for describing the change in $\sigma_i^{2(t)}$ over time. For example, in comparing certain types of financial products over time, it may be more reasonable to assume that there is a greater probability for small increases in $\gamma_i^{(t)}$ but occasional large decreases with small probability. In each case, standard Bayesian tools can still be invoked to fit such model extensions.

Bibliography

Bradley R. A. (1984) Paired comparisons: Some basic procedures and examples. In *Handbook of Statistics* 4, eds. P. R. Kirshnajah and P. K. Sen, Amsterdam: Elsevier, pp. 299–326.

Bradley R. A., and Terry, M. E. (1952) The rank analysis of incomplete block designs. 1. The method of paired comparisons. *Biometrika*, **39**, 324–45.

Carlin, B. P. and Chib, S. (1995) Bayesian model choice via Markov chain Monte Carlo methods. *J. Roy. Stat. Soc. B*, **57**, 473–484.

Capobianco, E. (1996) State-space stochastic volatility models: A review of estimation algorithms. *Appl. Stoch. Models Data Anal.*, **12**, 265–279.

Critchlow, D. E. and Fligner, M. A. (1991) Paired comparison, triple comparison, and ranking experiments as generalized linear models, and their implementation on GLIM. *Psychometrika*, **56**, 517–533.

David, H. A. (1988) *The method of paired comparisons* (2nd ed.), London: Chapman and Hall.

Davidson, R. R. (1970) On extending the Bradley-Terry model to accommodate ties in paired comparison experiments. *J. Am. Statist. Ass.*, **65**, 317–28.

Davidson, R. R. and Beaver, R. J. (1977) On extending the Bradley-Terry model to incorporate within-pair order effects. *Biometrics*, **33**, 693–702.

Elo, A. E. (1978) *The rating of chess players past and present*, New York: Arco Publishing.

Fahrmeir, L. and Tutz, G. (1994) Dynamic stochastic models for time-dependent ordered paired comparison systems. *J. Am. Statist. Ass.*, **89**, 1438–1449.

Geweke, J. (1992) Evaluating the accuracy of sampling-based approaches to calculating posterior moments. In *Bayesian Statistics 4*, (ed. J. M. Bernardo, J. O. Berger, A. P. Dawid, and A. F. M. Smith). Clarendon Press, Oxford, UK.

Glickman, M. E. (1993) *Paired comparison models with time-varying parameters.* PhD Dissertation, Harvard University Department of Statistics, Cambridge, USA.

Glickman, M. E. (1999) Parameter estimation in large dynamic paired comparison experiments. *Appl. Statist.*, **48**, 377–394.

Greenberg, M. G. (1965) A modification of Thurstone's law of comparative judgment to accommodate a judgment category of "equal" or "no difference." *Psychol. Bull.*, **64**, 108–112.

Henery, R. J. (1992) An extension to the Thurstone-Mosteller model for chess. *The Statistician*, **41**, 559–567.

Jacquier, E., Polson, N., and Rossi, P. (1994) Bayesian analysis of stochastic volatility models (with discussion) *J. Bus. Econ. Statist.*, **12**, 371–389.

Joe, H. (1990) Extended use of paired comparison models, with application to chess rankings. *Appl. Statist.*, **39**, 85–93.

Karim, M. and Zeger, S. L. (1992) Generalized linear models with random effects; Salamander mating revisited. *Biometrics*, **48**, 631–644.

Keene, R. and Divinsky, N. (1989) *Warriors of the mind: A quest for the supreme genius of the chess board.* Brighton (England): Hardinge Simpole.

Mosteller, F. (1951) Remarks on the method of paired comparisons: I. The least squares solution assuming equal standard deviations and equal correlations. *Psychometrika*, **16**, 3–9.

Oh, M. (1997) A Gibbs sampling approach to Bayesian analysis of generalized linear models for binary data. *Computat. Statist.*, **12**, 431–445.

Rao, P. V. and Kupper, L. L. (1967) Ties in Paired-Comparison Experiments: A Generalization of the Bradley-Terry Model. *J. Am. Statist. Ass.*, **62**, 194–204.

Simonton, D. K. (1997) Creative productivity: A predictive and explanatory model of career trajectories and landmarks. *Psychological Review*, **104**, 66–89.

Stern, H. (1992) Are all linear paired comparison models empirically equivalent? *Math. Soc. Sci.*, **23**, 103–117.

Thurstone, L. (1927) A law of comparative judgment. *Psychological Review*, **34**, 273–286.

Uhlig, H. (1997) Bayesian vector autoregressions with stochastic volatility. *Econometrica*, **65**, 59–73.

Zeger, S. L. and Karim, M. (1991) Generalized linear models with random effects: A Gibbs sampling approach. *J. Am. Statist. Ass.*, **86**, 79–86.