The syllabus for this four-hour exam is defined in the form of learning objectives, knowledge statements, and readings.

LEARNING OBJECTIVES set forth, usually in broad terms, what the candidate should be able to do in actual practice. Included in these learning objectives are certain methodologies that may not be possible to perform on an examination, such as applying the Hamiltonian Monte Carlo algorithm to solve for the parameters of a model or building an extended decision tree model with boosting or bagging, but we may ask the candidate to calculate a limited number of steps contained in those algorithms. We also may ask the candidate to respond to questions on the underlying concepts of the algorithms in the learning objectives and knowledge statements set forth below.

KNOWLEDGE STATEMENTS identify some of the key terms, concepts, and methods that are associated with each learning objective. These knowledge statements are not intended to represent an exhaustive list of topics that may be tested, but they are illustrative of the scope of each learning objective.

READINGS support the learning objectives. It is intended that the readings, in conjunction with the material on earlier examinations, provide sufficient resources to allow the candidate to perform the learning objectives. Some readings are cited for more than one learning objective. The CAS Syllabus & Examination Committee emphasizes that candidates are expected to use the readings cited in this *Syllabus* as their primary study materials.

Thus, the learning objectives, knowledge statements, and readings complement each other. The learning objectives define the behaviors, the knowledge statements illustrate more fully the intended scope of the learning objectives, and the readings provide the source material to achieve the learning objectives. Learning objectives should not be seen as independent units, but as building blocks for the understanding and integration of important competencies that the candidate will be able to demonstrate.

Note that the range of weights shown should be viewed as a guideline only. There is no intent that they be strictly adhered to on any given examination—the actual weight may fall outside the published range on any particular examination.

The overall section weights should be viewed as having more significance than the weights for the individual learning objectives. Over a number of years of examinations, absent changes, it is likely that the average of the weights for each individual overall section will be in the vicinity of the guideline weights. For the weights of individual learning objectives, such convergence is less likely. On a given examination, in which it is very possible that not every individual learning objective will be tested, there will be more divergence of guideline weights and actual weights. Questions on a given learning objective may be drawn from any of the listed readings, or a combination of the readings. There may be no questions from one or more readings on a particular exam.

After each set of learning objectives, the readings are listed in abbreviated form. Complete text references are provided at the end of this exam syllabus.

Items marked with a bold **OP** (Online Publication) are available at no charge and may be downloaded from the CAS website.

Please check the "*Syllabus* Update" section of the CAS Web Site for any changes to the *Syllabus*.

Thorough knowledge of calculus and probability is assumed. Given the material covered on this exam, we assume that the candidate has knowledge of linear algebra concepts at the level commonly assumed as a prerequisite to taking an undergraduate level course in regression analysis. Candidates are expected to have mastered the concepts in Exam MAS-I. For

those candidates who have obtained a waiver for Exam MAS-I through the transition rule that granted credit for Exam MAS-I by having credit for Exam S - Statistics and Probabilistic Models or through examinations administered by the Institute and Faculty of Actuaries (United Kingdom), Actuaries Institute (Australia), Actuarial Society of South Africa (ASSA), or the Institute of Actuaries of India, it is recommended to review and master the concepts in the paper "Generalized Linear Models" by Larsen[1] and the following Sections in *An Introduction to Statistical Learning, with Applications in R*: 2.1.4, 2.2.1, 2.2.2, 5.1, and 5.2. See Waivers of Examination page of the CAS website for a complete waiver explanation. While some problems may have an insurance or risk management theme, no prior knowledge of insurance terminology is expected.

A variety of tables along with standard notation for the mixed models will be provided to the candidate with the examination's reference materials. The tables include values for the standard normal distribution, Chi-square distribution, *t*-distribution, and F-distribution. Abridged inventories of discrete and continuous probability distributions will also be provided. Since they will be included with the examination, candidates will not be allowed to bring copies of the tables into the examination room.

---

[1] See Syllabus for Exam MAS-I for complete text reference.

A. Introduction to Credibility

Range of weight for Section A: 5-15 percent

Advances in statistical computing tools have made it practical to include a form of credibility weighting when building regression-type models. For example, what the statisticians call shrinkage in a Linear Mixed Effect Model is a form of least squares credibility weighting. These advanced techniques are covered extensively in Sections B and C. Also, one can view the penalized regression techniques covered in Exam MAS-I as another form of credibility weighting. Candidates should be familiar with the topics listed below as they can serve as a good introduction to those techniques and are still very much in practice today. Specifically, candidates should be familiar with limited fluctuation credibility and be able to calculate estimates using Bayesian credibility procedures. They should also be fluent with Bayesian and Bühlmann (least squares credibility) procedures both for discrete and continuous models.

LEARNING OBJECTIVES	KNOWLEDGE STATEMENTS
1. Understand the basic framework of credibility and be familiar with limited fluctuation credibility, including partial and full credibility	a. Limited fluctuation credibility, partial and full credibility
2. Understand the basic framework of Bühlmann credibility	b. Conjugate priors - Poisson/Gamma, Binomial/Beta
3. Calculate different variance components for Bühlmann credibility	c. Bühlmann credibility continuous
4. Calculate Bühlmann and Bühlmann-Straub credibility factor and estimates for frequency, severity, and aggregate loss	d. Bühlmann credibility discrete
5. Understand the basic framework of Bayesian credibility	e. Bayesian analysis discrete
6. Calculate Bayes estimate/Bayesian premium	f. Bayesian analysis continuous
7. Bayesian versus Bühlmann credibility for conjugate distributions	g. Nonparametric empirical Bayes
8. Calculate credibility estimates using the nonparametric empirical Bayes method	
Range of weight for Learning Objectives A.1 through A.8 collectively: 5-15 percent	

READINGS
• Tse, Chapters 6.1-6.3, 7.1-7.4, 8.1-8.2, and 9.1-9.2

B. Linear Mixed Models

Range of weight for Section B: 10-30 percent

This section covers linear models that use a form of credibility weighting for a designated subset of variables in the model called random effects. The candidates will be expected to understand the concepts of shrinkage for linear mixed models as well as how to accommodate models with correlated observations or models where the variance is not assumed to be constant for each observation or a function of the mean. Mixed models include both Gaussian (linear mixed models) and non-Gaussian (generalized linear mixed models or non-linear mixed models) models. This section will only cover linear mixed models. The candidate is expected to understand the linkage between shrinkage and credibility weighting, how to select the appropriate model to induce credibility weighting at the appropriate level when setting up the model structure, and how to account for correlation in the residuals.

LEARNING OBJECTIVES	KNOWLEDGE STATEMENTS
1. Understand the assumptions behind linear mixed Models and use that understanding to evaluate how to set up a linear mixed effect model design to best accomplish the goals of the modeling exercise	a. Characteristics of random and fixed effects explanatory variables for linear mixed models • How to identify a random effect variable • Interaction between fixed and random effect variables when calculating standard error of estimate b. Implications of correlation matrix choice by model form for linear mixed models • Independence assumption for observations • Repeated measures/longitudinal studies • Correlation forms for random vs. fixed effect variables • Hierarchal model structure implementation of treatment and design structure • Explicitly model variance as a function of an explanatory variable

LEARNING OBJECTIVES	KNOWLEDGE STATEMENTS
2. Understand the algorithms behind the numerical solutions for the linear mixed model to enable interpretation of output from the statistical software employed in modeling to make appropriate choices when evaluating modeling options	a. Restricted maximum likelihood b. The choice between restricted maximum likelihood and maximum likelihood c. Estimable vs. predictable functions d. Best linear unbiased predictor and interaction with fixed effects variables e. Shrinkage of best linear unbiased predictors f. Newton-Raphson vs. Fisher scoring vs. EM algorithm g. Bias in variance estimates of fixed effects h. Credibility adjusted degrees of freedom (Saitherwaite/Kenward Rodgers adjustments) i. Conditional vs. population estimate
3. Understand and be able to select the appropriate model structure and variable selection for a linear mixed model given the behavior of the data set to be modeled by interpreting the model diagnostics and or summary statistics on the variables available in the model along with any graphs depicting how the dependent variable behaves as a function of possible explanatory variables Range of weight for Learning Objectives B.1 through B.3 collectively: 10-30 percent	a. Units of replication b. Randomized block designs c. The implication of random effects for model prediction d. Interaction terms for fixed effects vs. random effect variables e. Model selection when covariance structure changes f. Covariance structure g. Selection of fixed vs. random effect class for mixed effect explanatory variables h. Explicitly model variance i. Marginal Model and Implied Marginal Model j. Residual graphs evaluating normality and constant variable assumptions k. Hypothesis tests for fixed and random effects l. Intraclass correlation coefficient m. Know when nested model comparisons are appropriate n. Application of AIC & BIC relative measures of goodness of fit o. Application of Scatter Plots and Box Plots as an aid to model design

LEARNING OBJECTIVES	KNOWLEDGE STATEMENTS

READINGS

- West, Chapter 1, 2 (excluding 2.9.6), 3 (excluding 3.4.1, 3.4.2, 3.4.4, 3.4.5, 3.6, 3.11), 4 (excluding 4.4.1, 4.4.2, 4.4.4, 4.4.5, 4.6, 4.11), 5 (excluding 5.4.1, 5.4.2, 5.4.4, 5.4.5, 5.6, 5.10), 6 (excluding 6.4.1, 6.4.2, 6.4.4, 6.4.5, 6.6, 6.10), 7 (excluding 7.4.1, 7.4.2, 7.4.4, 7.4.5, 7.6, 7.10), 8 (excluding 8.4.1, 8.4.2, 8.4.4, 8.4.5, 8.6, 8.10), and Appendix B, Additional Notes on Shrinkage (http://www-personal.umich.edu/~bwest/shrinkage.doc)

Chapters 1 and 2 contain an introduction to the modeling concepts underlying linear mixed effect models. Chapters 3 through 8 contain examples that illustrate how to build a linear mixed effect model to accommodate different circumstances. Adapting the general formulas for a Mixed Model to accommodate the specific nature of the problem at hand is a skill that the candidate should master. The excluded sections from Chapters 3-8 go into details on software that are outside of the scope of the syllabus. We include sections that demonstrate coding the models in R (*.4.3 in Chapters 3 through 8) since we will have case studies on the exam that use R to generate the modeling results. However, we do not expect candidates to master the details of coding linear mixed effect models in R. Candidates should focus on understanding the design choices made in modeling, the output from those packages, and how that output was interpreted rather than on the details of coding for the purpose of this exam. Comments in Chapters 3-8 on design choices and the type of hypothesis test to be employed at a given point in the modeling process expand on the introduction to modeling concepts covered in Chapters 1 and 2. These comments are a vital part of the reading from West.

The matrix notations employed in the readings for specifying linear model forms will be adopted for exam questions from this section and will be provided in the Exam MAS-II Tables supplementary packet that accompanies the examination's reference materials. Additionally, the format for model output will match that as provided by R software.

Similarly, exam questions from this section may contain parameter tables and diagnostic tables or plots of the type shown in the text. Candidates should understand how to interpret these tables. Candidates who become familiar with a statistical language capable of generating this type of output, such as R, will have an easier time understanding and applying the concepts covered in the syllabus material. In particular, candidates that work with the R code examples in the West textbook, along with the datasets provided, will have a better grasp of the material than that obtained by simply reading the textbook. However, for exam questions from this Section, candidates will not be explicitly tested on software code.

The book's website has made available the datasets and code introduced in the chapters. It can be found at http://www-personal.umich.edu/~bwest/almmussp.html.

C. Bayesian Analysis and Markov Chain Monte Carlo

Range of weight for Section C: 45-65 percent

This section introduces Bayesian Markov chain Monte Carlo (MCMC) methods and illustrates how to incorporate the following in a Bayesian regression model:

- Credibility considerations at the same time one is solving for the model parameters based on the behavior of the observations within the data set to temper the model parameters (partial pooling)
- Prior subject matter knowledge to supplement the information in a limited data set through a form of credibility weighting (regularizing priors)
- Hierarchical model structures (layers of data with varying credibility)
- Complex covariance structures that recognize the correlation between observations and/or variance results that are not a simple function of the mean

Note that the terminology employed by the text references does not always match actuarial terminology. Terms like regularizing priors or partial pooling, which are employed to statistically temper parameter estimates, would be called credibility weighting in much of the actuarial literature.

The candidate is expected to be able to apply Bayesian techniques through the MCMC algorithms, to understand a model, and to evaluate the resulting goodness of fit by interpreting the diagnostics that are described below. The candidate is also expected to understand why the Bayesian approach is different than the classical (frequentist) procedures.

LEARNING OBJECTIVES	KNOWLEDGE STATEMENTS
1. Understand the basis and basics of Bayesian analysis and incorporate that understanding when interpreting model results • Difference between Bayesian and classical (frequentist) procedures • How probability is used as a measure of uncertainty • Components of a Bayesian model • Use of simulation to create predictive distributions • Summarizing posterior distributions	a. Bayes' rule b. Subjectivity c. Likelihood d. Prior distribution e. Posterior distribution f. Posterior predictive distribution g. Credible, confidence, highest posterior density, and prediction intervals h. Subject matter knowledge in setting the prior distribution parameters
2. Evaluate the different options available when creating and using Bayesian models for a given modeling assignment. Understand how to set up a Bayesian MCMC model and evaluate how a given set of design choices affects the results of a model • Recognize benefits and limitations of different kinds and parameterizations of priors • Calculating posterior and posterior predictive distributions for single and multi-parameter models	a. Conjugacy b. Proper and improper priors c. Informative, non-informative, and regularizing priors d. Hyperpriors e. Exchangeability f. Transformations of parameters g. Sampling from the posterior distribution h. Regularization

LEARNING OBJECTIVES	KNOWLEDGE STATEMENTS
- Linear Regression - Additive models - Mixture models - Hierarchical models - Model covariance structures	i. Pooling and shrinkage j. Smoothing functions k. Varying effects of intercepts and slopes l. Gaussian Process Regression for continuous varying effects m. Model correlation effects through covariance structure n. Model non-constant variance through covariance structure o. Generating predictions beyond the model's training data
3. Understand Bayesian computation, how Markov Chain Monte Carlo methods are used, and how to evaluate model performance. Interpret and calculate diagnostics of simulation performance to evaluate when a given modeling approach should be used. - Simulation and sampling - Conditional sampling - Convergence assessment - Efficient samplers - Hamiltonian Monte Carlo	a. Shortcomings of grid and quadratic approximation b. Markov chains c. Gibbs sampler d. Metropolis and Metropolis-Hastings algorithms e. Warm-up / Burn-in f. Convergence in parameter estimate measurement g. Trace plot h. Trank plots i. Acceptance rate j. Within-sequence correlation k. Thinning l. Effective number of samples m. Potential scale reduction (R-Hat or Gelman-Rubin statistic) n. Advantage of Hamiltonian Monte Carlo o. Divergent transitions

LEARNING OBJECTIVES	KNOWLEDGE STATEMENTS
4. Understand how to apply model checking, evaluation, comparison, and expansion techniques as an aid to interpreting and evaluating model diagnostics • Know how to check model fit to data • Understand limitations of various tests • Understand and calculate measures of model predictive accuracy • Understand and calculate information criteria, their uses, and limitations (e.g., bias, dependence on prior, etc.) • Compare models via predictive performance measures • Understand how models can be expanded and what further checks may be needed Range of weight for Learning Objectives C.1 through C.4 collectively: 45-65 percent	a. Sensitivity analysis b. External validation c. Posterior predictive checking d. Predictor residual plots e. Counterfactual plots f. Marginal vs. population average posterior predictions g. Log predictive density h. Out-of-sample predictive accuracy i. Information criteria measures (AIC, DIC, WAIC) j. Effective number of parameters k. Cross-validation (LOO-CV) l. Pareto-Smoothed Importance Sampling Cross-Validation (PSIS-CV) m. Model averaging n. Simulate prior distribution

READINGS

• Ford
• McElreath, Chapters 1, 2 (excluding 2.1), 3, 4, 5, 7, 8.1, 9, 11, 12, 13, 14 (excluding 14.3, 14.4, and 14.5.2)

While candidates will not be explicitly tested on software code, they are encouraged to work the exercises in the material to get hands-on, practical experience building and interpreting Bayesian models.

For instruction on how to obtain the required software and configure the computing environment to run the examples in McElreath, candidates can read the 'Front Matter' section of the text and visit the author's website at http://xcelab.net/rm/software/.

For a copy of the R programs in the Ford study note, candidates can find them under mcmc_algorithms.R at https://github.com/pford221/mcmc_algorithms/.

D. Statistical Learning

Range of weight for Section D: 10-20 percent

This section introduces candidates to a sample of foundational statistical learning techniques. Both supervised and unsupervised techniques are detailed in the readings, and candidates should be able to distinguish between them. The supervised learning techniques are nonparametric in nature, meaning the model cannot easily be described in an equation form — they tend to model flexible, non-linear hypotheses well. The unsupervised learning techniques are useful for reducing the dimensions of the data to aid in the profiling of the data or to facilitate more efficient learning from the data. Candidates are expected to understand the mechanics of these algorithms and recognize their inherent strengths and weaknesses so as to be able to select the most appropriate procedure for the learning task at hand.

LEARNING OBJECTIVES	KNOWLEDGE STATEMENTS
1. Understand the computations behind K-nearest neighbors (KNN) and be able to explain how it works in practice and its relationship with Bayes classifier.	a. Classification versus regression for supervised learning b. Bayes classifier c. KNN decision boundary versus Bayes decision boundary
2. Understand the computations involved in building decision trees, the purpose of tree pruning, and how extensions such as bagging, random forest, and boosting can improve the prediction accuracy of tree-based methods.	a. Recursive binary splitting for decision trees b. Pruning for decision trees c. Comparison of decision trees versus linear models d. Advantages and disadvantages of decision trees e. Bagging and OOB Error f. Similarity and differences between bagging and random forest g. Sequential learning via boosting h. Gini/entropy application for splitting
3. Understand the purpose of, and the computations behind, principle components analysis (PCA) and be able to interpret related software outputs.	a. Loading vector and scores for principle component b. Effect of scaling on PCA c. Proportion of variance explained by PCA and scree plots d. Combining many dimensions (variables) into fewer e. Compare purpose of PCA to K-means

LEARNING OBJECTIVES	KNOWLEDGE STATEMENTS
4. Be familiar with the purpose of, and the computations behind, clustering procedures and be able to interpret related software outputs. Range of weight for Learning Objectives D.1 through D.4 collectively: 10-20 percent	a. K–means clustering algorithm b. Agglomerative hierarchical clustering algorithm c. Dendrogram d. Dissimilarity measure

READINGS

- James et al., Chapters 1 (Background reading only), 2.2.3, 8, and 10. Exam questions will not be sourced directly from Chapter 1.

Programming labs found in sections 2.3, 8.3, and 10.4 demonstrate the implementation of topics covered in the sections above with R, but no new concepts are introduced here. These labs show how to load and work with datasets made available in the ISLR package in R. For information on how to install and load R packages, please refer to section 3.6.1. Examination questions will not explicitly test software code, but a careful review of these sections will greatly help the candidates understand and apply these concepts.

Complete Text References for Exam MAS-II

Text references are alphabetized by the citation column.

Citation	Abbreviation	Learning Objective	Source
Ford, P., "MCMC Algorithms," CAS Study Note, Version 0.8, November 2019.	Ford	C3	**OP** **NEW**
James, G., et al., *An Introduction to Statistical Learning, with Application in R*, 1st ed. 2013, Corr. 8th printing, Springer, 2017. Notes: • Although page iv of the text identifies it as the corrected 8th printing, the publisher refers to it as the corrected 7th printing. Candidates should use the version found in the Complete Online Text References web page for this exam. • Although the 2nd edition of the text is available, the Spring 2023 examinations will be based on the first edition.	James et al.	D1-D4	**OP**
McElreath, R., *Statistical Rethinking: A Bayesian Course with Examples in R and Stan*, 2nd edition, CRC Press, March 2020.	McElreath	C1-C4	**B**
Tse, Y., *Nonlife Actuarial Models, Theory Methods and Evaluation*, Cambridge University Press, 2009.	Tse	A1-A8	**B**
West, B. T.; Welsh, K. B.; and Galecki, A. T., *Linear Mixed Models: A Practical Guide Using Statistical Software*, 2nd Edition, CRC Press, 2015. Candidates can study from either the 2nd or 3rd edition but should note that the 3rd edition may contain additional material that is not on the syllabus.	West	B1-B3	**B**

Source Key

B	Book—may be purchased from the publisher or bookstore.
DSK	Material included in the 2023 Digital Study Kit.
NEW	Indicates new or updated material.
OP	All text references marked as Online Publications will be available on a web page titled Complete Online Text References.
SK	Material included in the 2023 Study Kit.
SKU	Material included in both the 2023 CAS Study Kit and the 2023 Update to the 2022 Study Kit.

Items printed in **red** indicate an update, clarification, or change.