Time-to-Default Analysis of Mortgage Portfolios

M. Galloway [†], A. Johnson [‡] and A. Shemyakin [§*]

Abstract

Fluctuation in mortgage default rates provides vital information to financial institutions and is a key indicator of the state of the economy. Using a decade's worth (2002–2010) of data on prime and subprime mortgage portfolios, we propose and compare two models for mortgage defaults. The first, the Weibull-Gamma segmentation model (WGS), was utilized by Fader and Hardie (2007) in forecasting customer retention. Though effective in that setting, Markov chain Monte Carlo simulations suggest that the WGS suffers from over-parameterization. The Weibull segmentation model (WS) provides a simplified alternative that accurately forecasts default rates while identifying latent classes of "risky" prime and subprime mortgages characterized by increased hazard rates.

Keywords: *mortgage default, time-to-default, mixture model, latent class, Bayesian estimation.*

1 Introduction

Mortgage defaults played a critical role in the 2007–2010 financial crisis. Often referred to as the "subprime mortgage crisis," this period saw an overwhelming number of defaults in the subprime sector of the mortgage market. Accordingly, subprime lending practices have received great attention in the literature. For example, Das and Stein (2009) discuss the impact of underwriting standards and Demyanyk and Van Hemert (2011) explore potential economical and institutional causes. Crucial to these conversations is an understanding of the inherent risk in subprime lending.

Here we propose and compare two Bayesian models of time-to-default (TD) for a mortgage portfolio: an infinite mixture Weibull-Gamma segmentation model (WGS) and a finite mixture Weibull segmentation model (WS). The WGS formally addresses two key considerations in mortgage defaults: (1) potential heterogeneity in default rates among mortgage holders; and (2) the existence of latent mortgage market classes with differing temporal risk patterns. For example, there may exist two latent classes of mortgage holders, one for which hazard rates increase with time and the other for which hazard rates decrease. The WS is a simplification of the WGS which addresses

[†] University of Minnesota, Minneapolis, MN, USA

[‡] Macalester College, St. Paul, MN, USA

[§] University of St. Thomas, St. Paul, MN, USA

[*] corresponding author a9shemyakin@stthomas.edu

(2) while assuming a common default rate among mortgage holders. Similar strategies have been utilized to model heterogeneous survival data in various other contexts (see, for example, Marin et al. (2005), Erisoglu et al. (2011), and Soyer and Xu (2010)). In fact, the WGS and WS are extensions of the models proposed by Fader and Hardie (2007), Hardie et al. (1998), and Fader et al. (2003) for addressing population heterogeneity of customer retention in subscription settings.

Specification of the WGS and WS requires the estimation of corresponding model parameters, θ. The convention in financial and marketing applications is to approximate θ via maximum likelihood estimation. This *frequentist* approach assumes a fixed θ and ignores practitioners' prior understanding of θ compiled from subject-matter expertise and experience. We present a *Bayesian* alternative that evaluates θ through a weighted combination of observed data and prior knowledge. Though not standard, Bayesian applications in survival analysis settings are increasingly popular (see Chen et al. (1985) and Kiefer (2006, 2008)). For example, a Bayesian reliability model for mortgage default risk is suggested in Soyer and Xu (2010) and a Bayesian state space model is discussed in Aktekin et al. (2013). Most relevant to our exploration of default rates is the work of Popova et al. (2008) that conducts a Bayesian analysis of loan prepayment.

Our development of the WGS and WS is motivated by our receipt of data from a major American commercial bank (as reported to Federal Financial Institution Examination Council). These data include the monthly number of defaults in prime and subprime mortgage portfolios pooled over several cohorts, ie. vintages. In this particular setting, our results suggest that the simpler WS is superior to the WGS in modeling TD. Mainly, while these results support the segmentation of mortgages into two latent classes, heterogeneity among default rates appears to be insignificant.

The paper is organized as follows. In Section 2 we provide relevant background information on the provided mortgage data and the foundational Weibull-Gamma time-to-default model. We specify our Bayesian WGS and WS models in Section 3 and apply these to the provided mortgage data in Section 4.

2 Background

2.1 Mortgage Default Data

A major American bank provided the authors with data on the aggregated prime and subprime portfolios of residential mortgages on the U.S. market in 2002–2010. Included for the prime and subprime portfolios separately, was an indication of the total number of mortgages m and the monthly counts of defaults $k = (k_1, k_2, ..., k_{90})$ pooled over several vintages. From k we calculated times-to-default for all m mortgage holders, $(t_1, t_2, ..., t_m)$, discretized to month. Given the finite period of observation (90 months), the majority of the t_i are right-censored. These data are summarized in Figure 1. The inherent risk of subprime vs prime mortgage lending is clear.

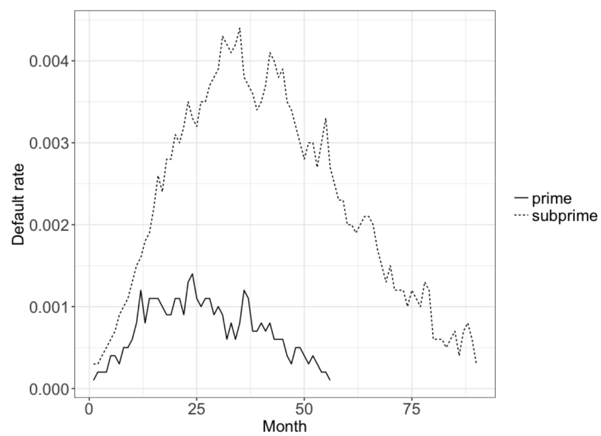

Figure 1: Proportion of mortgage holders that defaulted in each of the first 90 months of the holding.

2.2 The Weibull-Gamma Model

Our goal is to construct time-to-default models from which to predict long-term losses of the prime and subprime portfolios. In using pooled data, these models integrate out potential macroeconomic factors. Thus they provide a characterization of the bank's underwriting practices rather than specific economic circumstances. For vintage-specific analysis see Glennon and Nigro (2011) and Aktekin et al. (2013) which emphasizes the short-term role of covariates.

Let T denote the time-to-default for a mortgage holder in a given portfolio. A simple model for T can be derived from the discrete-time Beta-Geometric (BG) model for customer retention proposed by Fader and Hardie (2007):

$$\begin{aligned} T|\theta &\sim \text{Geo}(\theta) \\ \theta &\sim \text{Beta}(a, b) \end{aligned}$$

In the context of mortgage defaults, the BG model assumes the following:

(A1) Defaults are observed at discrete periods, i.e. $T \in \{1, 2, 3, \ldots\}$.

(A2) The probability of default for an individual mortgage holder, θ, remains constant throughout their holding. Thus $T|\theta \sim \text{Geo}(\theta)$ with distribution function

$$F_{BG}(t|\theta) = Pr(T \leq t|\theta) = 1 - (1-\theta)^t \quad \text{for } t \in \{1, 2, ...\} \ .$$

(A3) Variability of θ among mortgage holders can be characterized by $\theta \sim \text{Beta}(a, b)$ with probability density function $f(\theta|a, b) \propto \theta^{a-1}(1-\theta)^{b-1} \quad \text{for } \theta \in [0, 1]$.

Note that the BG is an infinite mixture of Geometric distributions, with parameter θ characterizing the random retention pattern of an individual mortgage holder. Thus the BG addresses potential population heterogeneity in default rates. Yet the BG assumptions oversimplify the reality of our setting. To begin, T is typically continuous, i.e. defaults can occur at any time. More importantly, depending on the circumstances of the mortgage portfolio and its holder, the hazard rate of default may increase or decrease throughout the duration of the holding. To accommodate these features, Fader and Hardie (2007) propose, but do not explore, a continuous-time Weibull-Gamma (WG) model:

$$\begin{aligned} T|\lambda, c &\sim \text{Weibull}(\lambda, c) \\ \lambda &\sim \text{Gamma}(r, \alpha) \end{aligned} \tag{1}$$

Specifically, the WG assumes the following:

(B1) Defaults can occur at any time, i.e. $T \geq 0$.

(B2) The risk of default for an individual may increase or decrease throughout their holding. Thus T can be characterized by $T|\lambda, c \sim \text{Weibull}(\lambda, c)$ with probability density function $f(t|\lambda, c) \propto t^{c-1}e^{-\lambda t^c}$ and distribution function

$$F(t|\lambda, c) = Pr(T \leq t|\lambda, c) = 1 - e^{-\lambda t^c} \quad \text{for } t \geq 0 \; .$$

The default risk factor of an individual mortgage holder, $\lambda > 0$, captures the *scale* of the Weibull with larger λ reflecting a higher overall risk of default. Further, *shape* parameter $c > 0$ reflects the magnitude and direction of the *hazard rate* (risk of default) over time. Specifically, $c = 1$ indicates a constant hazard rate and $c > 1$ ($c < 1$) indicates an increasing (decreasing) hazard rate.

(B3) The risk of default varies among mortgage holders. To this end, heterogeneity in λ is modeled by $\lambda \sim \text{Gamma}(r, \alpha)$ distribution with probability density function $f(\lambda|r, \alpha) \propto \lambda^{r-1}e^{-\alpha\lambda}$ for $\lambda > 0$.

It follows that the joint Weibull-Gamma distribution of (T, λ) is characterized by density function

$$f_{WG}(t, \lambda|c, r, \alpha) = f(t|\lambda, c)f(\lambda|r, \alpha) \; .$$

Thus integrating out λ produces a form of the Burr Type XII distribution:

$$F_{WG}(t|c, r, \alpha) = \int_0^t \int_0^\infty f_{WG}(z, \lambda|c, r, \alpha) d\lambda dz = 1 - \left(\frac{\alpha}{\alpha + t^c}\right)^r \; .$$

3 Latent Class Segmentation Models

3.1 Finite Mixtures

The WG framework under assumptions (B1)–(B3) provides a foundation for modeling time-to-default, yet it does not address the potential existence of mortgage market classes with differing temporal risk patterns. In reality, it is reasonable to assume that the mortgage population could be separated into two or more "segments" with differing hazard rates, some that increase over time and others that decrease. Since the segment membership of any given mortgage holder is unknown, these segments form latent classes.

Here we assume the existence of two latent mortgage classes, though our methodology can be extended to accommodate any number of classes. To this end, let $p \in (0, 1)$ be the segmentation weight for the first class of mortgages with shape parameter c_1, and $1 - p$ be the weight for the second class of mortgages with shape c_2. Then we assume that times-to-default $(T_1, T_2, ..., T_m)$ follow a mixture Weibull distribution

$$T_i|\lambda, c_1, c_2, p \overset{ind}{\sim} p\text{Weibull}(\lambda, c_1) + (1 - p)\text{Weibull}(\lambda, c_2) \tag{2}$$

with probability density function and distribution function

$$\begin{aligned} f_{WS}(t_i|\lambda, c_1, c_2, p) &= p\lambda c_1 t_i^{c_1 - 1} e^{-\lambda t_i^{c_1}} + (1 - p)a\lambda c_2 t_i^{c_2 - 1} e^{-\lambda t_i^{c_2}} \\ F_{WS}(t_i|\lambda, c_1, c_2, p) &= 1 - pe^{-\lambda t_i^{c_1}} - (1 - p)e^{-\lambda t_i^{c_2}} \end{aligned}$$

Recall that the observed times-to-default $t = (t_1, t_2, ..., t_m)$ are discretized by month and k_j denotes number of defaults in month $j \in \{1, \ldots, L\}$. The majority of mortgage holders did not default in the $L = 90$ month period, thus are right censored. To this end, let $\delta = (\delta_1, \delta_2, \ldots, \delta_m)$ be indicators of right-censorship (1 if censored and 0 otherwise) and note that the total number of right-censored mortgages is $\sum_{i=1}^{m} \delta_i = m - \sum_{j=1}^{L} k_j$. Thus a common technique in survival analysis of time-to-event data is to rewrite the probability density function to reflect right-censoring:

$$\tilde{f}_{WS}(t_i|\lambda, c_1, c_2, p) = f_{WS}(t_i|\lambda, c_1, c_2, p)^{1-\delta_i} \left(1 - F_{WS}(L|\lambda, c_1, c_2, p)\right)^{\delta_i} .$$

The likelihood function of (λ, c_1, c_2, p) follows:

$$\begin{aligned}
\mathcal{L}_{WS}(\lambda, c_1, c_2, p \mid t) &= \prod_{i=1}^{m} \tilde{f}_{WS}(t_i|\lambda, c_1, c_2, p) \\
&= \prod_{i=1}^{m} f_{WS}(t_i|\lambda, c_1, c_2, p)^{1-\delta_i} \cdot \prod_{i=1}^{m} (1 - F_{WS}(L|\lambda, c_1, c_2, p))^{\delta_i} \\
&= \left[\prod_{j=1}^{L} f_{WS}^{k_j}(j|\lambda, c_1, c_2, p) \right] (1 - F_{WS}(L|\lambda, c_1, c_2, p))^{m - \sum_{j=1}^{L} k_j} \ .
\end{aligned} \tag{3}$$

Note that the final line is a consequence of the discretization of time-to-default to months.

In finance and marketing applications, the convention is to estimate (λ, c_1, c_2, p) via maximum likelihood:

$$(\hat{\lambda}, \hat{c_1}, \hat{c_2}, \hat{p}) = \text{argmax}_{\lambda, c_1, c_2, p} \mathcal{L}_{WS}(\lambda, c_1, c_2, p \mid t) \ .$$

In contrast, we propose two Bayesian alternatives corresponding to the Weibull segmentation (WS) and Weibull-Gamma segmentation (WGS) time-to-default models.

3.2 The Weibull Segmentation Model

The mixture Weibull distribution (2) provides the foundation of the WS Bayesian time-to-default model. Specifically, for mortgages $i \in \{1, ..., m\}$, the WS assumes

$$T_i|\lambda, c_1, c_2, p \overset{ind}{\sim} p\text{Weibull}(\lambda, c_1) + (1-p)\text{Weibull}(\lambda, c_2)$$

with independent priors

$$\begin{aligned}
\lambda &\sim \text{Gamma}(0.00001, 0.00001) \\
c_i &\sim \text{Gamma}(0.00001, 0.00001) \ \text{ for } i \in \{1, 2\} \\
p &\sim \text{Beta}(0.5, 0.5) \ .
\end{aligned}$$

We have chosen non-informative priors here, but these could be modified to incorporate subject-matter expertise in future iterations. Letting $f(\lambda), f(c_1), f(c_2), f(p)$ denote the corresponding prior probability density functions and for likelihood $\mathcal{L}_{WS}$ defined by (3), the posterior distribution of (λ, c_1, c_2, p) given data t is characterized by probability density function

$$f_{\text{WS}_{\text{post}}}(\lambda, c_1, c_2, p|t) \propto \mathcal{L}_{WS}(\lambda, c_1, c_2, p \mid t) f(\lambda) f(c_1) f(c_2) f(p) \ .$$

3.3 The Weibull-Gamma Segmentation Model

The WS model assumes homogeneity in λ, ie. that default rates are constant across mortgage holders. The WGS relaxes this assumption, utilizing assumption (B3) of the WG model (1):

$$\begin{aligned}
T_i|\lambda, c_1, c_2, p &\overset{ind}{\sim} p\text{Weibull}(\lambda, c_1) + (1-p)\text{Weibull}(\lambda, c_2) \\
\lambda|r, \alpha &\sim \text{Gamma}(r, \alpha) \\
c_i &\sim \text{Gamma}(0.00001, 0.00001) \;\; \text{for } i \in \{1, 2\} \\
p &\sim \text{Beta}(0.5, 0.5) \\
r &\sim \text{InvGamma}(0.000000001, 0.000000001) \\
\alpha &\sim \text{InvGamma}(0.000000001, 0.000000001)\ .
\end{aligned}$$

The corresponding posterior distribution of $(\lambda, c_1, c_2, p, r, \alpha)$ given data t is characterized by probability density function

$$f(\lambda, c_1, c_2, p, r, \alpha|t) \propto \mathcal{L}_{WS}(\lambda, c_1, c_2, p \mid t) f(\lambda|r, \alpha) f(c_1) f(c_2) f(p) f(r) f(\alpha)\ . \tag{4}$$

To obtain the marginal posterior of (c_1, c_2, p, r, α) given t we can integrate out λ:

$$\begin{aligned}
f_{\text{WGS}_{post}}(c_1, c_2, p, r, \alpha|t) &= \int_0^\infty f(\lambda, c_1, c_2, p, r, \alpha|t) d\lambda\ . \\
&\propto \mathcal{L}_{WGS}(c_1, c_2, p, r, \alpha \mid t) f(c_1) f(c_2) f(p) f(r) f(\alpha)
\end{aligned}$$

where

$$\mathcal{L}_{WGS}(c_1, c_2, p, r, \alpha \mid t) = \prod_{j=1}^{L} f_{WGS}^{k_j}(j|c_1, c_2, p, r, \alpha)(1 - F_{WGS}(L|c_1, c_2, p, r, \alpha))^{m - \sum_{j=1}^{L} k_j} \tag{5}$$

for

$$\begin{aligned}
f_{WGS}(t|c_1, c_2, p, r, \alpha) &= \int_0^\infty f_{WS}(t|\lambda, c_1, c_2, p) f(\lambda|r, \alpha) d\lambda \\
F_{WGS}(t|c_1, c_2, p, r, \alpha) &= \int_0^t f_{WGS}(z|c_1, c_2, p, r, \alpha) dz = 1 - p\left(\frac{\alpha}{\alpha + t^{c_1}}\right)^r - (1-p)\left(\frac{\alpha}{\alpha + t^{c_2}}\right)^r\ .
\end{aligned}$$

4 Application to Default Data

4.1 Parameter Estimation

Fitting the WS and WGS to the provided time-to-default data for prime and subprime portfolios (separately) requires estimation of the corresponding model parameters. Let $f_{\text{post}}(\theta|t)$ denote a generic posterior probability density function of parameters θ. For example, the WS and WGS have $f_{\text{WS}_{\text{post}}}(\lambda, c_1, c_2, p|t)$ and $f_{\text{WGS}_{\text{post}}}(c_1, c_2, p, r, \alpha|t)$, respectively. The posterior expectation provides a point estimate of θ:

$$E[\theta|t] = \int \theta f_{\text{post}}(\theta|t) d\theta \ .$$

Yet when $f_{\text{post}}(\cdot)$ is analytically intractable, as is the case for the WS and WGS, finding a closed form solution to the posterior expectation is prohibitively difficult if not impossible. Thus posterior inference for θ requires MCMC techniques. Define Markov chain $\Phi = \left\{\theta^{(0)}, \theta^{(1)}, \ldots, \theta^{(N)}\right\}$ for $f_{\text{post}}(\theta|t)$. Monte Carlo averages calculated from the Markov chain output provide estimates of the posterior expectations:

$$\overline{\theta}_N := \frac{1}{N} \sum_{i=1}^{N} \theta^{(i)}.$$

Moreover, the strong law of large numbers guarantees that with probability 1, $\overline{\theta}_N \to E[\theta|t]$ as $N \to \infty$.

For each combination of portfolio setting (prime and subprime) and target model (WS and WGS) we construct a Markov chain Φ of length $N = 10^5$ using the `rjags` ("Just Another Gibbs Sampler") package for R (Plummer, 2016). Given the target model, `rjags` constructs Φ using a combination of Gibbs sampling, Metropolis-Hastings, adaptive rejection, and slice sampling techniques. For details on `rjags` methods see Coro (2017). Note that outside 'rjags' we also constructed alternative Markov chains using a Metropolis-Hastings random walk strategy. In light of the multivariate structure of the models, this algorithm proved difficult to tune and resulted in slow mixing Markov chains. Further, the results were similar to those provided by `rjags` thus are eliminated here.

4.2 WGS Results

Table 1 summarizes the Monte Carlo average estimates of WGS parameters (c_1, c_2, p, r, α) for the prime and subprime portfolios. A discussion is deferred to Section 4.4.

Table 1: Monte Carlo average parameter estimates (with naive standard errors) for the WGS model of time-to-default for subprime and prime portfolios.

	Subprime Portfolio		Prime Portfolio	
Parameter	Estimate	Naive standard error	Estimate	Naive standard error
c_1	2.239	$1.3 * 10^{-4}$	2.434	$2.3 * 10^{-4}$
c_2	0.340	$1.3 * 10^{-4}$	0.380	$1.4 * 10^{-4}$
p	0.201	$1.3 * 10^{-5}$	0.038	$6.2 * 10^{-6}$
r	443.1	3.9	3241.6	21.6
α	$2.481 * 10^{6}$	$2.1 * 10^{4}$	$1.210 * 10^{7}$	$7.6 * 10^{4}$

The overall fit of the WGS is satisfactory, reasonably capturing the tail behavior of the mortgage defaults. However, MCMC diagnostics provides insight into certaing problems of this model. Trace plots of the Markov chain output from which the prime portfolio parameter estimates are calculated are shown in Figure 2. Markov chain patterns in the subprime setting are similar. The trace plots for c_1 and p exhibit decent mixing behavior whereas the Markov chain for c_2 mixes slowly. Similar behavior was exhibited across multiple simulations, thus addressing this issue would require either a different model or more sophisticated Markov chain algorithm. Finally, consider the trace plots for α and r, the shape and scale parameters for the Gamma distribution of variability in mortgage risk. These indicate instability, slow mixing, and dependence across the Markov chain values of α and r. Mainly, the ratio r/α remains relatively constant. Relative to the prior expected value and variance of λ, $E(\lambda|r,\alpha) = r/\alpha \simeq$ constant and $Var(\lambda|r,\alpha) = r/\alpha^2 \simeq 0$. This Markov chain behavior reveals potential over-parametrization of the WGS, in turn suggesting the use of the simpler WS which assumes constant risk factor λ.

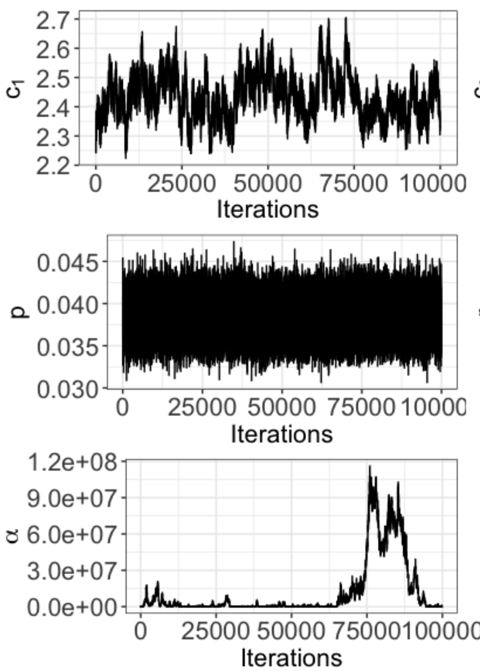

Figure 2: Markov chain trace plots corresponding to the WGS of prime portfolios.

4.3 WS Results

Table 2 summarizes the Monte Carlo average estimates of WS parameters (λ, c_1, c_2, p) for the prime and subprime portfolios. Figure 3 illustrates the WS at these parameter settings. These results are discussed and compared to those for the WGS in Section 4.4.

Table 2: Monte Carlo average parameter estimates (with naive standard errors) for the WS model of time-to-default for subprime and prime portfolios.

	Subprime Portfolio		Prime Portfolio	
Parameter	Estimate	Naive standard error	Estimate	Naive standard error
λ	0.00019	$9.8 * 10^{-8}$	0.00022	$6.5 * 10^{-5}$
c_1	2.234	$1.3 * 10^{-4}$	2.457	$2.6 * 10^{-4}$
c_2	0.272	$1.2 * 10^{-4}$	0.398	$2.4 * 10^{-4}$
p	0.201	$1.3 * 10^{-5}$	0.038	$6.2 * 10^{-6}$

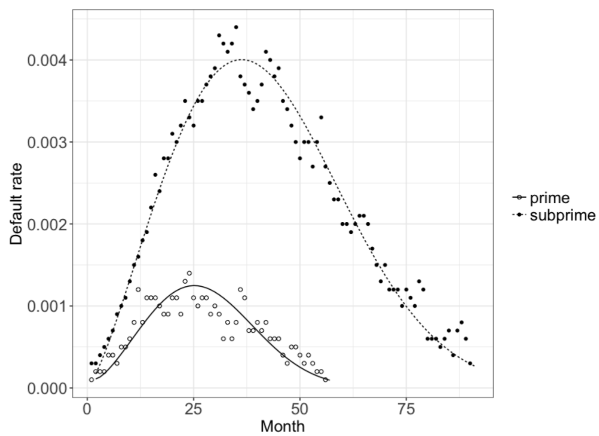

Figure 3: Estimated WS models for subprime and prime portfolios defaults superimposed on the raw data.

Finally, Figure 4 illustrates Markov chain trace plots corresponding to the WGS of prime portfolios. As with the WGS Markov chains, we observe decent mixing for (λ, c_1, p) yet slow mixing for c_2.

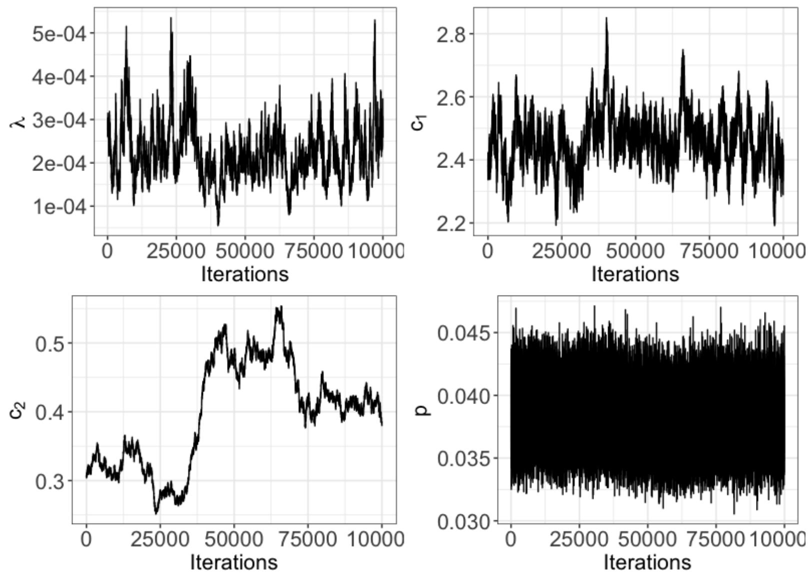

Figure 4: Markov chain trace plots corresponding to the WS of prime portfolios.

4.4 Discussion

The results above support the modeling of subprime and prime time-to-default via the Bayesian Weibull segmentation model (WS). In contrast, the over-parametrization of the WGS suggested by the MCMC simulation is an unexpected result of our study. Mainly, allowing for heterogeneity in default rates λ across mortgage holders does not appear to be necessary.

The parameter estimates themselves also offer unique insight into time-to-default patterns. Comparing Tables 1 and 2 we observe that c_1 and c_2, the parameters capturing temporal changes in hazard rates for the two latent segments of portfolio holders, are quite similar for the WS and WGS. Among both prime and subprime portfolios, the smaller segment ($p < 0.5$) is characterized by an increasing hazard rate ($c_1 > 1$). Increasing rates might be explained by additional stress on households that have signed a mortgage contract they cannot afford. In contrast, the larger segment ($(1-p) > 0.5$) is characterized by a decreasing hazard rate ($c_2 < 1$). This phenomenon may correspond to mortgage holders' growing equity and increasing financial losses in case of default. Both the WS and WGS suggest that this segment is also practically default-free, though it is hard to qualify this statement due to the latent nature of the segmentation.

Our results also suggest a key difference between prime and subprime portfolios. In the subprime portfolios, the risky segment with increasing hazard rates ($c_1 > 1$) constitutes roughly twenty percent of mortgage holders ($p \approx 0.2$). A similar 20-80 segmentation of mortgage holders into risky/non-risky classes was detected by Soyer and Xu (2010) in EPD (early default) data. Among prime portfolios, the risky segment constitutes less than five percent of mortgage holders ($p < 0.05$). A minor difference in the WS default rate parameters for prime portfolios ($\lambda \approx 0.00022$) and subprime portfolios ($\lambda \approx 0.00019$)is hardly significant. Thus our results support the suggestion that both prime and subprime portfolios may contain risky segments characterized by increasing default rates.

5 Acknowledgements

The authors wish to thank Ellen Klingner for her research assistance and also NSF CSUMS grant DMS0802959 which made this work possible. We are also very grateful to the referees for their many valuable comments.

References

Aktekin, T., Soyer, R., and Xu, F. (2013). Assessment of mortgage default risk via bayesian state space models. *Annals of Applied Statistics*, 7,3:1450–1473.

Chen, W. C., Hill, B. M., Greenhouse, J. B., and Fayes, J. V. (1985). Bayesian analysis of survival curves for cancer patients following treatment. *Bayesian Statistics*, 2:299–328.

Coro, G. (2017). *Gibbs Sampling with JAGS: Behind the Scenes.*

Das, A. and Stein, R. M. (2009). Underwriting versus economy: a new approach to decomposing mortgage losses. *Journal of Credit Risk*, 15(2):19–41.

Demyanyk, Y. and Van Hemert, O. (2011). Understanding the subprime mortgage crisis. *Review of Financial Studies*, 24(6):1848–1880.

Erisoglu, U., Erisoglu, M., and Erol, H. (2011). A mixture model of two different distributions approach to the analysis of heterogeneous survival data. *International Journal of Computational and Mathematical Sciences*, 5(2):75–79.

Fader, P. S. and Hardie, B. G. S. (2007). How to project customer retention. *Journal of Interactive Marketing*, 21:76–90.

Fader, P. S., Hardie, B. G. S., and Zeithammer, R. (2003). Forecasting new product trial in a controlled test market environment. *Journal of Forecasting*, 22(5):391–410.

Glennon, D. and Nigro, P. (2011). Evaluating the performance of static versus dynamic models of credit default: evidence from long-term Small Business Administration-guaranteed loans. *Journal of Credit Risk*, 7(2):3–35.

Hardie, B. G. S., Fader, P. S., and Wisniewski, M. (1998). An empirical comparison of new product trial forecasting models. *Journal of Forecasting*, 17:209–229.

Kiefer, N. M. (2006). The probability approach to default probabilities. Technical report, Cornell University.

Kiefer, N. M. (2008). Default estimation, correlated defaults, and expert information. Technical report. CAE Working paper #08-02, ISSN 1936-5098.

Marin, J. M., Rodriguez-Bernal, M. T., and Wiper, M. P. (2005). Using Weibull mixture distributions to model heterogeneous survival data. *Communications in Statistics : Simulation and Computation*, 34:673–684.

Plummer, M. (2016). *rjags: Bayesian Graphical Models using MCMC*. R package version 4-6.

Popova, I., Popova, E., and George, E. I. (2008). Assessment of mortgage default risk via bayesian reliability models. *Bayesian Analysis*, 3:393–426.

Soyer, R. and Xu, F. (2010). Assessment of mortgage default risk via bayesian reliability models. *Applied Stochastic Models Bus. Ind.*, 26:308–330.