



VNIVERSITAT DE VALÈNCIA

**Bayesian hierarchical modelling
for repeated measures and
longitudinal data**

by

Gabriel Calvo Bayarri

Programa de Doctorado en Estadística y Optimización

Supervised by

Carmen Armero Cervera

Luigi Spezia

at the

Universitat de València

November 2023

This thesis was supported by the ONCE Foundation, the Universia Foundation, and the Spanish Ministry of Education and Professional Training, grant FPU18/03101.

Acknowledgements

Gracias Carmen por introducirme en este nuevo mundo que desconocía y que poco a poco me ha ido encandilando: la investigación.

Gracias a tí y a Luigi por acompañarme en esta aventura apasionante.

Gracias a la Familia y a los Hermanos por mostrarme y conducirme hacia lo más importante: el Amor.

Gracias a los Amigos por soportarme, y aguantar mi estrés, y mis preocupaciones, y por estar cerca de mí.

Como dijo G.K Chesterton: *“El bien no está en encontrarnos con nuestros amigos, sino en haberlos encontrado.”* El verdadero Amigo es el que siempre permanece fiel. Gracias.

*Para venir a lo que no sabes,
has de ir por donde no sabes.*

Monte de Perfección. *San Juan de la Cruz*

Bayesian hierarchical modelling for repeated measures and longitudinal data

Gabriel Calvo Bayarri

Resumen

Introducción

El presente trabajo se apoya en tres pilares fundamentales: la estadística bayesiana, los datos longitudinales y de medidas repetidas, y la modelización. Estos tres elementos conforman el paraguas teórico de esta memoria. A continuación señalamos sus aspectos más importantes.

En primer lugar, el teorema de Bayes es esencial en la estadística bayesiana, que busca aprender sobre una población desconocida a partir de una muestra. La inferencia bayesiana utiliza el teorema de Bayes para actualizar la información sobre una cantidad de interés θ basada en los datos $\mathcal{D}$, representada por la ecuación $\pi(\theta | \mathcal{D}) \propto f(\mathcal{D} | \theta)\pi(\theta)$, donde $\pi(\theta | \mathcal{D})$ es la distribución a posteriori, $f(\mathcal{D} | \theta)$ es la función de verosimilitud, y $\pi(\theta)$ es la distribución previa.

En la estadística bayesiana, θ se modela con una distribución de probabilidad, brindando flexibilidad para representar parámetros, hipótesis, modelos, entre otros. Aunque la distribución posterior rara vez tiene una expresión analítica, se han desarrollado numerosas propuestas computacionales para simular o aproximar esta distribución, siendo los métodos de Monte Carlo de Cadenas de Markov (MCMC) de los más populares. Estos métodos generan una cadena de Markov cuya distribución límite es la distribución a posteriori buscada, facilitando la aproximación de la distribución a posteriori sin requerir una solución analítica explícita. Por esta razón, hemos trabajado o bien con un software específico llamado JAGS (Plummer, 2003), o hemos implementado los algoritmos de MCMC en R.

Por otro lado, la complejidad inherente a los datos de medidas repetidas y longitudinales reside en que suelen ser multivariantes y nos brindan información no solo de la población en general, sino también de cada individuo en particular. Además, este tipo de datos aparece con frecuencia en diversos campos del conocimiento, como la medicina, la salud pública, la educación, la economía, la psicología, la biología, la pesca y el deporte, entre otros. Por lo que pueden desempeñar un papel crucial en diferentes ramas del saber. Para conseguir una visualización clara de este tipo de datos, es habitual el uso de líneas para conectar observaciones consecutivas dentro de la unidad experimental. Este tipo de gráficas reciben el nombre de *spaghetti plot* o *profile plot*. Por supuesto, hemos utilizado ampliamente este tipo de representación de datos a lo largo del documento.

En nuestra investigación, abarcamos una amplia variedad de datos que provienen de diferentes campos, cada uno explorado en un capítulo distinto. El Capítulo 2 está dedicado a una investigación médica que involucra a 198 pacientes. Este capítulo investiga la intrincada conexión entre la presión intra-abdominal durante la cirugía lapa-

roscópica y el volumen intra-abdominal resultante. Mientras tanto, el Capítulo 3 cambia su enfoque hacia un análisis pesquero, específicamente en la industria pesquera de la sardina Europea en el mar Mediterráneo, desglosando los datos por país. Además, nuestra exploración se adentra en el ámbito del análisis deportivo en los Capítulos 5 y 6. Estos capítulos profundizan en el ámbito del baloncesto profesional y el baloncesto en silla de ruedas, respectivamente. El ámbito deportivo ofrece una fuente abundante de datos longitudinales, repleto de potencial inexplorado para un análisis longitudinal preciso.

Por último, esta memoria tiene como objetivo principal profundizar en el desarrollo metodológico de modelos para datos de medidas repetidas y datos longitudinales en el marco de la modelización jerárquica bayesiana. Estos modelos consideran la correlación dentro de cada unidad experimental y datos independientes entre diferentes unidades experimentales. Los modelos estudiados contienen básicamente cuatro elementos: la variable respuesta y , un posible proceso oculto o latente Z , un conjunto de parámetro e hiperparámetros θ y un conjunto de los efectos aleatorios y términos de correlación en serie ϕ . Se pueden estimar patrones de comportamiento general de la población, así como características específicas de individuos o grupos en estudio, teniendo en cuenta la heterogeneidad, errores de medición y correlación serial en los datos. Se asume, además, independencia condicional entre individuos, permitiendo expresar la distribución conjunta como un producto de distribuciones condicionales individuales. Este marco se aplica en el Capítulo 5 para investigar el fenómeno de la *hot hand* en el baloncesto. Además, se presenta una formulación simplificada sin procesos ocultos en (1.4), utilizada predominantemente en este trabajo, considerando una o dos variables de respuesta en los diferentes capítulos.

La investigación viene principalmente motivada por los numerosos desafíos del mundo real que generan los estudios con datos longitudinales y medidas repetidas. Este trabajo produce tanto propuestas e ideas aplicadas como avances metodológicos.

Estructura de la memoria

- **Capítulo 1. Introducción y marco teórico.** Este capítulo introduce el contexto teórico del presente trabajo y hace un resumen de los contenidos, propuestas metodológicas y avances que se abordan en el mismo.
- **Capítulo 2. Un modelo no lineal jerárquico bayesiano para medidas repetidas.** Este primer estudio presenta un modelo estadístico bayesiano jerárquico no lineal para representar la relación entre el volumen intraabdominal (IAV) y la presión intraabdominal (IAP) durante la cirugía laparoscópica.

La laparoscopia es un procedimiento quirúrgico que se realiza en el abdomen mediante pequeñas incisiones para visualizar órganos con una cámara o facilitar una intervención quirúrgica. Esta técnica se asocia con menor morbilidad que la cirugía abierta tradicional. Sin embargo, la presión abdominal elevada también puede provocar efectos adversos. Identificar el punto crítico en el que debe limitarse la insuflación es crucial para maximizar el espacio de trabajo quirúrgico y minimizar los efectos lesivos.

Estudios previos plantean que la relación entre IAP e IAV durante la insuflación neumoperitoneal sigue un patrón sigmoidal creciente. El objetivo de este trabajo es estimar un modelo que represente esta relación, incorporando características de los pacientes, para posteriormente realizar inferencias individuales y poblacionales sobre sus parámetros y puntos críticos asociados.

Se plantea un modelo no lineal de efectos mixtos, con una función logística de crecimiento. Esta función presenta una asíntota horizontal y puntos de inflexión. Sus parámetros se relacionan con covariables edad y sexo a través de un modelo jerárquico, el cual también incluye efectos aleatorios individuales. El análisis bayesiano permite estimar la distribución a posteriori conjunta de parámetros y efectos aleatorios. De ella se derivan distribuciones marginales a posteriori para realizar inferencia. También se pueden generar predicciones para nuevos pacientes.

Se analizaron datos de 198 pacientes, con 1361 observaciones IAV-IAP. La edad osciló entre 23 y 92 años, con mayor número de hombres. El modelo revela un comportamiento diferenciado en la relación IAV-IAP de hombres y mujeres, así como cambios leves con la edad. Se obtuvieron distribuciones a posteriori para los puntos críticos asociados a la curva logística, permitiendo predicción individual y poblacional. Por ejemplo, en el punto de desaceleración asintótica (ADP) la presión IAP estimada fue 10.06 mmHg (hombres) y 10.87 mmHg (mujeres) para la edad media muestral. Se realizó validación del modelo, detectando problemas de calibración para valores medios y medio-altos de IAP, donde se sobreestiman los valores de IAV.

Este modelo preliminar con una muestra pequeña de pacientes constituye el primer paso hacia una estrategia de medicina de precisión para protocolizar la insuflación neumoperitoneal en la cirugía laparoscópica. Sin embargo, se requieren más datos para representar adecuadamente la variabilidad poblacional. Futuros estudios deberán expandir el tamaño muestral e incluir otras covariables antropométricas y clínicas. Asimismo, se evaluará el uso de otras curvas de crecimiento no lineales. En conclusión, este es un primer modelo con capacidad predictiva, que con mejoras permitirá avanzar hacia una insuflación personalizada en una cirugía mínimamente invasiva.

- **Capítulo 3. Modelos longitudinales conjuntos bayesianos.** En el Capítulo 3 desarrollamos modelos longitudinales conjuntos bayesianos para evaluar la pesca de la sardina europea durante las últimas décadas en el mar Mediterráneo. Esta especie es posiblemente el pequeño pez pelágico más importante desde el punto de vista ecológico y económico. En los últimos años se ha observado una preocupante reducción del peso y de la esperanza de vida de este tipo de peces.

En los últimos años, la sardina europea ha experimentado un impactante proceso de fragilidad, con una reducción drástica en su peso corporal (alrededor de dos tercios) y una vida útil mucho más corta que hace 20 años, lo que podría ser un reflejo de la disminución en las capturas. Este declive es el resultado de un proceso complejo probablemente causado por varios factores. Se han propuesto muchas teorías en la comunidad científica: un aumento en los depredadores de sardinas, como el atún rojo y los delfines. Otro elemento importante es la reduc-

ción del plancton, su fuente de alimento. En los últimos años, estos microorganismos marinos se han reducido drásticamente y han perdido valor nutricional debido al calentamiento global, ya que generalmente se alimentan de nutrientes de aguas frías y profundas. Además, ha habido una reciente proliferación de floraciones de medusas, depredadores de huevos y larvas de sardinas y competidores por el plancton. El tamaño cada vez menor de las sardinas, el hecho de que la mayoría se vende fresca en lugar de congelada, la cambiante relación económica que genera sus ventas, dominada últimamente por grandes cadenas de distribución que virtualmente controlan la demanda, y las cambiantes preferencias del consumidor y las diferencias intergeneracionales han contribuido, en mayor o menor medida, a una caída en los precios y capturas de sardinas en el mercado que amenaza la viabilidad de las flotas pesqueras menos poderosas.

Este capítulo se centra en la evolución de las capturas de sardina en el mar Mediterráneo desde 1985 hasta 2018 según el país mediterráneo y el tipo de pesca practicada, artesanal e industrial. Proponemos tres modelos lineales mixtos longitudinales bayesianos para evaluar las diferencias en la evolución temporal de la pesca artesanal e industrial entre y dentro de los países.

En este caso, consideramos una respuesta bivalente: donde $y^{(A)}$ representa el logaritmo de las toneladas artesanales y $y^{(I)}$, la pesca industrial. Estas dos respuestas son condicionalmente independientes dados los parámetros, hiperparámetros, efectos aleatorios y términos de correlación serial. Se consideran tres modelizaciones alternativas. La primera opción es un modelo mixto con intercepts y pendientes comunes y sus subsiguientes efectos aleatorios individuales, con errores de medida distribuidos normalmente. El segundo modelo difiere solo en los errores, en lugar de ser independientes y distribuidos normalmente, son autoregresivos de orden 1. La última opción tiene los cuatro componentes: efectos fijos, efectos aleatorios, términos autoregresivos y errores normales. El punto clave es que asumimos una relación lineal entre los efectos aleatorios individuales de la respuesta artesanal y los de la respuesta industrial a través de dos matrices de varianzas-covarianzas.

Los resultados generales confirman que las series temporales de la pesca mediterránea son muy diversas y que esta heterogeneidad es persistente a lo largo del tiempo. Además, nuestros resultados ponen de manifiesto una correlación positiva entre la pesca artesanal y la industrial. Por último, el estudio observa una tendencia temporal decreciente en la cantidad de pesca durante las últimas décadas. Aunque las causas de esta característica también podrían estar relacionadas con motivaciones económicas (como una reducción en la demanda o la reorientación de las flotas hacia especies más comercialmente beneficiosas), podría indicar un riesgo potencial para la población de esta especie en el mar Mediterráneo.

- **Capítulo 4. Factor Bayes en modelos longitudinales.** Este capítulo es el más metodológico ya que explora el uso de la distribución *power posterior* para aproximar factores Bayes en modelos lineales mixtos longitudinales. Los factores Bayes son la herramienta bayesiana natural para selección de modelos, definidos como el ratio entre las probabilidades marginales o evidencias de dos modelos

competidores. Sin embargo, el cálculo de probabilidades marginales necesarias para los factores Bayes puede ser computacionalmente costoso y numéricamente complejo, especialmente en modelos lineales mixtos longitudinales donde intervienen efectos aleatorios asociados al individuo y términos de correlación en serie. Esto dificulta el cálculo de factores Bayes para este tipo de modelos.

Las funciones *power posterior* permiten expresar evidencias de los modelos como integrales unidimensionales en el rango unitario por medio del algoritmo de muestreo termodinámico. Este trabajo generaliza su uso a modelos lineales mixtos, aunque es fácilmente generalizable a otros modelos con una complejidad similar.

Se presentan dos estudios de simulación para evaluar la capacidad de esta metodología para identificar el modelo generador de los datos dentro de un conjunto de modelos competidores. En el primer estudio los datos provienen de un modelo simple con efectos aleatorios en el intercept que compete con otros 3 modelos. El modelo generador de los datos es correctamente identificado. El segundo estudio incluye datos de un modelo complejo con 3 fuentes aleatorias de variación: efectos aleatorios, un término autorregresivo y errores de medida normales. Este modelo compete con otros más simples con sólo dos de las tres fuentes de variación. Nuevamente el modelo generador de los datos es correctamente seleccionado entre los competidores. En ambos estudios simulados la metodología demuestra eficiencia para identificar el verdadero modelo generador de los datos entre opciones alternativas.

También se analiza un conjunto de datos reales sobre pesca de sardina en países del mar Mediterráneo entre 1970-2014. Se ajustan 3 modelos longitudinales para modelar la dinámica de las pesquerías y se calculan las probabilidades marginales de los modelos con la metodología de las *power posterior*. El modelo con mayor evidencia incluye efectos aleatorios y errores autorregresivos. Un modelo alternativo con tres fuentes de variación (efectos aleatorios, errores autorregresivos y errores normales) presenta evidencia levemente menor. El análisis destaca la relevancia del término autorregresivo en los modelos de la pesca de la Sardina Europea en el mar Mediterraneo.

Esta metodología se basa en el muestreo de Gibbs para cada *power posterior*, y así poder aproximar la evidencia de modelos bayesianos. Su implementación es ideal para modelos longitudinales con efectos aleatorios y términos de correlación en serie. Variaciones en el algoritmo son posibles usando “importance sampling” para evitar muestrear de la distribución a posteriori o distribuciones cercanas a la previa. Esto puede reducir levemente el coste computacional. También pueden realizarse pequeños cambios en la regla trapezoidal que aproxima la evidencia en escala logarítmica para reducir sesgos. En resumen, el uso de *power posterior* para aproximar probabilidades marginales facilita la selección de modelos en el contexto de modelos lineales mixtos longitudinales bayesianos, donde tradicionalmente el cálculo de la evidencia ha sido un desafío.

- **Capítulo 5. Un modelo bayesiano de Markov oculto para analizar el fenómeno hot hand.** El análisis de datos deportivos es un tema relevante de la estadística aplicada que ha ido creciendo en importancia en los últimos años. Este capítulo

lo presenta un novedoso modelo Bayesiano de cadena de Markov oculta longitudinal para analizar el controvertido fenómeno conocido como el *hot hand* en el rendimiento de tiros en partidos de baloncesto. Consideramos que un jugador o equipo de baloncesto tiene *hot hand* cuando su rendimiento durante un partido en un determinado periodo de tiempo es mejor de lo esperado. Este fenómeno ha generado una gran controversia con detractores que afirman su inexistencia mientras que otros autores señalan su evidencia.

El modelo considera cada partido jugado por un equipo durante una temporada como una unidad observacional. Se modelan dos subprocesos interconectados: uno oculto que representa dos posibles estados del equipo, frío (*cold*) o caliente (*hot*), mediante una cadena de Markov; y otro observable en términos de una variable Bernoulli longitudinal que evalúa el éxito (1) o fracaso (0) de cada tiro en relación al estado oculto. Las probabilidades de transición entre estados *cold* y *hot* en la cadena de Markov se modelan mediante modelos de regresión logística mixta, considerando efectos aleatorios específicos para cada partido. Similarmente, las probabilidades de encestar en cada estado se modelan mediante regresión logística mixta, incluyendo la distancia a la canasta y el tipo de tiro (tiro libre o tiro de campo) como covariables, y efectos aleatorios asociados al partido. Este modelo permite realizar inferencia sobre diversos indicadores del rendimiento del equipo, como probabilidades de transición, tiempos de ocupación en cada estado, distribución estacionaria o probabilidades de encestar según distancia en cada estado.

La modelización propuesta se aplicó a datos reales de tiros del equipo Miami Heat en la temporada 2005-2006 de la NBA, que ganaron el campeonato ese año. En total analizan 11042 tiros en 105 partidos, de los cuales 5922 fueron canastas. Los resultados indican que es casi tres veces más probable que el equipo permanezca en estado frío que cambie a estado caliente. La distancia a la canasta es un factor relevante, disminuyendo la probabilidad de encestar a mayor distancia. Hacer un tiro libre parece más fácil que encestar de la misma distancia con un lanzamiento normal. El modelo permite analizar el tiempo que pasa el equipo en cada estado, encontrando que pasan más tiempo en estado frío que caliente. Es decir, la distribución estacionaria de estimada es una probabilidad de 0.61 de estar en *cold* y 0.39 en *hot*. Asimismo, el modelo posibilita inferir la probabilidad de encestar según estado y distancia. Para tiros sencillos, en estado *cold* la probabilidad es de 0.5 mientras que en *hot* es casi 1. A partir de 10 pies en *cold* es casi imposible encestar, mientras que en *hot* hasta los 15 pies es muy probable. El modelo también permite visualizar la probabilidad de encestar durante los distintos partidos. Se observa gran regularidad durante toda la temporada, con poca variabilidad entre partidos en la probabilidad de acierto de triples.

Finalmente, concluimos que el fenómeno de *cold hand* parece tener más ocurrencias que el *hot hand* para el equipo analizado. El modelo propuesto parece una herramienta útil y con gran potencial para analizar en profundidad el rendimiento de equipos de baloncesto. Permite obtener valiosa información sobre rachas y rendimiento que no es posible mediante métodos tradicionales.

- **Capítulo 6. Combinando modelización bayesiana y optimización.** El último capítulo, que también tiene una motivación deportiva, presenta una novedosa metodología para seleccionar las mejores alineaciones de un equipo de baloncesto en silla de ruedas antes de un partido, basándose en métricas del rendimiento de los diferentes jugadores del equipo en encuentros previos. La propuesta combina modelización bayesiana, que analiza los datos de rendimiento de los jugadores a lo largo de la temporada, con un problema de programación lineal entera para determinar las alineaciones que maximizan el rendimiento del equipo sujeto a restricciones de juego. Una de las características principales es que esta propuesta incorpora incertidumbre, lo que propaga variabilidad a la solución óptima.

Esta metodología consta de 3 etapas: 1) Aproximar la distribución a posteriori de los parámetros del modelo longitudinal bayesiano que analiza las métricas de rendimiento de los jugadores. 2) Computar predicciones para un nuevo partido mediante la distribución predictiva a posteriori. 3) Obtener distribución a posteriori de las soluciones del problema de optimización de programación lineal entera que busca maximizar la suma de predicciones bajo ciertas restricciones.

Para ilustrar este procedimiento, se analizan datos de 18 partidos de la temporada 2022-2023 del equipo Doneck Dolphins Trier de 1a división de baloncesto en silla de ruedas de Alemania. Se consideran 3 métricas de desempeño individual: eficiencia del jugador (EFF), índice de rendimiento (PIR) y *Win Score*. El modelo bayesiano longitudinal se ajusta por separado para cada métrica. En el equipo, destaca Dirk Passivan, con un rendimiento algo irregular pero excepcionalmente bueno. Debido a las restricciones del sistema de clasificación funcional de jugadores (PCPS), un alto rendimiento individual no implica necesariamente siempre la inclusión en la alineación.

Los resultados permiten identificar claramente las composiciones con mayor probabilidad de ser las más óptimas. También se estiman las probabilidades de inclusión de cada jugador en la alineación óptima. Sobresalen Annabel Breuer, Correy Rossi, Dirk Passivan y Walter Vlaanderen. Esto demuestra la valiosa información derivada de los resultados. La metodología permite calcular probabilidades de compatibilidad entre jugadores sobre la cancha. También es posible estimar los resultados ante la ausencia de un jugador por suspensión o lesión, considerando específicamente las alineaciones sin dicho jugador.

Aunque el trabajo se enfoca en métricas combinadas de acciones positivas y negativas, puede complementarse con otros tipos de métricas. Es directamente aplicable a otros deportes adaptados con PCPS como fútbol, hockey o rugby en silla de ruedas. Incluso puede adaptarse a deportes sin PCPS reemplazando las restricciones de PCPS por limitaciones propias de cada disciplina. En conclusión, esta metodología tiene un gran potencial para ser utilizadas por entrenadores y cuerpos técnicos, proveyendo ideas valiosas sobre las composiciones ideales para partidos venideros, evaluando jugadores propios y contrarios, e identificando fortalezas y debilidades de diferentes esquemas.

- **Capítulo 7. Apuntes finales.** En el último capítulo de esta memoria se resumen las principales conclusiones obtenidas, las contribuciones tanto prácticas como

metodológicas y las líneas de trabajo futuro.

- **Apéndice A. Códigos de JAGS.** En el primer apéndice del trabajo aparecen todos los modelos en formato JAGS utilizados en la memoria. En particular, aquí podemos encontrar 6 modelos: el modelo no lineal jerárquico del Capítulo 2, las tres modelizaciones alternativas del Capítulo 3, el modelo oculto de Markov del capítulo 5 y el modelo lineal mixto que aparece en el Capítulo 6.
- **Apéndice B. Detalles técnicos.** El último apéndice está dividido en dos partes. La primera es una aclaración matemática del origen de la matriz de varianzas-covarianzas para la segunda modelización del Capítulo 3. En segundo lugar, se enumeran las densidades condicionales a posteriori completas necesarias para calcular las *power posterior* mediante el muestreo de Gibbs en el Capítulo 4 (implementado en el software R).

Conclusiones y trabajo futuro

En primer lugar, estamos convencidos de que la transparencia en el desarrollo científico es un requisito necesario para su crecimiento y mejora continua. Queremos mencionar el esfuerzo llevado a cabo para garantizar la reproducibilidad. Siempre que ha sido posible, todos los análisis que hemos realizado son reproducibles. Por esta razón, los códigos que implementamos y los conjuntos de datos utilizados en nuestro análisis están disponibles de forma gratuita en <https://github.com/gcalvobayarri>.

Respecto a las contribuciones a nivel práctico y metodológico de nuestro trabajo, en el frente práctico, destacamos tres puntos clave. El Capítulo 2 arroja luz sobre el papel crítico del género y la edad en el espacio intra-abdominal óptimo de los pacientes sometidos a una laparoscopia. En el Capítulo 3, se muestra un preocupante declive en las toneladas de Sardina Europea capturadas en la pesca artesanal e industrial en el Mediterráneo durante las últimas tres décadas. Además, se revela una relación lineal positiva entre estos dos tipos de pesca en cada país. Por último, en el Capítulo 5 se descubre algo parecido a un fenómeno de *cold hand* durante la temporada 2005-2006 del equipo de baloncesto Miami Heat de la NBA.

Sin embargo, más allá de estos descubrimientos prácticos, esta investigación contiene contribuciones metodológicas estadísticas básicamente en cuatro aspectos. En el Capítulo 3, se desarrolla un novedoso modelo longitudinal conjunto que establece una relación lineal entre los efectos aleatorios asociados con el mismo individuo en diferentes variables respuesta. El Capítulo 4 introduce una metodología para comparar modelos longitudinales con efectos aleatorios y términos de correlación en serie, calculando su verosimilitud marginal a través del método *power posterior*. El Capítulo 5 muestra la capacidad de los modelos longitudinales bayesianos con estructura de cadenas de Markov ocultas para describir el fenómeno de *hot hand* en el rendimiento del baloncesto a través de una cadena de Markov latente con dos estados, *cold* y *hot*. Por último, el Capítulo 6 integra de manera innovadora técnicas de modelado longitudinal bayesianas y optimización lineal entera. Estas dos técnicas se utilizan para analizar los datos de rendimiento temporal de los jugadores e identificar configuraciones de alineación óptimas para maximizar el rendimiento general, respectivamente. De esta

manera obtenemos una distribución de probabilidad para la solución de un problema de optimización lineal entera.

En cuanto a futuras investigaciones, en el Capítulo 2, se presenta un modelo para estimar con precisión el espacio ideal para realizar laparoscopias basado en información del paciente. En la segunda fase, mejoraremos el modelo anterior y recopilaremos mediciones antropométricas de nuevos pacientes, como el diámetro abdominal sagital y la relación cintura-cadera, con el objetivo de identificar de manera más precisa el volumen intraabdominal óptimo para la cirugía. En el Capítulo 4 nuestro próximo objetivo es comparar esta propuesta con enfoques existentes, como el uso del paquete `bridgesampling` en R, ya utilizado en el Capítulo 3. Además, en el Capítulo 5, al estudiar el fenómeno *hot hand* en el rendimiento de tiro, considerar la inclusión de un tercer estado intermedio podría enriquecer el análisis. Explorar cadenas de Markov de tiempo continuo también podría mejorar la modelización de este fenómeno. Finalmente, extender la metodología del Capítulo 6 a otros deportes podría ser muy interesante, especialmente en deportes de equipo para personas con discapacidad física. Creemos que el sistema de restricciones utilizado en el problema de optimización lineal también puede aplicarse a deportes de equipo que no utilicen un sistema de clasificación funcional. Por ejemplo, en el fútbol, podríamos considerar restricciones como el número necesario de jugadores, el número máximo y mínimo de delanteros, entre otros.

Contents

1	Introduction and framework	1
1.1	Introduction	1
1.2	Bayesian statistics	1
1.2.1	Markov Chain Monte Carlo methods	3
1.3	Repeated measures and longitudinal data	4
1.4	Modelling framework	6
1.5	Outline	8
2	A Bayesian hierarchical nonlinear model for repeated measures data	10
2.1	Introduction	10
2.2	Intra-abdominal volume and intra-abdominal pressure data	13
2.3	Logistic growth mixed-effects modelling	14
2.4	Posterior inferences and predictions	18
2.4.1	Posterior distribution	18
2.4.2	Posterior individual outcomes	20
2.4.3	Posterior population outcomes	22
2.5	Model checking	24
2.6	Conclusions	28
3	Bayesian joint longitudinal models	30
3.1	Introduction	30
3.2	Sardine fisheries in the Mediterranean Sea	34
3.3	Joint Bayesian longitudinal modelling	37
3.3.1	Longitudinal modelling for the total European sardine landings	38
3.4	Posterior inferences	40
3.4.1	Model comparison	43
3.4.2	Association and country-specific fishing patterns	46
3.5	Conclusions	48
4	Bayes factor in longitudinal models	50
4.1	Introduction	50
4.2	Bayesian longitudinal linear mixed models	53
4.3	Bayes factor and power posteriors	55
4.4	Simulation studies	57
4.4.1	Study 1: a balanced longitudinal data set	58

4.4.2	Study 2: an unbalanced longitudinal data set with serial correlation	60
4.5	Application - Sardine landings in the Mediterranean Sea	62
4.5.1	Modelling of sardine fisheries	63
4.5.2	Longitudinal modelling of European sardine landings	65
4.6	Conclusions	67
5	A Bayesian hidden Markov model for assessing the hot hand phenomenon	69
5.1	Introduction	69
5.2	Bayesian longitudinal hidden Markov modelling	71
5.3	Basketball team performance	74
5.3.1	Transition probabilities	74
5.3.2	t -step transition probabilities and stationary distribution . . .	75
5.3.3	Occupancy times	76
5.3.4	Sojourn times	77
5.3.5	Probabilities of making a basket	77
5.4	Miami Heat shooting in the 2005-2006 NBA season	78
5.4.1	Modelling the shooting performance of the Miami Heat team .	79
5.4.2	Posterior distribution	81
5.5	Conclusions	87
6	Combining Bayesian modelling and optimization	88
6.1	Introduction	88
6.2	Methodology	90
6.2.1	Modelling the performance	90
6.2.2	Prediction for a new game	90
6.2.3	Selecting the optimal composition	91
6.3	Doneck Dolphins Trier case	92
6.3.1	Data	92
6.3.2	Applying the proposed methodology	94
6.3.3	Posterior results	97
6.4	Conclusions	102
7	Final remarks	103
7.1	Final comments	103
7.2	Future research	105
A	JAGS codes	107
A.1	Model Section 2.3 - Logistic growth mixed-effects model	107
A.2	Model Section 3.3 - $\mathcal{M}_1$	108
A.3	Model Section 3.3 - $\mathcal{M}_2$	110
A.4	Model Section 3.3 - $\mathcal{M}_3$	112
A.5	Model Section 5.4 - Shooting performance model of the Miami Heat team	114
A.6	Model Section 6.3 - Longitudinal performance model of the Doneck Dolphins Trier team	116

B	Technical details	118
B.1	Autoregressive variance-covariance matrix	118
B.2	Complete posterior conditional distribution associated with the power posteriors	119
B.2.1	Model $\mathcal{M}_1$ - Study 2	119
B.2.2	Model $\mathcal{M}_2$ - Study 2	120
B.2.3	Model $\mathcal{M}_3$ - Study 2	121

List of Figures

1.1	Two examples of spaghetti plots with 10 experimental units.	6
2.1	Number of repeated measures in the men's group (top panel) and in the women's group (bottom panel). Each bar corresponds to a person and its ordinate describes the number of measurements of that person during the study. Patients are ordered according to their age from the youngest to the oldest.	14
2.2	<i>I</i> <i>AV</i> profiles (in L) according to <i>I</i> <i>AP</i> (in mmHg) for men (top panel) and women (bottom panel) in the sample.	15
2.3	Graphics of the logistic growth function $5/[1 + \exp\{-(-10 + x)\}]^{-1}$, the subsequent asymptotic value, and its <i>MAP</i> , <i>IP</i> , <i>ADP</i> , and <i>MDP</i> points.	17
2.4	Posterior mean and 95% credible interval of the <i>ADP</i> (<i>I</i> <i>AP</i> value) of the men (top panel) and the women (bottom panel) in the sample. Patients are ordered in the <i>x</i> -axis according to their age, from the youngest to the oldest.	21
2.5	Posterior predictive mean and 95% predictive interval for <i>I</i> <i>AV</i> in relation to <i>I</i> <i>AP</i> values. The top panel represents a man, while the bottom panel represents a woman. Both individuals are 64.65 years old, which corresponds to the sample mean.	22
2.6	Posterior predictive mean of the <i>I</i> <i>AV</i> with regard to <i>I</i> <i>AP</i> values for a man (top panel) and a woman (bottom panel) stratified by age.	23
2.7	Posterior distribution of the asymptotic <i>I</i> <i>AV</i> for men (left) and women (right).	24
2.8	Joint posterior distribution and contour lines of the <i>I</i> <i>AP</i> and <i>I</i> <i>AV</i> coordinates for the true <i>ADP</i> and posterior marginal distribution for each of both quantities for men (top panel) and women (bottom panel) aged 64.65 years (the sample mean). The horizontal and vertical lines represent the approximate posterior mean of <i>I</i> <i>AV</i> – <i>ADP</i> and the approximate posterior mean of <i>I</i> <i>AP</i> – <i>ADP</i> , respectively.	25
2.9	Histogram of the approximate <i>CPO</i> quantities for all the observations.	27
2.10	Density histogram of the approximate <i>PIT</i> quantities for all the observations.	27
2.11	Distribution of <i>PIT</i> values along the different values of <i>I</i> <i>AP</i>	27

3.1	Annual thousands of tons of European sardines caught by the main Mediterranean fishing countries from 1985 to 2018. Data come from Sea Around Us (www.seaaroundus.org).	32
3.2	Countries bordering the Mediterranean Sea. Photo by Ian Macky, PAT Atlas, https://ian.macky.net/pat/map/medi/mediblu.gif	35
3.3	Annual tons of European sardine (<i>Sardina pilchardus</i>) per country, from 1985 to 2018. From top to bottom: artisanal and industrial and fishing.	36
3.4	Approximate marginal posterior distribution of $\beta_1^{(A)}$ and $\beta_1^{(I)}$, respectively, for model $\mathcal{M}_3$	46
3.5	Artesanal vs. Industrial fishing. Approximate posterior mean of the random effects parameters from the model $\mathcal{M}_3$ and their estimated correlations (dashed lines): (a) random intercepts and (b) random trends.	47
3.6	Approximate marginal posterior distribution of ρ_0 and ρ_1 , respectively, for model $\mathcal{M}_3$	48
4.1	Spaghetti plot of the simulated longitudinal data from the first study, 5 individuals are considered in a complete balanced design for times $\{0, 1, \dots, 9\}$	59
4.2	Spaghetti plot of the unbalanced longitudinal data from the second study, 10 individuals are considered with a total of 441 observations.	61
4.3	Annual log tons of European sardines (<i>Sardina pilchardus</i>) caught per country, from 1970 to 2014.	63
4.4	Random intercepts vs. random slopes. Approximate posterior mean for the random intercepts and the slope effects by country according to $\mathcal{M}_2$	66
4.5	Posterior mean (line) and 95% credible interval (shaded area) of the autocorrelation function according to the posterior distribution of the autoregressive parameter ρ in $\mathcal{M}_2$	67
5.1	Shot chart of the Miami Heat during the 2005-2006 season, showing the location of shots on the court, and indicating the shot results and the exact position of the basket.	79
5.2	Posterior distribution histograms of the transition probabilities (a) $p^{(CC)}$ (in purple) and $p^{(CH)}$ (in yellow), and (b) $p^{(HC)}$ (in purple) and $p^{(HH)}$ (in yellow) integrating the random effects associated with the game.	83
5.3	Posterior distribution histograms of the occupancy times in the states <i>hot</i> (in yellow) and <i>cold</i> (in purple) in a game with $n = 120$ shots that starts (a) in the <i>cold</i> and (b) in the <i>hot</i> state.	84
5.4	Violin and box plot of the posterior probability distribution for a <i>cold</i> and a <i>hot</i> streak.	85
5.5	Posterior means and 95% credible interval for the probabilities of making a basket depending on the distance, in feet, in three different scenarios. Purple curve represents the case when the team is in the <i>cold</i> state, the yellow plot is for the <i>hot</i> state, and the green plot is the general case when the state of the team is unknown.	86

5.6	Posterior mean and 95% credible interval for the probability of making a basket from 25 feet distance in each of the games of the season when the state of the team is unknown.	86
6.1	Spaghetti plots of the three performance ratings per minute by player for the Doneck Dolphins Trier team during the 2022-2023 RBBL regular season.	95
6.2	Bar plots showing the estimated probabilities associated with the most optimal line-ups based on the three performance metrics for the seven compositions: L_1 , L_2 , L_3 , L_4 , L_5 , L_6 , and L_7	99
6.3	Estimated probabilities $P(I_i \mid \mathcal{D})$ of player i being included in the optimal line-up team based on EFF, PIR and Win Score metrics. . . .	100
6.4	Estimated probabilities of the inclusion of players Dejon Green, Lucas Jung, Natalie Passivan, Patrick Dörner and Svenja Erni in the line-up team given that players Annabel Breuer, Correy Rossi, Dirk Passivan and Walter Vlaanderen have already been chosen, based on EFF, PIR and Win Score metrics.	101

List of Tables

2.1	Posterior summaries (mean, standard deviation and 95% credible interval) for the parameters and hyperparameters of the logistic growth model with covariates gender and standardized age.	20
3.1	Summary of the approximated posterior distribution of the parameters and hyperparameters in models $\mathcal{M}_1$, $\mathcal{M}_2$ and $\mathcal{M}_3$	42
3.2	Approximated PED for models $\mathcal{M}_1$, $\mathcal{M}_2$ and $\mathcal{M}_3$ and mean (standard deviation) of ten replicates of the approximate evidence, in logarithmic scale, for models $\mathcal{M}_1$, $\mathcal{M}_2$ and $\mathcal{M}_3$	45
3.3	CPO 's, in logarithmic scale, for the different countries in the study computed from models $\mathcal{M}_1$, $\mathcal{M}_2$ and $\mathcal{M}_3$	45
4.1	Mean (standard deviation) of the ten replicates of the approximate log evidence for each model in (4.12) computed by means of the power posterior method.	60
4.2	Mean (standard deviation) of the ten replicates of the approximate log evidence for models in (4.13) and (4.14) computed by means of the power posterior method.	61
4.3	Mean (standard deviation) of the approximate log evidence for each model in (4.15) by means of the power posterior method.	65
4.4	Summary of the approximate posterior distribution of the parameters and hyperparameters in model $\mathcal{M}_2$	65
5.1	Posterior summaries (mean, standard deviation, 95% credible interval and $\hat{R}$) for the parameters and hyperparameters of the Miami Heat shooting performance <i>hot hand</i> model.	81
5.2	Posterior summaries (mean, standard deviation and 95% credible interval) for the stationary distribution of the hidden Markov chain associated with a generic game.	84
6.1	Name, functional classification value and sex of the players of the Don-neck Dolphins Trier team who play more than 40 minutes during the 2022-2023 RBBL season.	93

Notation

Symbol	Description
$\mathbf{y}$	Response vector
$\mathbf{y}_i$	Individual response vector
y_{ij}	Response for individual i at measurement j
$\mathbf{Z}$	Latent variables
$\boldsymbol{\theta}$	Parameters and hyperparameters vector
θ	Quantity of interest
$\bar{\theta}$	Vector of the quantities of interest
U	Length of $\bar{\theta}$
ϕ	Vector of the random effects and serial correlation
$\mathbf{b}$	Vector of the random effects
$\mathbf{b}_i$	Vector of the random effects associated with the individual i
b_{0i}	Random intercept associated with the individual i
b_{1i}	Random slope associated with the individual i
$b_i^{(CH)}$	Individual random effect associated with the transition from C to H
$b_i^{(HC)}$	Individual random effect associated with the transition from H to C
$\mathbf{w}_i$	Autoregressive vector for the individual i
w_{ij}	Autoregressive element for the individual i at measurement j
$\boldsymbol{\beta}$	Vector of the fixed effects
β_0	Vector of the common intercepts
$\boldsymbol{\beta}^{(CH)}$	Common coefficients vector associated with the transition from C to H
$\boldsymbol{\beta}^{(HC)}$	Common coefficients vector associated with the transition from H to C
β_0	Common intercept
$\boldsymbol{\alpha} = (\boldsymbol{\alpha}^{(C)}, \boldsymbol{\alpha}^{(H)})$	Vector of the fixed effects for the observed process in the BLHMM
$\mathbf{a} = (\mathbf{a}^{(C)}, \mathbf{a}^{(H)})$	Vector of the random effects for the observed process in the BLHMM
$\boldsymbol{\mu}$	Mean vector
$\boldsymbol{\mu}_i$	Mean vector for the individual i
μ_{ij}	Mean for the individual i at measurement j

N	Number of individuals
Q	Number of variables of interest
A, B	Random events
$\mathcal{D}$	Data
$\mathcal{D}_i$	Data associated with the individual i
f	Function
g	Function (in Metropolis-Hastings algorithm)
π	Distribution
$m(\mathcal{D})$	Marginal likelihood (or evidence)
s	Iteration
q_s	Value simulated from $g(\cdot)$ at the s -th iteration
κ	Metropolis-Hastings ratio
n_i	Number of measurements for the individual i
$\mathbf{t}_i$	Vector of time indicator for the individual i
t_{ij}	Observation at the j -th time for the individual i
$\mathcal{X}$	Design matrix for fixed effects
$\mathcal{Z}$	Design matrix for random effects
$\mathcal{X}_i$	Design matrix for fixed effects associated with the individual i
$\mathcal{Z}_i$	Design matrix for random effects associated with the individual i
$I_W(i)$	Sex indicator variable for the individual i
Age_i	Age of the individual i
Σ	Variance-covariance matrix
Σ_i	Variance-covariance matrix associated with the individual i
Σ_b	Variance-covariance matrix associated with the random effects
σ	Standard deviation
σ_w	Standard deviation associated with an autoregressive element
σ_0	Standard deviation associated with the random intercepts
σ_1	Standard deviation associated with the random slopes
ρ	Correlation
η	Autoregressive parameter
IAP_{ij}	Intra-abdominal pressure observed for the individual i at measurement j
a_i, b_i, c_i	Logistic growth parameters for the individual i
$\phi_i^{(a)}$	Random effect associated with the individual i in the a_i parameter
$\phi_i^{(b)}$	Random effect associated with the individual i in the b_i parameter
P_i	Transition matrix for the individual i
$p_i^{(CC)}$	Transition probability from state C to C for the individual i
$p_i^{(CH)}$	Transition probability from state C to H for individual i
$p_i^{(HH)}$	Transition probability from state H to H for individual i
$p_i^{(HC)}$	Transition probability from state H to C for the individual i
$P_i^{(t)}$	The t -step transition matrix for individual i
$p_i^{(CC)}(t)$	The t -step transition probability from state C to C for the individual i
$p_i^{(CH)}(t)$	The t -step transition probability from state C to H for the individual i

$p_i^{(HH)}(t)$	The t -step transition probability from state H to H for the individual i
$p_i^{(HC)}(t)$	The t -step transition probability from state H to C for the individual i
$\mathbf{X}_i$	Baseline covariates vector for the individual i
$\mathbf{X}_{ij}$	Baseline covariates vector for individual i at measurement j
$\mathcal{N}(\cdot, \cdot)$	Normal distribution
$\mathcal{B}(\cdot, \cdot)$	Bernoulli distribution
$\mathcal{U}(\cdot, \cdot)$	Uniform distribution
$\text{Be}(\cdot, \cdot)$	Beta distribution
$\gamma_i^{(C)}$	Success probability vector associated with state C for individual i
$\gamma_i^{(H)}$	Success probability vector associated with state H for individual i
$\gamma_{ij}^{(C)}$	Success probability associated with state C for individual i at game j
$\gamma_{ij}^{(H)}$	Success probability parameter associated with state H for i at game j
Δ_i	Limit distribution vector of the HMC for individual i
$\tau_i^{(k)}$	Sojourn time for the HMC of the individual i in state k
$\mathcal{M}_k$	Model k
β_1	Common trend parameter
ρ_0	Correlation associated with the random intercepts variance-covariance matrix
ρ_1	Correlation associated with the random trends variance-covariance matrix
I	Diagonal unit matrix
$\text{Cov}()$	Covariance
z_i	Indicator variable of player i
C_i	Functional classification of player i
b_{0j}	Random effect associated with the match j
σ_{0m}	Standard deviation parameter associated with the matches
β_W	Coefficient parameter associated with sex
β_C	Coefficient parameter associated with the functional classification
β_H	Coefficient parameter associated with the location of the match
L_k	Composition number k of the team
I_i	Event of including player i in line-up team

Chapter 1

Introduction and framework: Bayesian statistics, repeated measures, and longitudinal data

1.1 Introduction

As the opening chapter of this memory, we provide a comprehensive overview of the fundamental concepts and ideas that form the overarching framework of this thesis. In particular, we begin by presenting a brief introduction to Bayesian statistics, outlining its key principles and ideas. Then, we introduce the concepts of repeated measures and longitudinal data, illustrating their distinctive characteristics. Finally, we present the broad mathematical framework that underpins all of our work.

1.2 Bayesian statistics

In the very known paper entitled “An Essay towards solving a Problem in the Doctrine of Chances” (Bayes, 1763), specifically in its fifth proposition, the powerful law of

probability that we know today as Bayes theorem is described:

$$P(A | B) = \frac{P(B | A)P(A)}{P(B)},$$

being A and B two random events where $P(A) > 0$ and $P(B) > 0$. That is, the Bayesian rule allows calculating the conditional probability of one event given another.

Stigler (1986) assures that this theorem was ignored until it was rediscovered in a memoir by Laplace (1749-1827) when he was 25 years old. The influence of this memoir has been immense, especially from the 20th century to the present day.

This rule is the core of Bayesian statistics, which, in a very general way, has as its main goal, like any kind of statistics, to learn about an unknown population from a sample of that population. Indeed, Bayesian inference makes use of the Bayes theorem to update information about a quantity of interest (which we call θ) based on the data $\mathcal{D}$. This can be expressed mathematically as follows

$$\pi(\theta | \mathcal{D}) = \frac{f(\mathcal{D} | \theta)\pi(\theta)}{m(\mathcal{D})} \propto f(\mathcal{D} | \theta)\pi(\theta), \quad (1.1)$$

where $\pi(\theta | \mathcal{D})$ is the posterior distribution for θ given the data $\mathcal{D}$, $f(\mathcal{D} | \theta)$ is the likelihood function of θ when the data $\mathcal{D}$ have been observed, and $\pi(\theta)$ is the prior distribution of θ . Finally, $m(\mathcal{D})$ is known by different terms in the literature such as marginal likelihood (Newton and Raftery, 1994), predictive distribution value (Gelfand and Dey, 1994), marginal probability (Kass and Raftery, 1995), predictive probability (Lartillot and Philippe, 2006) or evidence (Friel, Hurn, and Wyse, 2014). In short, the marginal likelihood is the value of the joint probability distribution of the data $m(\mathcal{D})$ when θ is integrated out. That is, $m(\mathcal{D})$ is the normalising constant of the posterior density. This concept is crucial for Bayesian model comparison, as we will see in Chapter 4.

It is worth noting that in Bayesian statistics, the quantity of interest θ is modelled using a probability distribution. The unknown θ can represent parameters, hyperparameters, hypotheses, models, missing data, and more. This flexibility to model any

unknown element is one of the key advantages of the Bayesian paradigm.

However, the posterior distribution $\pi(\theta \mid \mathcal{D})$ for our quantity of interest θ rarely has an analytical expression. For this reason, a vast number of computational proposals have emerged to simulate or approximate this posterior distribution. Although some of these approaches, such as the Integrated Nested Laplace Approximation (INLA) developed by Rue, Martino, and Chopin (2009), are extremely interesting, according to Efron and Hastie (2021), the most popular computational methods (which we used in this memory) are the Markov Chain Monte Carlo (MCMC) methods (Gelfand and Smith, 1990; Tanner, 2012). Softwares like WinBUGS (Lunn et al., 2000), OpenBUGS (Spiegelhalter et al., 2014) or JAGS (Plummer, 2003) implement these techniques in order to approximate the target posterior distribution in a given model. In addition, this methodology can be programmed in the R language (Ihaka and Gentleman, 1996) as we do in the study developed in Chapter 4.

1.2.1 Markov Chain Monte Carlo methods

The fundamental idea behind MCMC methods is to build an ergodic Markov chain whose limit distribution is in fact the target posterior distribution of θ . By using MCMC to generate a sequence of samples from this Markov chain, one can approximately sample from the posterior distribution without requiring an explicit analytical solution.

In line with Casella and Robert (2010), we define the quantity of interest as a vector $\bar{\theta} = (\theta_1, \dots, \theta_U)$. To generate a Markov chain using the MCMC methods, we can use the basic Gibbs algorithm to move from $\bar{\theta}^{(s)}$ to $\bar{\theta}^{(s+1)}$, at each iteration s :

1. Simulate $\theta_1^{(s+1)}$ from $\pi_1(\theta_1 \mid \mathcal{D}, \theta_2^{(s)}, \dots, \theta_U^{(s)})$.
2. Simulate $\theta_2^{(s+1)}$ from $\pi_2(\theta_2 \mid \mathcal{D}, \theta_1^{(s+1)}, \theta_3^{(s)}, \dots, \theta_U^{(s)})$.
- $\vdots$
- U . Simulate $\theta_U^{(s+1)}$ from $\pi_U(\theta_U \mid \mathcal{D}, \theta_1^{(s+1)}, \theta_2^{(s+1)}, \dots, \theta_{U-1}^{(s+1)})$.

The conditional probability distributions $\pi_1, \dots, \pi_U$ are usually called full posterior conditional distributions, which must be known to compute the MCMC algorithm. However, when a full conditional distribution π_u ($u = 1, \dots, U$) has not a closed expression, it is possible to implement a Metropolis-Hastings algorithm at each step. This is based on the specification of a proposal distribution $g(\cdot)$ and consists of:

1. Simulate $q_{s+1} \sim g(\cdot)$.

2. Take

$$\theta_u^{(s+1)} = \begin{cases} q_{s+1} & \text{with probability } \min\{\kappa, 1\} \\ \theta_u^{(s)} & \text{otherwise,} \end{cases}$$

where κ is a ratio which assesses the suitability of the new simulated value q_{s+1} as opposed to the previous value $\theta_u^{(s)}$,

$$\begin{aligned} \kappa = & \frac{f(\mathcal{D} \mid \theta_1^{(s+1)}, \theta_2^{(s+1)}, \dots, \theta_{u-1}^{(s+1)}, q_{s+1}, \theta_{u+1}^{(s)}, \dots, \theta_{U-1}^{(s)})}{f(\mathcal{D} \mid \theta_1^{(s+1)}, \theta_2^{(s+1)}, \dots, \theta_{u-1}^{(s+1)}, \theta_u^{(s)}, \theta_{u+1}^{(s)}, \dots, \theta_{U-1}^{(s)})} \times \\ & \frac{\pi(\theta_1^{(s+1)}, \theta_2^{(s+1)}, \dots, \theta_{u-1}^{(s+1)}, q_{s+1}, \theta_{u+1}^{(s)}, \dots, \theta_{U-1}^{(s)})}{\pi(\theta_1^{(s+1)}, \theta_2^{(s+1)}, \dots, \theta_{u-1}^{(s+1)}, \theta_u^{(s)}, \theta_{u+1}^{(s)}, \dots, \theta_{U-1}^{(s)})} \times \frac{g(\theta_u^{(s)})}{g(q_{s+1})}. \end{aligned}$$

It should be noted that in order to use the Metropolis-Hastings algorithm, it is a necessary condition that the proposal distribution $g(\cdot)$ is known and that we are able to simulate values from it.

Finally, it is important to highlight that, as for any Markov chain, the MCMC algorithm requires the initialisation of the chain at $\bar{\theta}^{(1)}$.

1.3 Repeated measures and longitudinal data

Repeated measures data can be defined as repeated observations of several experimental units, subjects or individuals, under different experimental conditions (Lindstrom and Bates, 1990). In particular, longitudinal data are repeated measures taken over time for each of the experimental unit in the study.

These types of data are generally characterised as multivariate (Weiss, 2005), i.e. they are observations collected on different variables. Therefore, in longitudinal or repeated measures studies, it is common to have many response variables (possibly along with covariates) repeatedly observed. Moreover, in this context, it is important to note that the number of observations per individual may vary. As a consequence, one may encounter balanced data, where an equal number of measurements are collected for each experimental unit, and these measurements are gathered under identical conditions. As an alternative, we may encounter unbalanced data, in which the number of observations per individual varies or has been taken at different points in time or values of the covariates.

An additional intriguing aspect of longitudinal and repeated measures analysis is that it provides insights on two distinct levels or hierarchies. Firstly, the data collected from all subjects offers an understanding of the population pattern, and secondly, all observations from the same individual provide information regarding their individual behaviour.

To achieve comprehensive visualisation of these data, it is recommended to use line segments to connect consecutive observations within a subject. This type of graphic is commonly known as spaghetti plot or profile plot. An illustration of this type of visual representation is presented in Figure 1.1, which depicts the behaviour of ten simulated subjects in two different contexts. Firstly, Figure 1.1a displays a balanced data set with an equal number of measurements per individual under the same experimental conditions. Conversely, Figure 1.1b clearly illustrates unbalanced data.

Finally, repeated measures and longitudinal data play a crucial role in various knowledge fields (Weiss, 2005), including medicine, public health, education, economics, psychology, biology, fisheries, and sports. In our work, we address these data types across different chapters. Chapter 2 delves into a medical study involving 198 patients and explores the relationship between intra-abdominal pressure during laparoscopy surgery and the produced intra-abdominal volume. Meanwhile, Chapter 3 focuses on a fishery analysis, specifically examining European Sardine fishing in the




Figure 1.1: Two examples of spaghetti plots with 10 experimental units.

Mediterranean Sea by country. Additionally, Chapters 5 and 6 shift the focus to sports analysis, specifically professional basketball and wheelchair basketball. The field of sports provides a rich source of longitudinal data, presenting untapped potential for accurate longitudinal analysis.

1.4 Modelling framework

Repeated measures models, particularly those involving longitudinal data (Diggle et al., 2002), are complex statistical models. They are specifically designed for the analysis of panel or repeated measures data in various fields such as medical research, social science, sport science and economics among others. These models generally account for correlated data within each experimental unit and independent data between different experimental units, allowing for the estimation of general behavioural patterns of the population as well as specific characteristics of individuals or groups under study. Common sources of variability in the data include the heterogeneity of individuals in the population, measurement errors, and serial correlation, which accounts for fluctuations in observations of the same individual. Moreover, the response behaviour of the subjects can be influenced by latent processes, which can definitely impact on the pattern of the response variable, although they are unobservable.

We consider the random vector response $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_i, \dots, \mathbf{y}_N)$ for Q response

variables and N individuals participating in the study, where $\mathbf{y}_i = (\mathbf{y}_i^{(1)}, \dots, \mathbf{y}_i^{(Q)})$ is the response for individual i ($i = 1, \dots, N$). A generic Bayesian model for $\mathbf{y}$ can be formulated through the joint probability distribution:

$$f(\mathbf{y}, \mathbf{Z}, \boldsymbol{\theta}, \boldsymbol{\phi}) = f(\mathbf{y} \mid \mathbf{Z}, \boldsymbol{\theta}, \boldsymbol{\phi})f(\mathbf{Z} \mid \boldsymbol{\theta}, \boldsymbol{\phi})f(\boldsymbol{\phi} \mid \boldsymbol{\theta})\pi(\boldsymbol{\theta}), \quad (1.2)$$

where $f(\mathbf{y} \mid \mathbf{Z}, \boldsymbol{\theta}, \boldsymbol{\phi})$ is the conditional distribution of $\mathbf{y}$ given an unobserved process of the model or a set of latent variables $\mathbf{Z}$, the vector $\boldsymbol{\phi}$ of random effects and serial correlation terms, and the vector $\boldsymbol{\theta}$ of parameters and hyperparameters of the model. In addition, $f(\mathbf{Z} \mid \boldsymbol{\theta}, \boldsymbol{\phi})$ denotes the conditional distribution of $\mathbf{Z}$ given $\boldsymbol{\phi}$ and $\boldsymbol{\theta}$, $f(\boldsymbol{\phi} \mid \boldsymbol{\theta})$ is the conditional distribution of $\boldsymbol{\phi}$ given $\boldsymbol{\theta}$, and $\pi(\boldsymbol{\theta})$ is the prior distribution of $\boldsymbol{\theta}$.

Henceforth, conditional independence among individuals is assumed given $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$. Therefore, the joint distribution of $\mathbf{y}$, $\mathbf{Z}$, $\boldsymbol{\theta}$, and $\boldsymbol{\phi}$ in (1.2) can be expressed as the product of the subsequent conditional distributions for each individual:

$$f(\mathbf{y}, \mathbf{Z}, \boldsymbol{\theta}, \boldsymbol{\phi}) = \left(\prod_{i=1}^N f(\mathbf{y}_i^{(1)}, \dots, \mathbf{y}_i^{(Q)} \mid \mathbf{Z}, \boldsymbol{\theta}, \boldsymbol{\phi}) \right) f(\mathbf{Z} \mid \boldsymbol{\theta}, \boldsymbol{\phi}) f(\boldsymbol{\phi} \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta}). \quad (1.3)$$

This framework is applied in Chapter 5 to investigate the hot hand phenomenon in basketball (with $Q = 1$). Specifically, a latent process $\mathbf{Z}$ in the form of hidden Markov chains is used to model the shooting performance of professional teams.

However, it is not always necessary to incorporate hidden processes in the model. Instead, it may suffice to model the response variable depending on $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ only. If $\mathbf{Z}$ is not included and conditional independence of the individuals is taken into account, the resulting expression of (1.3) is

$$f(\mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\phi}) = \left(\prod_{i=1}^N f(\mathbf{y}_i^{(1)}, \dots, \mathbf{y}_i^{(Q)} \mid \boldsymbol{\theta}, \boldsymbol{\phi}) \right) f(\boldsymbol{\phi} \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta}). \quad (1.4)$$

This is the formulation we used most in this work. We consider a single response variable ($Q = 1$) for the models of Chapter 2, Chapter 4 and Chapter 6. Finally, in

Chapter 3 we analyse two response variables ($Q = 2$) jointly modelled to assess the dynamics of the industrial and artisanal fishing in the Mediterranean Sea.

1.5 Outline

After introducing the key topics covered in this thesis, the remaining chapters are structured as follows. Chapter 2 assesses the relationship between the intra-abdominal pressure insufflated during laparoscopic surgery and the corresponding intra-abdominal volume generated. This is accomplished through a Bayesian nonlinear growth mixed-effects model. We also discuss the plausibility of this model in representing these repeated measures by conducting a model checking process, and suggest modelling improvements.

Chapter 3 presents three different Bayesian longitudinal models that jointly analyse the annual fishery pattern of the industrial and artisanal European Sardine fisheries by all of the Mediterranean countries. Useful information is extracted from the three proposal models. In order to select the model that better suits the data, model comparison is conducted employing various methodologies.

In Chapter 4, we explore the use of power posteriors in longitudinal linear mixed models to approximate marginal likelihoods. The purpose of this exploration is to enable a comparison of these complex models. Furthermore, we discuss the effectiveness of this methodology through two simulation studies and a real data set regarding the landings in the Mediterranean Sea of the European Sardines.

The following two chapters are related to sports data analytics. This is a relevant topic in applied statistics that has been growing in importance in the recent years. In particular, in Chapter 5 we present a Bayesian longitudinal hidden Markov model that analyses the hot hand phenomenon in consecutive basketball shots, each of which can be either missed or made. Two possible states (cold or hot) are assumed in the hidden Markov chains of events, and the probability of success for each shot is modelled by considering both the corresponding hidden state and the distance to the basket.

This model is applied to a real data set, the performance of the Miami Heat team in the season 2005-2006 of the USA National Basketball Association. On the other hand, Chapter 6 introduces a mathematical tool that calculates the best team line-ups by using historical performance data. Our proposed methodology involves Bayesian longitudinal modelling combined with integer linear programming to optimise the line-up of wheelchair basketball teams. We consider three distinct performance ratings for each player and incorporate uncertainty from the posterior predictive distribution of the longitudinal model into the optimisation process. To validate our approach, we use a real data set from a team competing in the Rollstuhlbasketball Bundesliga, called the Doneck Dolphins Trier.

In the last chapter, Chapter 7, we present the main conclusions of our research and provide suggestions for future research directions. Finally, we have included two appendices at the end of this thesis. Appendix A contains all the JAGS codes used in this work, whereas Appendix B presents the complete conditional densities required to compute the methodology based on the power posterior in Chapter 4, along with a more detailed explanation of the variance-covariance matrix for model $\mathcal{M}_2$ in Chapter 3.

Chapter 2

A Bayesian hierarchical nonlinear model to analyse the intra-abdominal volume during pneumoperitoneum for laparoscopic surgery

2.1 Introduction

This chapter is derived from our article (Calvo et al., 2021) published on the Statistics and Operations Research Transactions (SORT) journal. It focuses on a medical study that aims to evaluate the connection between the intra-abdominal volume generated at the start of laparoscopic surgery and the quantity of insufflation within the abdomen.

Laparoscopy is an operation carried out in the abdomen or pelvis through small incisions with the help of a camera. It is performed by insufflating CO_2 into the abdomen that yields a working space, i.e., pneumoperitoneum, and passing surgical instruments through small incisions using a camera to have external visual control of

the procedure (Neugebauer et al., 2010). Laparoscopy has been gaining ground since its inception because it is associated with less morbidity than the traditional method performed through a single, larger skin incision (Pache et al., 2017).

The introduction of CO_2 into the abdomen is performed by laparoscopic insufflators, inserted in the patient's abdominal wall. The CO_2 insufflation pressure, i.e., intra-abdominal pressure (IAP), is set manually on the insufflator by the surgical team. IAP is measured in millimeters of mercury (mmHg), and the usual figures during laparoscopic surgery range between 12 and 15 mmHg. Although international guidelines recommend working with the lowest IAP value that ensures an adequate working space, the standard practice is still to initially set the IAP value without further adjustments regardless of the amount of generated intra-abdominal volume (IAV) (Neudecker et al., 2006), measured in litres (L). Operating at such high IAP increases perioperative morbidity since it leads to decrease abdominal blood perfusion, greater postoperative pain, peritoneal injury, and increased risk of pulmonary complications.

The abdominal compartment shows an anisotropic behaviour during pneumoperitoneum. Initially, marginal gains in volume in response to pressure increments are proportional. In other words, the abdominal compliance (C_{abd}) which defines the change in volume determined by a change in pressure, follows an approximately linear relationship (Mulier et al., 2009). According to biomechanics laws, the yield stress point is eventually reached, after which applying additional pressure leads to diminishing gains in volume (Forstemann et al., 2011). Identifying this critical point at which insufflation should be limited is crucial to maximizing surgical working space while minimising injurious IAP effects.

The abdomen pressure–volume dynamics during pneumoperitoneum has been discussed in previous papers (Díaz-Cambronero et al., 2019; Díaz-Cambronero et al., 2020; Mazzinari et al., 2020; Mazzinari et al., 2021). These studies suggest the adequacy of an increasing sigmoidal model for describing the relationship between both variables. One aim of this chapter is to estimate such a model to gain knowledge about the relationship between IAP and IAV , especially about the parameters that

determine the different growth stages of the process in accordance with the specific characteristics of the individuals in the target population. The second goal is to discuss the quality of the fit of the model to the data. This is a relevant question since the logistic growth curve was already used in the past to tackle this kind of problem. The hypothesis now is that, in a personalised medicine environment, patient responses to insufflation can be estimated and predicted so that an ideal *IAP* value could be determined to optimise *IAV* with the lowest risks of potential negative effects.

The statistical framework of this study are nonlinear growth mixed-effects models, also known as hierarchical nonlinear growth models. They have a long and important scientific tradition for describing biological, medical, and environmental growth phenomena such as pharmacokinetics, epidemiology (Lindsey, 2001), physiological-response processes (Peek et al., 2002), or forestry (Fang and Bailey, 2001) among others. One of the major appeals of these models is that their parameters contain direct and intuitive information on the process under study. This fact generates a multifaceted knowledge about the phenomena in question of great scientific value (Davidian, 2008).

The data for the study come from a repeated measures design. In our case, the variable of interest *IAV* is measured for each individual with regard to different *IAP* values. This design generates two types of data: data from the same individual and data from several individuals. Random effects in these models are essential elements to link the different observations of the same individual with the within-individual variation (Laird and Ware, 1982).

Bayesian statistics simplifies the implementation and interpretation of mixed effects models. The conditional formulation of this type of models which explicitly includes random effects in the conditional mean allows obtaining not only individual inference, but also population inference. This is due to the simplicity of integrating out the random effects of the model, i. e. moving from the conditional formulation of the model to the marginal formulation (Lee and Nelder, 2004). This feature of Bayesian models is of utmost importance in the case of growth models because it expresses in a natural probabilistic way all information about the parameters and other relevant fea-

tures of the growth process through the respective posterior distribution. Furthermore, model checking can be conducted in a straightforward way to detect systematic bias in the model. This is particularly important for medical applications to avoid patients from receiving a sub-optimal medical treatment.

The current chapter is organised as follows. Section 2.2 presents the data and contains a brief description of them that emphasises the peculiarities of the repeated-means design through the number of observations per individual and their *I*AV trajectories according to the *I*AP values. Section 2.3 introduces and formulates the statistical modelling. Section 2.4 accounts for posterior inferences and prediction. Section 2.4.1 discusses the posterior distribution of the inferential process. Sections 2.4.2 and 2.4.3 contain, respectively, some relevant results of clinical interest at specific individual levels and in general terms for different population groups. Section 2.5 deals with model checking by means of the cross-validated predictive density. The chapter ends with an overview of the results and some conclusions.

2.2 Intra-abdominal volume and intra-abdominal pressure data

The data for the current modelling are obtained from a previously published individual patient meta-analysis (Mazzinari et al., 2021) that included experimental information from three previous clinical studies. The data set includes 198 patients, 118 men and 80 women, and a total of 1 361 observations. The study follows a repeated measures design, with varying numbers of observations per individual. The number of observations per individual ranges from a minimum of one to a maximum of 15. Figure 2.1 illustrates the distribution of repeated measures among men and women, organized by age. Notably, women generally have fewer measurements across all age groups compared to men.

The data have a very wide age range. The youngest patient is 23 years old and the oldest 92, with a mean age of 64.65 years. In the men's group, the minimum and




Figure 2.1: Number of repeated measures in the men’s group (top panel) and in the women’s group (bottom panel). Each bar corresponds to a person and its ordinate describes the number of measurements of that person during the study. Patients are ordered according to their age from the youngest to the oldest.

maximum also are 23 and 92, respectively, and their average is 64.49 years. Women have a minimum age of 34 and a maximum of 85, and their mean is 64.87 years.

IAP values range between 0 and 16 mmHg, and IAV values between 0.5 and 13 L. Figure 2.2 shows a spaghetti plot of IAV , in L, for men and women. Men and women show a fairly similar pattern of the IAV with IAP , although a wider range of values is observed in men, due to higher values of IAP . In both groups there are individuals with different behaviours but men behave more homogeneously than women.

2.3 Logistic growth mixed-effects modelling

Let y_{ij} be the response random variable in the nonlinear mixed-effects model. Variables y_{ij} describe the IAV value for each individual i , $i = 1, \dots, N$, with standardised




Figure 2.2: IAV profiles (in L) according to IAP (in mmHg) for men (top panel) and women (bottom panel) in the sample.

intra-abdominal pressure values IAP_{ij} , $j = 1, \dots, n_i$. Throughout this document, n_i represents the number of measurements for individual i .

The model considered in this study is in line with expression (1.4). Additionally, an important assumption we make in this model is that the observations of each individual are conditionally independent, given θ and ϕ . Thus, the individual likelihood of (θ, ϕ) , $f(\mathbf{y}_i | \theta, \phi)$, is the product of each individual contribution:

$$f(\mathbf{y}_i | \theta, \phi) = \left(\prod_{j=1}^{n_i} f(y_{ij} | \theta, \phi) \right). \quad (2.1)$$

The variables y_{ij} are probabilistically defined in terms of a conditional normal distribution as follows

$$f(y_{ij} | \theta, \phi) = \mathcal{N}(\mu_{ij}, \sigma^2), \quad (2.2)$$

where μ_{ij} is the mean of the response variable y_{ij} . The mean can be expressed in terms

of the conditional logistic growth function

$$(\mu_{ij} \mid a_i, b_i, c_i, IAP_{ij}) = \frac{a_i}{1 + \exp\{-(b_i + c_i IAP_{ij})\}}, \quad (2.3)$$

with parameters a_i , b_i , and c_i determining the growth of the function, and σ^2 the unknown variance associated with the random measurement error of the normal (2.2).

The logistic growth model for μ_{ij} has important features which are very valuable to better understand the relationship between IAP and IAV (Davidian, 2008):

- It is an increasing sigmoid function (see Figure 2.3), or S -curve, whose name comes from its shape and was introduced by the mathematician Pierre-François Verhulst in the 19th century to study the growth of populations in autocatalytic chemical reactions (Cramer, 2004).
- The asymptotic value of IAV when IAP goes to infinity is a_i .
- The inflection point (IP), where the curve changes from being concave downward to concave upward, and therefore it is the point at which the acceleration of the process switches from positive to negative, is $-b_i/c_i$ for IAP . The value of IAV at this point is $a_i/2$.
- The maximum acceleration and deceleration points, MAP and MDP respectively, have the following IAP and IAV coordinates: $((-(\ln(2 + \sqrt{3}) + b_i)/c_i, a_i/(3 - \sqrt{3}))$ and $((-\ln(2 - \sqrt{3}) + b_i)/c_i, a_i/(3 + \sqrt{3}))$.
- The asymptotic deceleration point (ADP) is calculated by equalling the fourth derivative to 0. It is located after the maximum deceleration point, and it indicates the point in which the acceleration is negative but close to 0. Therefore, it is expected that the increase of the function is not of much practical interest. Its value for IAP is $-(\ln(5 - 2\sqrt{6}) + b_i)/c_i$, whereas that for IAV is $a_i(3 + \sqrt{6})/6$.

To illustrate these concept, Figure (2.3) shows the graph of the logistic growth model $y = 5/[1 + \exp\{-(-10 + x)\}]^{-1}$ with generic variables x and y , and the

location on the graph of the special points described above.




Figure 2.3: Graphics of the logistic growth function $5/[1 + \exp\{-(-10 + x)\}]^{-1}$, the subsequent asymptotic value, and its *MAP*, *IP*, *ADP*, and *MDP* points.

Hierarchical modelling for parameters a_i and b_i was based on expert information and connected with two covariates: standardized age and gender. Parameter c_i was associated with the covariate gender only. This choice was made to avoid the interaction between *IAP* and age. Furthermore, a_i and b_i included a random effect specifically associated with each individual that allows connecting all their repeated observations. We have not included any random effect in the modelling of the parameter c_i . This because it would generate a random interaction term with the *IAP* values that would be difficult to understand and justify. Taking these issues into consideration, our model (2.2) and (2.3) is completed as follows

$$a_i = \beta_0^{(a)} + \phi_i^{(a)} + \beta_W^{(a)} I_W(i) + \beta_A^{(a)} Age_i, \quad (2.4)$$

$$b_i = \beta_0^{(b)} + \phi_i^{(b)} + \beta_W^{(b)} I_W(i) + \beta_A^{(b)} Age_i, \quad (2.5)$$

$$c_i = \beta_0^{(c)} + \beta_W^{(c)} I_W(i), \quad (2.6)$$

where $\beta_0 = (\beta_0^{(a)}, \beta_0^{(b)}, \beta_0^{(c)})$ stands for the common intercept with the men group being the reference group; $I_W(i)$ is the indicator variable taking the value of 1 if individual i is a woman and 0 otherwise; $\beta_W = (\beta_W^{(a)}, \beta_W^{(b)}, \beta_W^{(c)})$ and $\beta_A = (\beta_A^{(a)}, \beta_A^{(b)})$ are the vectors of regression coefficients associated with individual i being a woman

and their standardized age, respectively. Note that, in this case, the model just includes random effects, i. e., $\phi = (\phi_1^{(a)}, \dots, \phi_N^{(a)}, \phi_1^{(b)}, \dots, \phi_N^{(b)})$. Random effects $\phi_i^{(a)}$ and $\phi_i^{(b)}$, $i = 1, \dots, n$, are assumed conditional independent given σ_a^2 and σ_b^2 and normally distributed: $(\phi_i^{(a)} | \sigma_a^2) \sim \mathcal{N}(0, \sigma_a^2)$ and $(\phi_i^{(b)} | \sigma_b^2) \sim \mathcal{N}(0, \sigma_b^2)$.

The Bayesian model is completed with the elicitation of a prior distribution for the parameters and hyperparameters $\theta = (\beta_0, \beta_W, \beta_A, \sigma, \sigma_a, \sigma_b)$ of the model. We assume prior independence between them and select the uniform distribution $\mathcal{U}(0, 10)$ as the prior for all standard deviation terms. The elicited marginal prior distribution for $\beta_0^{(a)}$ and $\beta_0^{(c)}$ is $\mathcal{U}(0, 20)$ and $\mathcal{U}(0, 10)$, respectively. These uniform distributions are sufficiently large to cover the whole range of possible values of both parameters. The normal distribution $\mathcal{N}(0, 10^2)$ is selected for $\beta_0^{(b)}$, $\beta_W^{(a)}$, $\beta_W^{(b)}$, $\beta_W^{(c)}$, $\beta_A^{(a)}$, and $\beta_A^{(b)}$ to allow the parameters to move freely between a wide range of positive and negative values.

2.4 Posterior inferences and predictions

2.4.1 Posterior distribution

The relevant quantities in the inferential process are the parametric vector θ and the set of random effects associated with the individuals in the sample $\phi = (\phi_1, \dots, \phi_n)$, where $\phi_i = (\phi_i^{(a)}, \phi_i^{(b)})$. Let $\mathcal{D}$ be the observed *IAV* and *IAP* data of all individuals in the sample as well as their age and gender. The posterior distribution $\pi(\theta, \phi | \mathcal{D})$, contains all of the relevant information of the problem and it is usually the starting point of all relevant inferences. It was approximated by means of MCMC simulation methods through the JAGS software. For the estimated model, we ran three parallel chains with 1,000,000 iterations with a previous burn-in of 500,000. Chains were also thinned by storing every 1,000th iteration to reduce the autocorrelation in the sample. Convergence to the joint posterior distribution was guaranteed by visualising every autocorrelation function plot by means of `mcmcplots` package (Curtis, Goldin, and Evangelou, 2018) for the R software and assuring an effective number of independent

simulation draws greater than 100. For the sake of reproducibility we have generated a synthetic data set, which is available as supplementary material along with the R code for the analysis at https://github.com/gcalvobayarri/intra_abdominal_volume_model.git.

Table 2.1 summarizes $\pi(\boldsymbol{\theta}, \boldsymbol{\phi} \mid \mathcal{D})$. The posterior means of $\beta_0^{(a)}$ and $\beta_0^{(b)}$ provide an approximate overall assessment of the baseline values of a_i and b_i for male patients. Regarding the asymptotic value a_i , it decreases to 0.418 in the female group (although this estimation has a lot of uncertainty), and shows a slight positive trend with age. Differences between individuals are relevant as it can be seen from the estimation of the standard deviation of the random effect in a_i , i. e. 1.67. The parameter b_i has an approximate basal value of 1.08 in the men group, which decreases to -0.27 in the women group. Age also has a positive estimation and the random effect associated with individuals are also important for b_i , especially because this term appears on an exponential scale, with a negative sign in the ratio of the growth curve. Finally, the posterior mean for the c_i term is about 2.26 in the men group and decreases in 0.264 units in the group of women. The posterior mean of the standard deviation associated with the measurement error is not very large but it does have a very high accuracy. Since the *IAP* value of the *IP*, *ADP*, *MAP* and *MDP* of individual i depends on b_i and c_i proportionally to $-b_i/c_i$, and that the estimated coefficient associated with age is positive for b_i , the relationship of the *IP*, *ADP*, *MAP* and *MDP* for *IAP* coordinate with age is negative but barely important.

As we said before, the posterior distribution is the starting point for the analysis of the different outcomes of interest in the study. In the following, we will present different results that may be useful to better understand the relationship between *IAV* and *IAP* both at individual and population level. This will allow us to answer in the next sections the scientific questions raised by the study. But first we would like to make a brief comment on the variable selection process discussed above for parameter c_i of the growth model. In this context, we considered different modelling approaches for c_i with regard to covariate gender. The Deviance Information Criterion (Spiegelhalter

Table 2.1: Posterior summaries (mean, standard deviation and 95% credible interval) for the parameters and hyperparameters of the logistic growth model with covariates gender and standardized age.

Parameters	Logistic growth model		
	mean	sd	$CI_{0.95}$
$\beta_0^{(a)}$	5.729	0.377	(4.968, 6.452)
$\beta_W^{(a)}$	-0.418	0.259	(-0.927, 0.095)
$\beta_A^{(a)}$	0.101	0.124	(-0.145, 0.349)
σ_a	1.670	0.090	(1.501, 1.860)
$\beta_0^{(b)}$	1.080	0.181	(0.730, 1.440)
$\beta_W^{(b)}$	-0.270	0.125	(-0.517, -0.028)
$\beta_A^{(b)}$	0.134	0.054	(0.026, 0.241)
σ_b	0.650	0.041	(0.572, 0.736)
$\beta_0^{(c)}$	2.260	0.120	(2.029, 2.503)
$\beta_W^{(c)}$	-0.264	0.082	(-0.431, -0.101)
σ	0.490	0.011	(0.469, 0.513)

et al., 2002) was used for model comparison and according to this criterion the best model was the one with only the gender covariate and a common population term in the parameter c_i as stated before.

2.4.2 Posterior individual outcomes

The basic inferential process allows the Bayesian methodology to obtain information both individually and in terms of the target population.

In the following, we focus on ADP . The mean of the IAP value of ADP for individual i , ADP_i , depends on b_i and c_i , which in turn depends on (θ, ϕ_i) . As a consequence, we can compute the posterior distribution of the true ADP_i of each individual i in the sample from the subsequent posterior distribution $\pi(\theta, \phi_i \mid \mathcal{D})$. Figure 2.4 shows the posterior mean and the 95% credible interval for that quantity of the individuals in the sample ranked by age. The first thing that is striking in both graphs is the great difference the range of the credibility interval in the group of men and that of women. This is mainly explained by the differences in the number of repeated observations for each of them.




Figure 2.4: Posterior mean and 95% credible interval of the ADP (IAP value) of the men (top panel) and the women (bottom panel) in the sample. Patients are ordered in the x -axis according to their age, from the youngest to the oldest.

Prediction of observations for new individuals in the target population is an important issue that Bayesian statistics approaches in a natural way. The posterior predictive distribution of the random variable $y_{N+1,j}$ depends on the conditional model in (2.2) and the posterior distribution $\pi(\boldsymbol{\theta}, \boldsymbol{\phi}_{N+1} \mid \mathcal{D})$. The random variable $y_{N+1,j}$ records the IAP value for a new individual, $N+1$, of the population based on their IAP , age, and gender values denoted as $\mathbf{X}_{N+1,j}$.

The posterior predictive distribution is computed as follows:

$$f(y_{N+1,j} \mid \mathbf{X}_{N+1,j}, \mathcal{D}) = \int f(y_{N+1,j} \mid \mathbf{X}_{N+1,j}, \boldsymbol{\theta}, \boldsymbol{\phi}_{N+1}) \pi(\boldsymbol{\theta}, \boldsymbol{\phi}_{N+1} \mid \mathcal{D}) d(\boldsymbol{\theta}, \boldsymbol{\phi}_{N+1}). \quad (2.7)$$

The posterior distribution $\pi(\boldsymbol{\theta}, \boldsymbol{\phi}_{N+1} \mid \mathcal{D})$ factorizes in terms of the marginal posterior distribution $\pi(\boldsymbol{\theta} \mid \mathcal{D})$ and the conditional distributions for the random effects $(\phi_{N+1}^{(a)} \mid \sigma_a^2) \sim N(0, \sigma_a^2)$ and $(\phi_{N+1}^{(b)} \mid \sigma_b^2) \sim N(0, \sigma_b^2)$. Figure 2.5 shows the posterior




Figure 2.5: Posterior predictive mean and 95% predictive interval for IAV in relation to IAP values. The top panel represents a man, while the bottom panel represents a woman. Both individuals are 64.65 years old, which corresponds to the sample mean.

predictive mean and the 95% predictive interval for the IAV value of a new individual of the target population with respect to their IAP and in relation to their gender. Both groups behave similarly. The stabilisation of the values of IAV in both groups can be clearly seen, as well as the variability associated with the predictive processes, which is always greater in comparison with the estimation processes themselves. Finally, Figure 2.6 shows the posterior predictive mean for the response IAV variable with regard to the IAP values of individuals with different ages and gender. Of course, as we observed with the approximate posterior distribution of $\beta_A^{(b)}$ in Table 2.1, a positive relationship between IAV and age can be observed in the graphic but very mild, and possibly not very relevant for practical purposes.

2.4.3 Posterior population outcomes

Random effects connect the different repeated measures of the same individual in the statistical model and allow computation of individual-specific outcomes. We would




Figure 2.6: Posterior predictive mean of the IAV with regard to IAP values for a man (top panel) and a woman (bottom panel) stratified by age.

also like to be able to have not only that individual information, but also outcomes that can provide general information about the target population. This aim implies to work with the marginal formulation of the model in (2.2) and (2.3) which we would obtain by integrating out the random effects as follows

$$f(y_{ij} | \mathbf{X}_{ij}, \boldsymbol{\theta}) = \int f(y_{ij} | \mathbf{X}_{ij}, \boldsymbol{\theta}, \phi) f(\phi | \boldsymbol{\theta}) d\phi. \quad (2.8)$$

This marginal formulation depends only on the parameter and hyperparameters of the model $\boldsymbol{\theta}$ and is the basis for the computation of any feature of this marginal model. For simplicity, we only focus on the true asymptotic IAV value and the true asymptotic deceleration point ADP for a patient with an average age.

Figure 2.7 shows the posterior distribution of the asymptotic IAV for men and women aged 64.65 years (the mean of the sample). There is not much difference between the two distributions. An estimation of the asymptotic IAV in the group of men is 5.60 L, whereas in the group of women it is 5.25 L. Figure 2.8 shows the joint posterior distribution (in terms of contour lines) of the ADP pressure point and the




Figure 2.7: Posterior distribution of the asymptotic IAV for men (left) and women (right).

subsequent volume value for men and women aged 64.65 years (the sample mean) as well as the marginal distributions of both quantities. Posterior mean for the ADP 's pressure and volume is 10.06 mmHg and 5.05 L in men aged 64.65, and 10.87 mmHg and 4.74 L in the group of women with the same age, respectively. A similar analysis is possible for MAP , IP and MDP . However, their posterior results for both coordinates (IAP and IAV) are proportional to those of ADP and their information would be repetitive.

2.5 Model checking

Model checking is an essential component of any statistical analysis which has generated an extensive literature within the Bayesian reasoning (Vehtari and Ojanen, 2012). Our interest in this subject focuses on assessing (following the philosophy in Gelman et al. (2014)) whether the possible shortcomings of our model have a relevant effect on the derived results. We approach model checking via posterior predictive distributions following the ideas by Box (1980), who states that prediction (and not estimation) enables “criticism of the entertained model in light of current data”.




Figure 2.8: Joint posterior distribution and contour lines of the IAP and IAV coordinates for the true ADP and posterior marginal distribution for each of both quantities for men (top panel) and women (bottom panel) aged 64.65 years (the sample mean). The horizontal and vertical lines represent the approximate posterior mean of $IAV - ADP$ and the approximate posterior mean of $IAP - ADP$, respectively.

We check our model through the cross-validated predictive density (Gelfand, Dey, and Chang, 1992) defined as the conditional posterior density of a future IAV value $y_{ij}^{(pred)}$

$$f(y_{ij}^{(pred)} | \mathcal{D}^{-(ij)}) = \int f(y_{ij}^{(pred)} | \mathbf{X}_{ij}, \boldsymbol{\theta}, \boldsymbol{\phi}) \pi(\boldsymbol{\theta}, \boldsymbol{\phi} | \mathcal{D}^{-(ij)}) d(\boldsymbol{\theta}, \boldsymbol{\phi}),$$

where $\mathcal{D}^{-(ij)}$ are all data in $\mathcal{D}$ except for the observation y_{ij} (leave-one-out procedure).

The fundamental idea underlying this proposal assumes that if the estimated model

is correct, each observation can be considered as a random value from the cross-validated predictive density (Chen, Shao, and Ibrahim, 2012). In this framework, we consider two complementary characteristics of such predictive distribution assessed at each observed value y_{ij} . These quantities are the conditional predictive ordinate (*CPO*) and the cross-validated probability integral (*PIT*), and are defined as:

$$CPO_{ij} = f(y_{ij} \mid \mathcal{D}^{-(ij)})$$

$$PIT_{ij} = P(Y_{ij} \leq y_{ij} \mid \mathcal{D}^{-(ij)}).$$

CPO_{ij} values correspond to the ordinates in the y_{ij} of the cross-validated predictive density. Large CPO_{ij} values support the selected model because indicate a good tuning between the data and the model. PIT_{ij} is the posterior probability that the replicated (ij) -th observation is less or equal the subsequent observed value. When the model is well calibrated these probabilities follow a uniform distribution $\mathcal{U}(0, 1)$.

The direct implementation of these quantities is computationally very expensive because we need to approximate as many posterior distributions as the elements in $\mathcal{D}$. In practice, the application of self-normalized importance sampling allows approximating CPO 's and PIT 's from draws of the posterior distribution $\pi(\boldsymbol{\theta}, \boldsymbol{\phi} \mid \mathcal{D})$ computed through the complete data $\mathcal{D}$ (Gelfand, 1996; Ntzoufras, 2011). In our application, the computation of CPO 's and PIT 's was done by means of the R software from the posterior outputs obtained with JAGS.

Figures 2.9 and 2.10 show the histograms of CPO 's and PIT 's respectively. The information provided in both cases suggests that the model can be improved. Most of the approximate CPO 's are small, near 0. Furthermore, the distribution of the PIT values is not uniform, the values accumulate close to 0.

Figure 2.11 shows how PIT values are distributed along IAP . Theoretically, PIT 's should be uniformly distributed at each IAP point. However, from $IAP = 6.5$ (vertical red line) PIT tends to accumulate in 0. This indicates that this model overpredicts observations of IAP for medium and high values of the covariate IAP .




Figure 2.9: Histogram of the approximate CPO quantities for all the observations.




Figure 2.10: Density histogram of the approximate PIT quantities for all the observations.




Figure 2.11: Distribution of PIT values along the different values of IAP .

2.6 Conclusions

Precision medicine tenets state that different interventions have distinct effects in different people and that this variability can, at least in part, be characterised and predicted (Senn, 2016). In this study we tried to lay the foundation for the mathematical modelling of the abdomen behaviour during pneumoperitoneum insufflation. We have parameterised such model to achieve predictive capability based on a few simple baseline characteristics. This is the first step in a precision medicine approach to pneumoperitoneum insufflation for laparoscopic surgery. This process can be potentially scaled up and recursively performed throughout the duration of the surgical intervention to ensure that even if conditions change, we could be able to provide an optimal surgical field to the surgeon while exposing the patient the lowest possible pressure.

With this procedure, we would like to achieve an optimal surgical workspace while minimising the pressure administered to the patient. In other words, each subject would receive a titrated pressure according to his characteristics. Also, the ability to predict where the marginal gain in volume diminishes by deriving critical points on the parameterized curve have an especially interesting clinical potential.

Bayesian inference can provide a suitable inferential framework in this context. First of all, Bayesian hierarchical models are useful to elicit and formulate the different sources of variation and uncertainty of the problem and incorporate suitable terms into the model to account for them. In this particular case, the model includes nonlinear effects through a logistic growth function. As model fitting relies on MCMC methods, inference about particular elements of interest in the model becomes feasible. For example, the logistic growth model has a known parametric form from which some crucial critical points can be derived analytically but inference on these points is far from straightforward. However, the output produced by MCMC during model fitting can be exploited to compute the posterior marginals of these particular points as well as those of the other model parameters. This provides extra information that can be used during the laparoscopic surgery. Inference about these critical points under other

inferential frameworks would not be so straightforward.

The most important critical point in our study is *ADP* as this controls how much air is insufflated during surgery. From a clinical point of view, when operating on new patients, *ADP*'s predictive distribution can help physicians to provide adequate insufflation during laparoscopic surgery.

As illustrated, model checking is critical for health applications of statistical methodology as this will allow detecting potential bias in the proposed model. We believe that the use of the model checking techniques should be widely adopted when relying on statistical models for medical practice to detect and avoid systematic biases in medical treatment (Obermeyer et al., 2019).

The study presented here illustrates a preliminary analysis in which 198 patients have been enrolled. In the future, we aim to conduct a larger trial so that a wider range of patients is represented. Furthermore, other covariates will be recorded and included into the model to reduce the uncertainty about the estimates and predictions, and increase the accuracy of insufflation. We plan to refine our model with anthropometric measurements. We are recording not only height and weight, but also waist and hip circumference, and abdomen sagittal height to update our model with this new data. Furthermore, data from medical imaging such as abdominal volume estimation based on routine preoperative computerized tomography images or ultrasonic assessment on the abdominal wall thickness and fat component can be explored as alternative covariates. Finally, models with different assumptions will be considered such as non-homoscedastic models with increasing variability, different types of curves such as the Gompertz curve (Funatogawa and Funatogawa, 2018), or even the inclusion of random effects to assess the possible variability among the different studies.

Chapter 3

Bayesian joint longitudinal models for assessing the exploitation rates of sardine stock in the Mediterranean Sea

3.1 Introduction

Small pelagic fish species are key elements of the Mediterranean pelagic ecosystem due to their high bulk biomass at the mid-trophic level, which provides an important energy connection between the lower and the upper trophic levels (Albo-Puigserver et al., 2015). They are shown to be essential in the coupling between the pelagic and the demersal environment, as they are prey for pelagic predators such as tuna, cetaceans, pelagic birds, and demersal predators such as hakes (Mellon-Duval et al., 2017; Navarro et al., 2017). Small pelagic fish usually live in dense shoals, making gear such as mid-water pelagic trawls and purse seines particularly efficient for their capture.

Fluctuations in populations of small pelagics can therefore have serious ecological

and socio-economic consequences. The population dynamics of these species could be strongly influenced by natural environmental fluctuations (bottom-up control) and mortality (top-down control) (Checkley Jr, Asch, and Rykaczewski, 2017). Catches in the Mediterranean Sea are dominated by small pelagics, representing nearly 49% of the harvest (Saraux et al., 2014). Among them, the European sardine (*Sardina pilchardus*) is possibly the most important species both ecologically and economically.

In recent years, the European sardine has undergone a shocking process of fragility, with a drastic reduction in body weight (about two thirds) and a much shorter life span than 20 years ago, which has been reflected in a consequent reduction in catches. This decline is the result of a complex process probably caused by many factors, the combination of which has led to this negative process (Brosset et al., 2017; Brigolin et al., 2018). Many theories were raised within the fishery scientific community: an increase in hungry sardine-eating predators such as bluefin tuna, whose population is increasing in the Mediterranean Sea, and dolphins. Another important element is the reduction of plankton, their food source. In the last few years, these marine micro-organisms have been drastically reduced and have become less nutritious. This is a result of global warming, as plankton generally feed on nutrients from cold, deep waters. In addition, there has also been a recent proliferation of jellyfish blooms, predators of eggs and larval stages of sardines and also competitors for plankton. The shrinking size of sardines, the fact that most sardines are sold fresh rather than frozen, the changing economic relationship that generates their sales, dominated lately by large retail chains that virtually control their demand, and changing consumer preferences and intergenerational differences have all contributed, to a greater or lesser extent, to a fall in sardine prices and catches in the market that threatens the viability of the less powerful fishing fleets.

Figure 3.1 describes the quantity, in thousands of tonnes (tt), of sardine caught by the main Mediterranean fishing countries from 1985 to 2018. This graph shows very interesting information. Two distinct behaviours can be reported. On the one hand, we observe that from 1985 to around 2005 the quantity of sardines fished experienced a

notable decrease, going from around 350 to 250 tt. In 2006 and 2007 we can appreciate a slight upturn, but from that date onward there is again a decrease that seems to have stabilised in recent years at around 225 tt.




Figure 3.1: Annual thousands of tons of European sardines caught by the main Mediterranean fishing countries from 1985 to 2018. Data come from Sea Around Us (www.seaaroundus.org).

At regional level, strong declines in sardine landings have been reported in different Mediterranean areas (Coll et al., 2019). However, no study has been performed to assess sardine landings dynamics at a large-scale. Although working at a fine scale and regional level (e.g. Mediterranean eco-regions or subareas) has several advantages, large-scale analysis is also essential to provide a broader and complementary perspective of non-localised drivers of fisheries where sardine fish stocks are shared among countries, and reduce the bias of landings misallocation (Stergiou et al., 2016). Understanding fish population dynamics is a major goal of fisheries ecology.

Pennino et al. (2017) stated that Mediterranean fisheries are highly diverse and geographically varied, not only because of the existence of different marine environments, but also because of diverse socio-economic situations, and fisheries status. The two most important types of commercial fishing in the Mediterranean are artisanal and industrial, defined in terms of small-scale and large-scale commercial fisheries,

respectively. Artisanal fishing is carried out by small boats with a low level of mechanisation owned by family businesses, working close to the coast, and whose products are mainly intended to supply the local market. The European Union (EU) defines this type of fishing as that carried out by vessels of less than 12m in length and with certain fishing gear. All fishing carried out by vessels of greater length or with gears not considered as artisanal by the EU would fall under the denomination of industrial fishing. Although industrial fishing is the most abundant, artisanal fishing is not only a cultural heritage but also the livelihood of many families.

This chapter deals with a study that focuses on the evolution of sardine catches in the Mediterranean Sea from 1985 to 2018 according to the Mediterranean fishing country and the type of fishing practised. The main objective is to detect and analyse general patterns of the joint temporal evolution of artisanal and industrial fisheries, as well as relevant individual characteristics of the different Mediterranean countries. This is a quantitative overview that provides a different perspective on the problem than a purely biological, economic or ecological one.

The statistical framework of this study are the longitudinal linear mixed models. Each Mediterranean country was considered as an individual and the annual sardine landings formed a set of individual time series (1985 to 2018), one for each country and type of fishing. The assumption that all of the information about the response variable is generated by a function common to all individuals of the target population is not realistic. Ignoring individual heterogeneity among countries could lead to inconsistent and inefficient estimates of the parameters of interest (Pinheiro and Bates, 2000). As a result, our longitudinal models will include not only common population information and measurement error terms, as described by Diggle et al. (2002), but also random effects to account for specific country characteristics.

The contributions of this chapter are twofold. On the methodological side, we developed an approach in the fishery framework for jointly modelling several longitudinal response variables. This approach connects individual information by considering correlation parameters that measure the linear relationships between the random

effects of the different responses within each individual. Secondly, on the applied fishery side, we demonstrated, among all the heterogeneous behaviours of the countries, a decline in landing quantity for both the artisanal and the industrial fishing in the Mediterranean Sea over the past three decades. In addition, we found a positive synergistic relationship between the two types of fisheries within each country.

This chapter continues with Section 3.2 that presents the data on artisanal and industrial sardine catches from 1985 to 2018 corresponding to the Mediterranean countries that participate in the study. Section 3.3 discusses the general joint Bayesian statistical framework and introduces three modelling approaches that seem appropriate for analysing the study's data. Section 3.4 deals with the approximation of the subsequent posterior distributions via MCMC. It also presents a discussion comparing the three proposed models, and a final subsection that analyses the possible synergistic relationships between the two types of fisheries in the different countries studied. The chapter ends with some concluding remarks.

3.2 Sardine fisheries in the Mediterranean Sea

Annual landings data (tonnes) from 1985 to 2018 of European sardine caught by artisanal and industrial methods by Mediterranean countries have been extracted from Sea Around Us (www.seaaroundus.org). Countries participating in the study are Albania, Algeria, Bosnia and Herzegovina, Croatia, France, Greece, Italy, Montenegro, Morocco, Slovenia, Spain and Turkey (See Figure 3.2).

Data from European countries recognised as sovereign from 2010 by the international community (Bosnia and Herzegovina, Croatia, Montenegro, and Slovenia) were imputed from 1985-2010 (Zeller and Pauly, 2016) based on information from Exclusive Economic Zones (EEZ), which was linked to these countries after reconstructing data by means of several sources.

Figure 3.3 shows two spaghetti plots. The top figure shows the annual amount of sardines caught by artisanal fishing from 1985 to 2018 in those Mediterranean coun-



Figure 3.2: Countries bordering the Mediterranean Sea. Photo by Ian Macky, PAT Atlas, <https://ian.macky.net/pat/map/medi/mediblu.gif>

tries that retain this type of fishery. The bottom figure displays the same quantity but for the annual amount of sardines caught by industrial methods. The first thing that strikes is how few sardines are caught by traditional methods compared to those caught by industrial fishing. There is only one notable situation, such as that of Algeria, which has a positive evolution of its traditional fisheries. Note that not all countries have both types of fisheries (Bosnia and Herzegovina has no registered industrial fisheries and Albania has not registered artisanal fisheries).

The behaviour of the countries is very different both in relation to the quantity of sardines caught by industrial methods at the beginning of the study (year 1985) and in their evolution over time. Countries such as Italy and Algeria are particularly striking. Italy started out as the country that fished the most sardines at the beginning of the study, exceeding the 150 tonne barrier the first years. From the mid-1980s onwards, it experienced a significant decrease, fishing less than 50 tonnes per year from 1998 onwards. Although it seems to have recovered in recent years, it is still below the annual level of 50 tonnes of sardines fished. Algeria started the study in the group of low fishing countries but experienced strong growth until the end of the




Figure 3.3: Annual tons of European sardine (*Sardina pilchardus*) per country, from 1985 to 2018. From top to bottom: artisanal and industrial and fishing.

1980s. Since then it has maintained a fairly stable behaviour until around 2008 (time of the great financial crisis) when it has seen a drastic reduction in its sardine catches. Spain has also exhibited a consistent decline in industrial fishing activity since around 1995, without showing any signs of recovery since then. A weak recovery seems to have been observed in some countries in recent years. Bosnia and Herzegovina is the country that fishes the least quantity of sardines. It only reports industrial fishing but has experienced a notable expansion in the most recent years. Turkey also increased fishing significantly until 1995, and it is currently, along with Algeria, Croatia, Greece and Morocco one of the largest sardine fisheries. France and Slovenia have lowered

substantially their catches in the recent years.

3.3 Joint Bayesian longitudinal modelling

Let $(y_{ij}^{(1)}, y_{ij}^{(2)})$ be the bivariate longitudinal random vector describing the response of individual i , $i = 1, \dots, N$ at measurement j , where $j = 1, 2, \dots, n_i$. Let t_{ij} be the time covariate for individual i at measurement j . Consider the random vector $(\mathbf{y}^{(1)}, \mathbf{y}^{(2)}) = \{(\mathbf{y}_i^{(1)}, \mathbf{y}_i^{(2)}), i = 1, \dots, N\}$, where $(\mathbf{y}_i^{(1)}, \mathbf{y}_i^{(2)}) = \{(y_{ij}^{(1)}, y_{ij}^{(2)}), j = 1, \dots, n_i\}$.

Bayesian joint longitudinal models (BJLMs) for the random vector $(\mathbf{y}^{(1)}, \mathbf{y}^{(2)})$ are specified via the joint probability distribution from expression (1.4). Note that in this case the number of response variables is $Q = 2$:

$$f(\mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\phi}) = \left(\prod_{i=1}^N f(\mathbf{y}_i^{(1)}, \mathbf{y}_i^{(2)} \mid \boldsymbol{\theta}, \boldsymbol{\phi}) \right) f(\boldsymbol{\phi} \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta}).$$

An additional feature of these models is the assumption of independency between the two response variables given the parameters, hyperparameters, random effects and serial correlation terms:

$$f(\mathbf{y}_i^{(1)}, \mathbf{y}_i^{(2)} \mid \boldsymbol{\theta}, \boldsymbol{\phi}) = f(\mathbf{y}_i^{(1)} \mid \boldsymbol{\theta}, \boldsymbol{\phi}_i) f(\mathbf{y}_i^{(2)} \mid \boldsymbol{\theta}, \boldsymbol{\phi}_i). \quad (3.1)$$

This BJLM works in the framework of the normal distribution, so the i -th component of the sampling probability model is:

$$\begin{aligned} f(\mathbf{y}_i^{(1)} \mid \boldsymbol{\theta}, \boldsymbol{\phi}_i) &= \mathcal{N}(\boldsymbol{\mu}_i^{(1)}, \boldsymbol{\Sigma}_i^{(1)}) \\ f(\mathbf{y}_i^{(2)} \mid \boldsymbol{\theta}, \boldsymbol{\phi}_i) &= \mathcal{N}(\boldsymbol{\mu}_i^{(2)}, \boldsymbol{\Sigma}_i^{(2)}), \end{aligned} \quad (3.2)$$

where $\boldsymbol{\mu}_i^{(\cdot)}$ is the conditional mean of $\mathbf{y}_i^{(\cdot)}$ given $(\boldsymbol{\theta}, \boldsymbol{\phi}_i)$ and $\boldsymbol{\Sigma}_i^{(\cdot)}$ its conditional variance-covariance matrix whose (j, l) th element is $\text{Cov}(y_{ij}^{(\cdot)}, y_{il}^{(\cdot)} \mid \boldsymbol{\theta}, \boldsymbol{\phi})$. As in Verbeke and Davidian (2009) and Armero et al. (2018), we assume a Bayesian share-parameter

framework to jointly model both responses that uses random effects to generate an association structure between both longitudinal measures.

3.3.1 Longitudinal modelling for the total European sardine landings

Let $y_{ij}^{(A)}$ represent the total tonnage of sardines caught in country i during year t_{ij} using artisanal methods, and let $y_{ij}^{(I)}$ represent the total tonnage of sardines caught using industrial methods. Both quantities are measured on a logarithmic scale. This study considers a time scale based on the calendar, with $t_i = (0, \dots, 34)$ denoting the range of years. Time zero ($t_{i1} = 0$) corresponds to the initial year of the study, specifically, 1985.

We will consider three BLJM (i.e., models $\mathcal{M}_1$, $\mathcal{M}_2$, and $\mathcal{M}_3$) for analysing the evolution of artisanal and industrial fisheries within the general methodological framework explained in the previous section. Model $\mathcal{M}_1$ is a basic homoscedastic linear mixed model with the natural covariate time and random effects. Models $\mathcal{M}_2$ and $\mathcal{M}_3$ account for serial correlation in terms of autoregressive elements, whose inclusion is justified by a strong dependence of each year's fishing activity on previous years and the large number of observations per country. In particular, $\mathcal{M}_2$ and $\mathcal{M}_3$ introduce autoregressive terms in the conditional means and the variance-covariance matrices of the sampling probability distribution in (3.2), respectively. The models are defined as follows:

Model $\mathcal{M}_1$

$$\begin{aligned} f(\mathbf{y}_i^{(A)} \mid \boldsymbol{\theta}, \boldsymbol{\phi}_i) &= \mathcal{N}(\mathbf{1} \beta_0^{(A)} + \mathbf{1} b_{0i}^{(A)} + \mathbf{t} \beta_1^{(A)} + \mathbf{t} b_{1i}^{(A)}, \sigma^{(A)^2} \mathbf{I}), \\ f(\mathbf{y}_i^{(I)} \mid \boldsymbol{\theta}, \boldsymbol{\phi}_i) &= \mathcal{N}(\mathbf{1} \beta_0^{(I)} + \mathbf{1} b_{0i}^{(I)} + \mathbf{t} \beta_1^{(I)} + \mathbf{t} b_{1i}^{(I)}, \sigma^{(I)^2} \mathbf{I}), \end{aligned} \quad (3.3)$$

where $\mathbf{1}$ is a vector of ones of dimension 35, $\mathbf{t} = (0, 1, \dots, 34)$, $\beta_0^{(A)}$ and $\beta_0^{(I)}$ the common intercepts for artisanal and industrial fishing, respectively as well as $\beta_1^{(A)}$ and $\beta_1^{(I)}$ for the subsequent common slopes. Individual random intercepts $b_{0i}^{(A)}$ and $b_{0i}^{(I)}$, and random slopes $b_{1i}^{(A)}$ and $b_{1i}^{(I)}$ are conditionally independent and normally distributed as

$(b_{0i}^{(A)}, b_{0i}^{(I)} \mid \Sigma_0) \sim \mathcal{N}(0, \Sigma_0)$ and $(b_{1i}^{(A)}, b_{1i}^{(I)} \mid \Sigma_1) \sim \mathcal{N}(0, \Sigma_1)$. Variance-covariance matrix Σ_0 includes the correlation coefficient ρ_0 between $b_{0i}^{(A)}$ and $b_{0i}^{(I)}$ as well as its respective variances $\sigma_0^{(A)^2}$ and $\sigma_0^{(I)^2}$. Analogously, ρ_1 and both variances $\sigma_1^{(A)^2}$ and $\sigma_1^{(I)^2}$ for Σ_1 . This model is homoscedastic. Therefore the conditional variability associated with the random measurements is always equal and constant, $\sigma^{(A)^2} \mathbf{I}$ and $\sigma^{(I)^2} \mathbf{I}$ for artisanal and industrial, respectively, where $\mathbf{I}$ represents the 35×35 diagonal unit matrix.

Model $\mathcal{M}_1$ is completed with the elicitation of a prior distribution for the parameters and hyperparameters of the model. We assume a vague informative prior scenario, which gives the maximum prominence to the data, as well as prior independence. We specify normal distributions with large standard deviation for the common regression coefficients, $\pi(\beta_0) = \pi(\beta_1) = \mathcal{N}(0, 10^2)$ and uniform distribution $\mathcal{U}(0, 100)$ for all standard deviation terms, $\sigma^{(A)}, \sigma^{(I)}, \sigma_0^{(A)}, \sigma_0^{(I)}, \sigma_1^{(A)}$, and $\sigma_1^{(I)}$. Prior distributions for correlation terms ρ_0 and ρ_1 are chosen as uniform $\mathcal{U}(-1, 1)$.

Model $\mathcal{M}_2$

$$\begin{aligned} f(\mathbf{y}_i^{(A)} \mid \boldsymbol{\theta}, \boldsymbol{\phi}_i) &= \mathcal{N}(\mathbf{1} \beta_0^{(A)} + \mathbf{1} b_{0i}^{(A)} + \mathbf{t} \beta_1^{(A)} + \mathbf{t} b_{1i}^{(A)}, \boldsymbol{\Sigma}^{(A)}), \\ f(\mathbf{y}_i^{(I)} \mid \boldsymbol{\theta}, \boldsymbol{\phi}_i) &= \mathcal{N}(\mathbf{1} \beta_0^{(I)} + \mathbf{1} b_{0i}^{(I)} + \mathbf{t} \beta_1^{(I)} + \mathbf{t} b_{1i}^{(I)}, \boldsymbol{\Sigma}^{(I)}), \end{aligned} \quad (3.4)$$

In this model, the means of both conditional distributions are expressed as in model $\mathcal{M}_1$. The element (j, l) of the conditional variance-covariance matrix $\Sigma^{(\cdot)}$ is defined as follows:

$$\text{Cov}(y_{ij}^{(\cdot)}, y_{il}^{(\cdot)} \mid \boldsymbol{\theta}, \boldsymbol{\phi}_i) = \sigma^{(\cdot)^2} \sum_{k=1}^{\min(j,l)} \eta^{(\cdot)^{(t_{ij}-t_{ik})+(t_{il}-t_{ik})}}, \quad (3.5)$$

where $\eta^{(\cdot)}$ is the autoregressive coefficient. We chose this specific matrix, instead of considering matrices as described in Hedeker and Gibbons (2006), to avoid potential confusion between the random intercept and the errors at the initial time. Additionally, we have selected a uniform prior distribution, $\mathcal{U}(-1, 1)$, for the autoregressive coefficient.

cients $\eta^{(A)}$ and $\eta^{(I)}$. Thus, $\mathcal{M}_2$ is heteroscedastic at the time level, but by enforcing $|\eta^{(A)}|$ and $|\eta^{(I)}|$ to be less than 1, we place limits on the values of the elements in the variance-covariance matrix. Appendix B.1 provides more details on the derivation of (3.5) and the limit of its values. The prior distribution for the remaining parameters and hyperparameters in this model follows the same approach as model $\mathcal{M}_1$.

Model $\mathcal{M}_3$

$$\begin{aligned} f(\mathbf{y}_i^{(A)} \mid \boldsymbol{\theta}, \phi_i) &= \mathcal{N}(\mathbf{1} \beta_0^{(A)} + \mathbf{1} b_{0i}^{(A)} + \mathbf{t} \beta_1^{(A)} + \mathbf{t} b_{1i}^{(A)} + \mathbf{w}_i^{(A)}, \sigma^{(A)^2} \mathbf{I}), \\ f(\mathbf{y}_i^{(I)} \mid \boldsymbol{\theta}, \phi_i) &= \mathcal{N}(\mathbf{1} \beta_0^{(I)} + \mathbf{1} b_{0i}^{(I)} + \mathbf{t} \beta_1^{(I)} + \mathbf{t} b_{1i}^{(I)} + \mathbf{w}_i^{(I)}, \sigma^{(I)^2} \mathbf{I}). \end{aligned} \quad (3.6)$$

This model is similar to model $\mathcal{M}_1$ but includes a latent autoregressive term $\mathbf{w}_i^{(\cdot)} = (w_i(t_{i1})^{(\cdot)}, \dots, w_i(t_{in_i})^{(\cdot)})$ in the conditional mean, where each $w_i(t_{ij})^{(\cdot)}$, $j = 2, \dots, n_i$, is a realisation at time t_{ij} from a Gaussian process with mean $\eta^{(\cdot)} w_i(t_{ij-1})^{(\cdot)}$ and variance σ_w^2 , where $(w_i(t_{i1})^{(\cdot)} \mid \sigma_w^{(\cdot)}) \sim \mathcal{N}(0, \sigma_w^2)$ (Diggle et al., 2002). That is, $\mathbf{w}_i^{(\cdot)}$ is a vector of time correlated noises. In fact, the variance-covariance matrix of $\mathbf{w}_i^{(\cdot)}$ given $\boldsymbol{\theta}$ is that described in (3.5) of model $\mathcal{M}_2$. Again, the prior distribution for the parameters and hyperparameters of this model is as in model $\mathcal{M}_1$ to which we add the marginal prior distribution, $\mathcal{U}(-1, 1)$ for the autoregressive coefficients $\eta^{(A)}$ and $\eta^{(I)}$, and $\mathcal{U}(0, 100)$ for the standard deviations, $\sigma_w^{(A)}$ and $\sigma_w^{(I)}$, of the autoregressive term in the conditional means.

3.4 Posterior inferences

The posterior distribution associated with models $\mathcal{M}_1$, $\mathcal{M}_2$ and $\mathcal{M}_3$ was approximated via MCMC sampling through JAGS software (version 4.0.5). In particular, we used the `runjags` package developed by Denwood (2016) to parallelise the simulation process, running three parallel chains for each model with a total of 5,000,000 iterations and a burn-in of 1,000,000 iterations. To reduce autocorrelation in the sample, we also thinned the chains by storing every 5,000th iteration. The computations were per-

formed on a computer with an Intel i9 processor, 32GB of RAM, and running Windows 10 Enterprise LTSC. The full inferential analysis, performed by an R code, and the data are publicly available as supplementary material in the GitHub repository at https://github.com/gcalvobayarri/European_sardine_analysis.git.

Table 3.1 shows a basic description of the approximated posterior distribution of the parameters and hyperparameters of the models under study. The first thing that strikes us is the consistency of the results obtained with the three models. All of them present similar results for their outcomes in relation to the initial behaviour of the different countries, their time trends, and the strong impact of the autoregressive terms in models $\mathcal{M}_2$ and $\mathcal{M}_3$. Very few differences can be observed in the estimates of the regression coefficients in both types of fisheries. In all of them, a smooth but negative time trend is discernible which indicates a decreasing pattern in the amount of sardine fishing in recent decades. This is a consistent situation with the behaviour shown by the data in Figure 3.1.

Posterior estimation of the variances associated with the random effects of the intercept and the slope of the artisanal and industrial fisheries are also very similar, mainly for models $\mathcal{M}_2$ and $\mathcal{M}_3$. It is worth noting the great heterogeneity shown by the different Mediterranean countries at the beginning of the study in both types of fisheries as well as the fact that the greatest heterogeneity is shown by the artisanal fishery, which moves in much smaller dimensions than the industrial fishery. The association between the amount of artisanal and industrial fishing in the different countries, expressed through the information on ρ_0 and ρ_1 , is not very strong either at the beginning or during the study period, but the results indicate a positive association between them. This is an important topic that will be fully described later.

The correlation $\eta^{(\cdot)}$ parameters involved in the autoregressive terms of models $\mathcal{M}_2$ and $\mathcal{M}_3$ are not comparable. In the case of $\mathcal{M}_2$ they are associated with the variance-covariance matrices of the sampling model while in $\mathcal{M}_3$ they are on the conditional mean. In both cases, they play a relevant role in the estimated model, with positive posterior means very close to one: 0.913 and 0.938 in $\mathcal{M}_2$, and 0.953 and 0.947 in

Table 3.1: Summary of the approximated posterior distribution of the parameters and hyperparameters in models $\mathcal{M}_1$, $\mathcal{M}_2$ and $\mathcal{M}_3$.

	$\mathcal{M}_1$			$\mathcal{M}_2$			$\mathcal{M}_3$		
	Mean	Sd	0.95CI	Mean	Sd	0.95CI	Mean	Sd	0.95CI
$\beta_0^{(A)}$	5.759	1.253	(3.197,8.239)	5.365	1.340	(2.379, 7.794)	5.375	1.169	(2.936, 7.590)
$\beta_1^{(A)}$	-0.022	0.026	(-0.074,0.031)	-0.029	0.031	(-0.090,0.031)	-0.029	0.031	(-0.090, 0.031)
$\sigma_0^{(A)}$	3.837	1.144	(2.333,6.780)	3.848	1.198	(2.258,6.969)	3.781	1.075	(2.269, 6.434)
$\sigma_1^{(A)}$	0.083	0.026	(0.049,0.144)	0.081	0.034	(0.019,0.158)	0.080	0.036	(0.016, 0.162)
$\sigma^{(A)}$	0.662	0.027	(0.612,0.718)	0.522	0.022	(0.480,0.567)	0.181	0.049	(0.063, 0.0.259)
$\eta^{(A)}$	-	-	-	0.913	0.053	(0.804,0.994)	0.953	0.038	(0.861, 0.998)
$\sigma_w^{(A)}$	-	-	-	-	-	-	0.455	0.036	(0.389, 0.528)
$\beta_0^{(I)}$	9.249	0.768	(7.668,10.777)	9.019	0.783	(7.215,10.363)	9.105	0.662	(7.652,10.291)
$\beta_1^{(I)}$	-0.032	0.024	(-0.079,0.017)	-0.033	0.025	(-0.081, 0.018)	-0.035	0.023	(0.080,0.010)
$\sigma_0^{(I)}$	2.398	0.692	(1.488,4.111)	2.198	0.717	(1.294,3.927)	2.109	0.603	(1.285,3.620)
$\sigma_1^{(I)}$	0.074	0.022	(0.045,0.127)	0.070	0.025	(0.032,0.130)	0.070	0.025	(0.030,0.130)
$\sigma^{(I)}$	0.486	0.018	(0.453,0.524)	0.341	0.013	(0.317,0.368)	0.064	0.037	(0.003,0.136)
$\eta^{(I)}$	-	-	-	0.938	0.037	(0.862, 0.997)	0.947	0.035	(0.872, 0.998)
$\sigma_w^{(I)}$	-	-	-	-	-	-	0.327	0.019	(0.288, 0.362)
ρ_0	0.237	0.325	(-0.420, 0.807)	0.393	0.318	(-0.319, 0.889)	0.369	0.313	(-0.335, 0.867)
ρ_1	0.422	0.285	(-0.245, 0.853)	0.404	0.371	(-0.473, 0.930)	0.397	0.387	(-0.514, 0.942)

$\mathcal{M}_3$ for artisanal and industrial fisheries, respectively.

The variability associated with the measurement error in $\mathcal{M}_1$ is large (posterior means 0.662 and 0.486 for artisanal and industrial fisheries, respectively). These values are somewhat lower in the case of $\mathcal{M}_2$ due to the inclusion of the autoregressive term in the conditional matrix. The relevant difference happens with $\mathcal{M}_3$ in which the inclusion of the autoregressive term in the conditional mean takes away a large amount of the measurement error variability.

3.4.1 Model comparison

All of the three proposed models seem reasonable and show consistent and robust results. However, we would like to be able to choose one of the three as the most appropriate and compatible with the data. Model checking is an essential issue of any statistical analysis. The existing literature on this subject is abundant and multifaceted due to the great quantity and disparity of criteria proposed (Vehtari and Ojanen, 2012).

It is not our aim to study this issue in depth because it is beyond the scope of this chapter. Here, we would like to briefly discuss this subject and obtain some results that will help us to compare our three models. In this regard, we will focus on three methods that are quite popular in the Bayesian reasoning: the penalised expected deviance (PED), the Bayes factor (BF) and the conditional predictive ordinate (CPO). PED was proposed by Plummer (2008). It is based on the deviance information criterion (DIC), commented in Chapter 2, that combines the expected deviance as a measure of fit and the effective number of parameters of the model as a measure of their complexity. The complexity penalty used by PED is higher than that of DIC so it is to be understood that PED will generally favour simpler models than DIC.

The Bayes factor is regarded as the conceptual solution to model selection problems but presents many practical problems. According to Kass and Raftery (1995) and Berger and Pericchi (1996), it is defined in terms of prior predictive distributions evaluated on the observed data (evidence) which can be interpreted as the support provided by the data in favour of the subsequent model. We will use the R package

`bridgesampling` (Gronau, Singmann, and Wagenmakers, 2020) that uses bridge sampling (Meng and Wong, 1996; Meng and Schilling, 2002) to approximate the evidence from each of the three models examined. The bridge sampling method is centered around identifying a function enabling the calculation of the Bayes factor of two distributions, of which we know everything except their integration constant, i.e., their marginal likelihood. This particular software package determines the marginal likelihood by comparing a distribution with a constant of 1 to the one for which we lack information about the integration constant. Consequently, the ratio it computes essentially represents the unknown constant, the marginal likelihood.

Finally, as in Chapter 2, we use the cross-validated predictive density to assess our models. In our case, it is defined as the conditional posterior density of a future artisanal and industrial tonnage of sardines caught of country i in an hypothetical replicated experiment

$$f(\mathbf{y}_i^{pred(A)}, \mathbf{y}_i^{pred(I)} \mid \mathcal{D}^{-(i)}) = \int f(\mathbf{y}_i^{pred(A)}, \mathbf{y}_i^{pred(I)} \mid \boldsymbol{\theta}, \boldsymbol{\phi}) \pi(\boldsymbol{\theta}, \boldsymbol{\phi} \mid \mathcal{D}^{-(i)}) d(\boldsymbol{\theta}, \boldsymbol{\phi})$$

where $\mathcal{D}^{-(i)}$ are all data in $\mathcal{D}$ except for the observations of country i (leave-one-out procedure). Particularly, we consider the *CPO* for country i , i.e. CPO_i , defined as the value of $f(\mathbf{y}_i^{pred(A)}, \mathbf{y}_i^{pred(I)} \mid \mathcal{D}^{-(i)})$ at the observed $(\mathbf{y}_i^{(A)}, \mathbf{y}_i^{(I)})$ data. Large CPO_i values are supportive of the model as they indicate a good agreement between the data and the model.

Table 3.2 shows the PED value for models $\mathcal{M}_1$, $\mathcal{M}_2$ and $\mathcal{M}_3$ as well as the sample and standard deviation of ten replicates of the approximated evidence, in the logarithmic scale. Both criteria point to $\mathcal{M}_3$ as the best model, followed in both cases by $\mathcal{M}_2$. It seems clear that the two models with autoregressive elements perform better than the basic model, the PED is clearly in favour of $\mathcal{M}_3$, but the differences are narrower in the case of the logarithm of the evidence.

Table 3.3 includes the *CPO*, in logarithmic scale, for each of the countries in the study associated with models $\mathcal{M}_1$, $\mathcal{M}_2$ and $\mathcal{M}_3$. This latter model records the

Table 3.2: Approximated PED for models $\mathcal{M}_1$, $\mathcal{M}_2$ and $\mathcal{M}_3$ and mean (standard deviation) of ten replicates of the approximate evidence, in logarithmic scale, for models $\mathcal{M}_1$, $\mathcal{M}_2$ and $\mathcal{M}_3$.

	$\mathcal{M}_1$	$\mathcal{M}_2$	$\mathcal{M}_3$
PED	1226	804	161
log(Evidence)	-802.52 (0.37)	-531.91 (0.01)	-517.93 (1.02)

highest *CPO* values in all of the countries except Montenegro. In this country, the highest *CPO* is obtained with $\mathcal{M}_2$, although the differences between the *CPO* of this country with both models is quite small. $\mathcal{M}_1$ has the smallest *CPO* values in all of the countries with a large difference compared to ones from $\mathcal{M}_2$ and $\mathcal{M}_3$. Spain is in all models by far the most different country from the rest, followed by Slovenia and Croatia. The *CPO* values obtained suggest that the model most compatible with the data is also $\mathcal{M}_3$.

The inclusion of an autoregressive term in the model clearly seems to be a good decision, and better on the conditional mean than on the variance. If we had to choose one model it would be $\mathcal{M}_3$.

Table 3.3: *CPO*'s, in logarithmic scale, for the different countries in the study computed from models $\mathcal{M}_1$, $\mathcal{M}_2$ and $\mathcal{M}_3$.

	Log <i>CPO</i>		
	$\mathcal{M}_1$	$\mathcal{M}_2$	$\mathcal{M}_3$
Algeria	-35.45	-28.30	-25.29
Croatia	-125.88	-86.74	-76.46
France	-65.40	-46.63	-45.44
Greece	-53.00	-31.54	-30.06
Italy	-51.68	-34.94	-31.74
Montenegro	-50.86	-37.40	-37.94
Morocco	-38.19	-29.56	-27.31
Slovenia	-97.49	-95.91	-95.28
Spain	-353.48	-169.58	-163.76
Turkey	-189.13	-44.60	-35.94

3.4.2 Association and country-specific fishing patterns

As a complement to the statistical analysis carried out, we now turn to $\mathcal{M}_3$ and present some of the derived results. We recall that $\beta_1^{(A)}$ and $\beta_1^{(I)}$ is, respectively, the regression slope associated with time in the conditional mean of the total tonnage (in logarithmic scale) of sardines caught by artisanal and industrial methods. The subsequent posterior mean and standard deviation of $\beta_1^{(A)}$ is -0.029 and 0.031 and -0.035 and 0.023 for $\beta_1^{(I)}$. Figure 3.4 shows the marginal posterior distribution of $\beta_1^{(A)}$ and $\beta_1^{(I)}$ and shows in the shaded part the posterior probability that the subsequent β_1 is negative, 0.912 and 0.857, for the artisanal and industrial fishery, respectively. These values give strong support to a decreasing trend for both types of fisheries.




Figure 3.4: Approximate marginal posterior distribution of $\beta_1^{(A)}$ and $\beta_1^{(I)}$, respectively, for model $\mathcal{M}_3$.

As mentioned before, an interesting issue of our study is the possible association between the specific characteristics of artisanal and industrial fisheries in the different countries of the study: do both types of fisheries have a synergistic relationship and generate a mutually cooperative relationship? Or, on the contrary, are they negatively associated, and therefore it is to be expected that high values for industrial fisheries are linked to low values for artisanal fisheries and vice versa?.

Figure 3.5 shows the mean of the posterior distribution of the random effects associated with the intercept and slope of each of the countries in the study. Slovenia, Montenegro and Spain started the study with low levels associated with small-scale fisheries, below the level of the other countries in the study. These countries behave

very differently with respect to industrial fishing: above, at the average and below the rest of the countries in the study. This is the case for Spain, Slovenia and Montenegro, respectively. Algeria has a level of artisanal fishing well above the rest but its level of industrial fishing is in the average of the countries studied. The rest of the countries have specific characteristics above the average of the countries in the study, both in artisanal and industrial fishing. Italy and Algeria stand out as the countries with a higher level associated with industrial and artisanal fishing than the rest. In particular, it is worth noting that Spain is the country with the greatest imbalance, in favour of industrial fishing over artisanal fishing at the beginning of the study.




Figure 3.5: Artesanal vs. Industrial fishing. Approximate posterior mean of the random effects parameters from the model $\mathcal{M}_3$ and their estimated correlations (dashed lines): (a) random intercepts and (b) random trends.

The evolution of artisanal and industrial fisheries has not been the same. Spain, Italy, Greece, France and, especially, Croatia exhibit decreasing dynamics for artisanal fisheries while in the other countries, Algeria, Morocco, Montenegro and Turkey, the trends are positive. It seems that artisanal fisheries are in recession in the economically stronger countries. Industrial fisheries show a positive trend in almost all countries except France and Italy. The countries with the highest pattern of industrial fishing are Turkey and Algeria. Although the association between sardine catches and artisanal

and industrial fishing is not very strong, it is generally positive, so there seems to be a certain synergy between both types of fishing, although not in all of the countries, as in the case of Croatia that shows an apparent negative relationship.

In addition, Figure 3.6 illustrates the approximate posterior distribution of the correlation parameters ρ_0 and ρ_1 for model $\mathcal{M}_3$. Note that each of these parameters measures the linear relationship between the random effects associated with the artisanal and industrial responses. It can be seen that the marginal posterior distributions obtained are mostly positive. Particularly the probability of being positive are 0.882 and 0.848, for ρ_0 and ρ_1 , respectively.




Figure 3.6: Approximate marginal posterior distribution of ρ_0 and ρ_1 , respectively, for model $\mathcal{M}_3$.

3.5 Conclusions

We presented a Bayesian joint longitudinal modelling to assess the temporal evolution of the European sardine landings, industrial and artisanal, in the Mediterranean countries from 1985 to 2018. These models dealt with general patterns as well as individual characteristics of the countries in the study.

Overall results confirmed that Mediterranean fisheries were highly diverse from the beginning of the time series and this heterogeneity still remains over time. This general result is in agreement with a recent study that analysed the convergence of the Mediterranean countries in terms of several ecological indicators (i.e. Marine Trophic Index, Fishing in Balance Index and Expansion Factor) during 1950-2010 (Pennino

et al., 2017), showing strong temporal persistence in their fishery behaviours. Furthermore, the results demonstrate a general decline in European sardine fishing quantities in the Mediterranean Sea. Additionally, there is a positive correlation observed between the landing dynamics of artisanal and industrial fishing within each country.

As mentioned in the introduction, the decline of sardine fisheries is a multifaceted issue that would need different approaches from different points of view. Our study has specific but interesting objectives because it provides quantitative information on the disparity in the evolution of industrial and artisanal sardine fishing in the different Mediterranean countries. Bayesian modelling is very flexible and would allow for the inclusion of additional information through baseline or temporal covariates. In this sense, it would be interesting to be able to feed our modelling with information related to the economic characteristics of the different countries and their fishing fleets, as well as geographical and environmental information on the areas where the sardine shoals are located.

Chapter 4

Bayes factor in longitudinal models through power posteriors

4.1 Introduction

Model selection is a key issue in Bayesian statistics, as noted in Chapter 1, and it has generated a large scientific literature. Several proposals have been made from different perspectives to carry out model comparison. This is a multifaceted topic with a high philosophical content that will undoubtedly continue to improve scientific knowledge in the future.

As mentioned earlier, Bayesian methodology uses a conception of probability that allows assigning probability distributions to any kind of quantity with uncertainty. The posterior distribution of these quantities is the basis of the natural Bayesian procedure for model selection through the posterior distribution for each of the candidate models. When comparing two competing models, the ratio of their corresponding posterior distributions (posterior odds) is obtained through the product of the prior odds and the Bayes factor, the latter defined in terms of the ratio of the marginal likelihood of each of the two models (Kass and Raftery, 1995; Berger and Pericchi, 1996). This is a conceptually simple and powerful procedure with a key limitation in that it is indeterminate when working with improper prior distributions, i. e., those distributions

that do not integrate to unity.

The computation of marginal likelihoods for assessing Bayes factors can be demanding and complex. In fact, this situation is exacerbated in models with random effects such as linear mixed models (LMMs). The marginal likelihood function for these models involves the integral with respect to the prior distribution of the conditional likelihood function that depends on both the parameters and the random effects. The dimension of this integral increases with the size of the random effect set that, in turn, increases with the number of individuals in the sample.

The challenge in calculating marginal likelihoods in complex models has produced a large literature and a fruitful scientific debate, which in turn generated different proposals for its computation. All of them have interesting properties, but none of them has managed to close the issue definitively, mainly due to computational problems. Although it is not our aim to present in this chapter an exhaustive list of all of the proposals, we would like to provide brief comments on some of the most popular ones. One of the first procedures to estimate the marginal likelihood was the Laplace's method (Tierney and Kadane, 1986). Then, Newton and Raftery (1994) approximated the marginal likelihood through the posterior harmonic mean of the likelihood. This is a simple but computationally unstable approach (Raftery et al., 2007). Chib (1995) and Chib and Jeliazkov (2001) developed an algorithm which computes marginal likelihoods from the MCMC outputs by making use of additional iterations. Importance sampling ideas have been proposed for estimating the evidence such as annealed importance sampling (Neal, 2001) or bridge sampling (Meng and Wong, 1996). Nested sampling proposed by Skilling (2006) represents the marginal likelihood in terms of one-dimensional integral over $[0, 1]$. This is an interesting approach based on simulated values from the prior distribution subject to constraints in the conditional likelihood but it involves many challenges when working with multidimensional models and prior distributions poorly informative.

The power posterior (Lartillot and Philippe, 2006, Friel and Pettitt, 2008, and Friel, Hurn, and Wyse, 2014) is a proposal developed in the context of MCMC algorithms

in which the logarithm of the marginal likelihood is evaluated numerically through the path sampling algorithm of Gelman and Meng (1998), also known as thermodynamic integration.

This chapter generalises the use of the power posterior for approximating marginal likelihoods into complex models which include random effects and serial correlation terms in their conditional formulations. In particular, we extend the use of power posteriors to LMMs, although it can be easily generalised to any other models with random effects, and we discuss its behaviour through two simulated studies and a real study on European sardine (*Sardina pilchardus*, Walbaum, 1792) landings in the Mediterranean Sea. The data and the code used in this study are available at https://github.com/gcalvobayarri/Bayes_factor_longitudinal_models.git.

This chapter is organised as follows. Section 2 introduces the basic Bayesian linear mixed models (BLMMs) and describes two popular generalisations based on autoregressive terms. Section 3 reviews Bayes factors and marginal likelihoods and presents the power posterior for models with random effects. Section 4 assesses the behaviour of the power posterior in two simulated longitudinal studies. The first one deals with data generated from a simple LMM model that competes with three basic models. The second data set was simulated from a model with a high level of complexity, in the sense that it includes three types of random variation: random effects, an autoregressive term and normally distributed measurement errors. This model is compared with two simpler competing models with only two of the three sources of variation discussed above. Section 5 is devoted to the selection of a Bayesian longitudinal model to analyse sardine fisheries in different countries of the Mediterranean Sea. The chapter concludes with a small discussion emphasising the usefulness of the power posterior methodology in the calculation of marginal likelihoods for comparing LMMs, and highlights the most interesting results of the three studies carried out in the work.

4.2 Bayesian longitudinal linear mixed models

Let $\mathbf{y}_i = (y_{i1}, \dots, y_{in_i})'$ be the random vector describing the response of individual i , $i = 1, \dots, N$, recorded at times $\mathbf{t}_i = (t_{i1}, \dots, t_{in_i})'$, and consider $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_N)'$. A Bayesian linear mixed model (BLMM) for $\mathbf{y}$ is specified via the formulation of the joint probability distribution expressed in (1.4).

The simplest BLMM assumes the following conditional normal distribution for $\mathbf{y}_i$:

$$f(\mathbf{y}_i \mid \boldsymbol{\theta}, \boldsymbol{\phi}) = \mathcal{N}(\boldsymbol{\mu}_i = \mathcal{X}_i \boldsymbol{\beta} + \mathcal{Z}_i \mathbf{b}_i, \boldsymbol{\Sigma}_i = \sigma^2 \mathbf{I}_{n_i}), \quad (4.1)$$

where $\mathcal{X}_i$ and $\mathcal{Z}_i$ are the design matrices for the fixed effects $\boldsymbol{\beta}$ and the random effects $\mathbf{b}_i$, respectively, and $\boldsymbol{\Sigma}_i$ is the variance-covariance matrix defined in terms of the identity matrix $\mathbf{I}_{n_i}$ of size n_i and the common variance σ^2 . Here, the vector of random effects $\boldsymbol{\phi}$ in (4.1) is $\mathbf{b} = (\mathbf{b}_1, \dots, \mathbf{b}_N)'$ whose components are assumed conditional i.i.d. $(\mathbf{b}_i \mid \boldsymbol{\Sigma}_b) \sim \mathcal{N}(0, \boldsymbol{\Sigma}_b)$. This conditional model (4.1) implies conditional independence not only between the observations of the different individuals but also of all observations of the same individual.

Common generalisations of (4.1) when the number of observations per individual is large can include temporal elements that account for serial correlation (Diggle et al., 2002), frequently in terms of autoregressive processes in the conditional mean $\boldsymbol{\mu}_i$ or in the conditional variance-covariance matrix $\boldsymbol{\Sigma}_i$. In this sense, we consider two different modelling approaches based on a first order autoregressive ($AR(1)$) process. The first proposal introduces an $AR(1)$ term in the random measurement (Chi and Reinsel, 1989; Hedeker and Gibbons, 2006) which, as a result, reformulates the (j, l) element of $\boldsymbol{\Sigma}_i$ in (4.1) as follows

$$\text{Cov}(y_{ij}, y_{il} \mid \boldsymbol{\theta}, \boldsymbol{\phi}_i) = \frac{\sigma^2}{1 - \rho^2} \rho^{|j-l|}, \quad (4.2)$$

where ρ , $|\rho| < 1$, is the coefficient of the autoregressive term of a stationary $AR(1)$. Because of the stationarity condition, ρ is also equal to the conditional autocorrelation

at lag 1. More generally, ρ^h is the conditional autocorrelation at lag h or $-h$, with $h \geq 0$. This model preserves the conditional independence between individuals but not between the observations of the same individual. Furthermore, a reformulation of this model is possible by conditioning the response of individual i in the j -th measurement (y_{ij}) on their $(j - 1)$ -th measurement in the following manner

$$f(y_{ij} \mid y_{ij-1}, \boldsymbol{\theta}, \boldsymbol{\phi}) = \mathcal{N}(\mu_{ij} + \rho(y_{ij-1} - \mu_{ij-1}), \sigma^2), \quad j = 2, \dots, n_i \quad (4.3)$$

being $f(y_{i1} \mid \boldsymbol{\theta}, \boldsymbol{\phi}) = \mathcal{N}(\mu_{i1}, \sigma^2/(1 - \rho^2))$.

The second proposal introduces a latent autoregressive element $\mathbf{w}_i = (w_{i1}, \dots, w_{in_i})'$ in the conditional mean $\boldsymbol{\mu}_i$ of (4.1) as follows:

$$\boldsymbol{\mu}_i = \mathcal{X}_i \boldsymbol{\beta} + \mathcal{Z}_i \mathbf{b}_i + \mathbf{w}_i, \quad (4.4)$$

where each w_{ij} , $j = 2, \dots, n_i$, is the j -th realisation from a Gaussian process with mean ρw_{ij-1} and variance σ_w^2 , where $(w_{i1} \mid \rho, \sigma_w) \sim \mathcal{N}(0, \sigma_w^2/(1 - \rho^2))$ (Diggle et al., 2002), that is $\mathbf{w}_i$ is a vector of correlated noises. This model augments the vector of random effects associated with individual i to $\boldsymbol{\phi}_i = (\mathbf{b}_i, \mathbf{w}_i)$. The two elements in $\boldsymbol{\phi}_i$ are conditionally independent given $\boldsymbol{\theta}$.

The full specification of the Bayesian model is completed with the elicitation of a prior distribution $\pi(\boldsymbol{\theta})$ for $\boldsymbol{\theta}$. This includes parameters in $\boldsymbol{\beta}$ and $\boldsymbol{\Sigma}_i$ as well as hyperparameters in the covariance-matrix $\boldsymbol{\Sigma}_b$ and in the autoregressive terms.

Following Bayes theorem in (1.1), we can calculate the posterior distribution of $(\boldsymbol{\theta}, \boldsymbol{\phi})$ based on the information provided by the data $\mathcal{D}$, which consists of the observations of $\mathbf{y}$. This computation forms the cornerstone of the Bayesian inferential process, allowing us to derive posterior distributions for key outcomes of the problem. It is expressed as

$$\pi(\boldsymbol{\theta}, \boldsymbol{\phi} \mid \mathcal{D}) \propto f(\mathcal{D} \mid \boldsymbol{\theta}, \boldsymbol{\phi}) f(\boldsymbol{\phi} \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta}), \quad (4.5)$$

where $f(\mathcal{D} \mid \boldsymbol{\theta}, \boldsymbol{\phi})$ represents now the likelihood function of $(\boldsymbol{\theta}, \boldsymbol{\phi})$ for data $\mathcal{D}$ which is expressed as the product of each individual's contribution to the likelihood:

$$f(\mathcal{D} \mid \boldsymbol{\theta}, \boldsymbol{\phi}) = \prod_{i=1}^N f(\mathcal{D}_i \mid \boldsymbol{\theta}, \boldsymbol{\phi}) = \prod_{i=1}^N \frac{\exp\left(-\frac{1}{2}(\mathcal{D}_i - \boldsymbol{\mu}_i)' \boldsymbol{\Sigma}_i (\mathcal{D}_i - \boldsymbol{\mu}_i)\right)}{\sqrt{(2\pi)^{n_i} |\boldsymbol{\Sigma}_i|}}. \quad (4.6)$$

Note that the posterior distribution $\pi(\boldsymbol{\theta}, \boldsymbol{\phi} \mid \mathcal{D})$ in (4.5) involves $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ together. We have chosen to express the prior information on them as the product $f(\boldsymbol{\phi} \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta})$ and not as $\pi(\boldsymbol{\theta}, \boldsymbol{\phi})$. This is a philosophically debatable topic and it is not our intention to enter into this issue here. We simply chose that expression because we recognised it as the most standard one.

4.3 Bayes factor and power posteriors

Henceforth, we assume to have a set of alternative BLMMs $\mathcal{M}_k$ with parameter and hyperparameter $\boldsymbol{\theta}_k$ and random effects $\boldsymbol{\phi}_k$ ($k = 1, \dots, K$). The marginal or predictive density $m_k(\mathbf{y})$ for $\mathbf{y}$, introduced in Chapter 1 is defined as

$$\begin{aligned} m_k(\mathbf{y}) &= \int f(\mathbf{y}, \boldsymbol{\theta}_k, \boldsymbol{\phi}_k) d(\boldsymbol{\theta}_k, \boldsymbol{\phi}_k) \\ &= \int f(\mathbf{y} \mid \boldsymbol{\theta}_k, \boldsymbol{\phi}_k) f(\boldsymbol{\phi}_k \mid \boldsymbol{\theta}_k) \pi(\boldsymbol{\theta}_k) d(\boldsymbol{\theta}_k, \boldsymbol{\phi}_k). \end{aligned} \quad (4.7)$$

This distribution evaluated on the data, $m_k(\mathcal{D})$, can be interpreted as the support provided by the data in favour of model $\mathcal{M}_k$.

The main tool used for choosing between two alternative models (i. e., $K = 2$), $\mathcal{M}_1$ and $\mathcal{M}_2$, in the Bayesian methodology is the Bayes factor (Kass and Raftery, 1995; Berger and Pericchi, 1996) of model $\mathcal{M}_1$ against model $\mathcal{M}_2$. It is defined in terms of their marginal likelihood functions as follows

$$B_{12} = \frac{m_1(\mathcal{D})}{m_2(\mathcal{D})},$$

and measures the strength with which the data support $\mathcal{M}_1$ with regard to $\mathcal{M}_2$.

In the following, we will generalise the power posterior to the case of longitudinal models involving not only parameters and hyperparameters $\boldsymbol{\theta}$ but also random effects ϕ . Accordingly, we define the power posterior of $(\boldsymbol{\theta}, \phi)$ as

$$\pi_\tau(\boldsymbol{\theta}, \phi \mid \mathcal{D}) \propto f(\mathcal{D} \mid \boldsymbol{\theta}, \phi)^\tau f(\phi \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta}), \quad (4.8)$$

where $\tau \in [0, 1]$ is an auxiliary temperature variable that modulates the effect of the likelihood. Let the “power marginal likelihood” be

$$m(\mathcal{D} \mid \tau) = \int f(\mathcal{D} \mid \boldsymbol{\theta}, \phi)^\tau f(\phi \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta}) d(\boldsymbol{\theta}, \phi). \quad (4.9)$$

When $\tau = 0$, $m(\mathcal{D} \mid \tau = 0) = \int f(\phi \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta}) d(\boldsymbol{\theta}, \phi) = 1$, whereas when $\tau = 1$, $m(\mathcal{D} \mid \tau = 1)$ equals the marginal likelihood $m(\mathcal{D})$, our target quantity. Moreover, derived from the equation (4.9) we obtain the following result:

$$\begin{aligned} \mathbb{E}_{(\boldsymbol{\theta}, \phi \mid \mathcal{D}, \tau)}(\log(f(\mathcal{D} \mid \boldsymbol{\theta}, \phi))) &= \int \log(f(\mathcal{D} \mid \boldsymbol{\theta}, \phi)) \pi_\tau(\boldsymbol{\theta}, \phi \mid \mathcal{D}) d(\boldsymbol{\theta}, \phi) \\ &= \frac{1}{m(\mathcal{D} \mid \tau)} \int \log(f(\mathcal{D} \mid \boldsymbol{\theta}, \phi)) f(\mathcal{D} \mid \boldsymbol{\theta}, \phi)^\tau f(\phi \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta}) d(\boldsymbol{\theta}, \phi) \\ &= \frac{1}{m(\mathcal{D} \mid \tau)} \int \left[\frac{d}{d\tau} f(\mathcal{D} \mid \boldsymbol{\theta}, \phi)^\tau f(\phi \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta}) \right] d(\boldsymbol{\theta}, \phi) \\ &= \frac{1}{m(\mathcal{D} \mid \tau)} \frac{d}{d\tau} \left[\int f(\mathcal{D} \mid \boldsymbol{\theta}, \phi)^\tau f(\phi \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta}) d(\boldsymbol{\theta}, \phi) \right] \\ &= \frac{1}{m(\mathcal{D} \mid \tau)} \frac{d}{d\tau} m(\mathcal{D} \mid \tau) \\ &= \frac{d}{d\tau} \log(m(\mathcal{D} \mid \tau)). \end{aligned}$$

As a consequence, we have

$$\int_0^1 \mathbb{E}_{(\boldsymbol{\theta}, \phi \mid \mathcal{D}, \tau)}(\log(f(\mathcal{D} \mid \boldsymbol{\theta}, \phi))) d\tau = \log\left(\frac{m(\mathcal{D} \mid \tau = 1)}{m(\mathcal{D} \mid \tau = 0)}\right) = \log(m(\mathcal{D})). \quad (4.10)$$

By means of the expectations, expression (4.10) provides numerical stability to the estimation of the log of the marginal likelihood (Friel and Pettitt, 2008).

For any value $\tau \in [0, 1]$, an estimate of the expectation $\mathbb{E}_{(\boldsymbol{\theta}, \phi \mid \mathcal{D}, \tau)}(\log(f(\mathcal{D} \mid$

θ, ϕ)) and an estimate of the variance $V_{(\theta, \phi | \mathcal{D}, \tau)}(\log(f(\mathcal{D} | \theta, \phi)))$ are obtained by using the MCMC draws generated from the power posterior $\pi_\tau(\theta, \phi | \mathcal{D})$. Running $M + 1$ MCMC algorithms with a temperature ladder such as $0 = \tau_0 < \tau_1 < \dots < \tau_M = 1$, Friel and Pettitt (2008) provide an estimate of $\log m(\mathcal{D})$ by applying the trapezoidal rule, which in our case works as follows

$$\begin{aligned} \log(m(\mathcal{D})) &\approx \sum_{m=0}^{M-1} (\tau_{m+1} - \tau_m) \times \\ &\times \frac{1}{2} \left(E_{(\theta, \phi | \mathcal{D}, \tau_m)}(\log(f(\mathcal{D} | \theta, \phi))) + E_{(\theta, \phi | \mathcal{D}, \tau_{m+1})}(\log(f(\mathcal{D} | \theta, \phi))) \right). \end{aligned} \quad (4.11)$$

The temperature ladder from $\tau = 0$ to $\tau = 1$ defines a path from the prior distribution to the posterior distribution, hence the name of path sampling (Gelman and Meng, 1998).

4.4 Simulation studies

Now, we explore the behaviour of the power posterior for approximating the evidence in BLMMs. Two separate simulation studies are conducted to evaluate the capability of the power posterior to identify the model that generated the data within a set of potential competitors. In the first study, we work with a small data set generated by a simple model with random effects and normally distributed measurement errors. In the second case, we examine a larger data set generated by a more complex model, which includes random effects, measurement errors, and an autoregressive term.

We provide the complete posterior densities required for the MCMC process to implement the power posterior methodology in Appendix B.2. These densities correspond to the most general models, particularly those utilised in the second simulation study mentioned in this chapter.

4.4.1 Study 1: a balanced longitudinal data set

We assume a scenario based on a small balanced set of simulated longitudinal data. Specifically, each individual response $\mathbf{y}_i$ has the same number of measurements $n_i = 10$, for any $i = 1, \dots, N$ with $N = 5$. Moreover, all individual observations are recorded at identical times $\mathbf{t} = (t_1 = 0, \dots, t_j = j - 1, \dots, t_{10} = 9)$.

In the first step, we generate the data from a simple model. Then, we approximate the evidence of this model for the simulated data. After that, we compute the marginal likelihood of the three competing longitudinal models. Finally, we compare the marginal likelihood of the generator model with the evidence of three other competing longitudinal models for the same data set.

In particular, we consider the generator BLMM (4.1) with design matrices $\mathcal{X}_i = \mathcal{Z}_i = \mathbf{1}$, where $\mathbf{1}$ is the n_i -dimensional vector of ones. We represent by $\mathbf{1} \beta_0$ the common intercept and by $\mathbf{1} b_{0i}$ the individual random intercepts, which are conditionally normally distributed, $(b_{0i} \mid \sigma_0^2) \sim \mathcal{N}(0, \sigma_0^2)$. We set the values $\beta_0 = 2$, $\sigma = 0.5$, and $\sigma_0 = 1.5$ for the parameters and hyperparameter of the model (see Figure 4.1).

Four competing longitudinal models $\mathcal{M}_1, \mathcal{M}_2, \mathcal{M}_3$ and $\mathcal{M}_4$, ordered from the least to the most complex, are taken into account for analysing the simulated data, i.e.,

$$\begin{aligned} \mathcal{M}_1 : f(y_{ij} \mid \boldsymbol{\theta}_1) &= \mathcal{N}(\beta_0, \sigma^2), \\ \mathcal{M}_2 : f(y_{ij} \mid \boldsymbol{\theta}_2, \boldsymbol{\phi}_2) &= \mathcal{N}(\beta_0 + b_{0i}, \sigma^2), \\ \mathcal{M}_3 : f(y_{ij} \mid \boldsymbol{\theta}_3, \boldsymbol{\phi}_3) &= \mathcal{N}(\beta_0 + b_{1i}t_j, \sigma^2), \\ \mathcal{M}_4 : f(y_{ij} \mid \boldsymbol{\theta}_4, \boldsymbol{\phi}_4) &= \mathcal{N}(\beta_0 + b_{0i} + b_{1i}t_j, \sigma^2). \end{aligned} \quad (4.12)$$

The model that generated the data is $\mathcal{M}_2$. It includes a random effect associated with the intercept (b_{0i}) of each individual and a normal measurement error. Model $\mathcal{M}_1$, with only normal measurement errors, is the simplest model. $\mathcal{M}_3$ and $\mathcal{M}_4$ includes normal measurement errors, $\mathcal{M}_3$ also includes random slopes (b_{1i}) associated with individuals, and $\mathcal{M}_4$ considers random intercepts and random slopes.

All four models assume conditional independence between individuals and within observations from the same individual as well as homogeneity of variances. Individual random effects b_{0i} and random slopes b_{1i} are mutually independent and conditionally normally distributed as $(b_{0i} | \sigma_0^2) \sim \mathcal{N}(0, \sigma_0^2)$ and $(b_{1i} | \sigma_1^2) \sim \mathcal{N}(0, \sigma_1^2)$.

The elicitation of the subsequent prior distribution $\pi(\theta)$ in all models is based on both prior independence among the parameters and a noninformative prior scenario: a normal distribution for the common intercept $\pi(\beta_0) = \mathcal{N}(0, 10^2)$ and wide uniforms for the standard deviation parameters $\pi(\sigma) = \pi(\sigma_0) = \pi(\sigma_1) = \text{U}(0, 10)$.




Figure 4.1: Spaghetti plot of the simulated longitudinal data from the first study, 5 individuals are considered in a complete balanced design for times $\{0, 1, \dots, 9\}$.

For the computation of the evidence via the power posterior in each of the four models in (4.12), the temperature variable τ was discretised according to $0 = \tau_0 < \tau_1 < \dots < \tau_{199} = 1$ with $\tau_r = (r/199)^5$ and $r = 0, \dots, 199$. This temperature ladder ensures that a high proportion of values are close to 0, improving the convergence of the algorithm. Each power posterior is computed by the Gibbs sampling, and implemented in the R environment, version 4.0.5, (R Core Team, 2021) running a chain with 50,000 iterations for each discrete value of τ . Then, we calculate an estimate of the model evidence applying the trapezoidal rule (4.11). In addition, to quantify the variability of the process, we repeat the algorithm ten times per model.

Table 4.1 shows for each model the mean and the standard deviation of the ten values of the approximate logarithm of the evidence. The evidence for each of the

Table 4.1: Mean (standard deviation) of the ten replicates of the approximate log evidence for each model in (4.12) computed by means of the power posterior method.

$\mathcal{M}_1$	$\mathcal{M}_2$	$\mathcal{M}_3$	$\mathcal{M}_4$
$-103.01(0.14)$	$-51.63(0.11)$	$-91.83(0.14)$	$-68.21(0.32)$

four models is clearly ordered. The model with the highest log evidence (-51.63) is actually the true sampling model $\mathcal{M}_2$, followed by the complete linear mixed model $\mathcal{M}_4$ which includes a common population intercept and two types of individual random effects, b_{0i} and b_{1i} . In contrast, the model with the lowest log marginal likelihood value (-103.01) is the fixed effects model $\mathcal{M}_1$. The standard deviation for $\mathcal{M}_4$ is the highest (0.32), possibly because it is the most complex of the four models. Note also that the variability associated with the replicate process is relatively low in all cases.

4.4.2 Study 2: an unbalanced longitudinal data set with serial correlation

Now, we consider a scenario where we examine an unbalanced set of simulated longitudinal data (Figure 4.2). The BLMM that has generated the data includes a common intercept β_0 , an individual random slope b_{1i} for each individual ($i, i = 1, \dots, 10$), where $(b_{1i} \mid \sigma_1) \sim \mathcal{N}(0, \sigma_1^2)$, and an autoregressive latent element $\mathbf{w}_i = (w_{i1}, \dots, w_{in_i})'$ in the conditional mean. The autoregressive latent element is defined in (4.4) with parameters σ_w^2 and ρ . Additionally, there is a common conditional variance σ^2 . This model, which we will call $\mathcal{M}_1$, can be written as:

$$\mathcal{M}_1 : f(\mathbf{y}_i \mid \boldsymbol{\theta}_1, \phi_1) = \mathcal{N}(\mathbf{1}\beta_0 + \mathbf{t}_i b_{1i} + \mathbf{w}_i, \mathbf{1}\sigma^2). \quad (4.13)$$

We simulate from this model using the values $\beta_0 = 2$, $\sigma_1 = 0.5$, $\sigma = 2$, $\rho = 0.8$, and $\sigma_w = 1.5$ for the parameters and hyperparameter of the model. Observation times of the response variable for each individual i , $\mathbf{t}_i = (t_{i1}, \dots, t_{in_i})'$, were generated in a doubly random manner, i.e., both the number of observations (between 10 and 70) and all times when each observation is recorded (between 0 and 20). The total number of observations for the 10 individuals is 441.




Figure 4.2: Spaghetti plot of the unbalanced longitudinal data from the second study, 10 individuals are considered with a total of 441 observations.

We compare $\mathcal{M}_1$ with the two alternative models $\mathcal{M}_2$ and $\mathcal{M}_3$ defined as follows:

$$\begin{aligned}\mathcal{M}_2 : f(\mathbf{y}_i \mid \boldsymbol{\theta}_2, \boldsymbol{\phi}_2) &= \mathcal{N}(\mathbf{1}\beta_0 + \mathbf{1}b_{0i} + \mathbf{t}_i b_{1i}, \mathbf{1}\sigma^2), \\ \mathcal{M}_3 : f(\mathbf{y}_i \mid \boldsymbol{\theta}_3, \boldsymbol{\phi}_3) &= \mathcal{N}(\mathbf{1}\beta_0 + \mathbf{t}_i b_{1i}, \Sigma_{AR}).\end{aligned}\quad (4.14)$$

Both competing models $\mathcal{M}_2$ and $\mathcal{M}_3$ are somewhat simpler than $\mathcal{M}_1$. Model $\mathcal{M}_2$ exchanges the autoregressive term for an individual random effect in the intercept b_{0i} , with $(b_{0i} \mid \sigma_0) \sim \mathcal{N}(0, \sigma_0^2)$, and model $\mathcal{M}_3$ has only a random slope effect but with an autoregressive element in the measurement errors.

We complete the Bayesian models by eliciting a prior distribution for the parameters and hyperparameters. We assume prior independence among them and select a uniform distribution $\mathcal{U}(0, 10)$ for the standard deviation parameters σ , σ_0 , σ_1 and σ_w , and a $\mathcal{U}(-1, 1)$ for the autoregressive parameter ρ . The normal distribution $\mathcal{N}(0, 10^2)$ is chosen for the common intercept β_0 .

Table 4.2: Mean (standard deviation) of the ten replicates of the approximate log evidence for models in (4.13) and (4.14) computed by means of the power posterior method.

$\mathcal{M}_1$	$\mathcal{M}_2$	$\mathcal{M}_3$
-1123.54 (0.48)	-1158.45 (1.17)	-1131.04 (1.13)

Table 4.2 shows, for each model, the mean and the standard deviation of the 10 replicates of the approximate logarithmic evidence. Model $\mathcal{M}_1$, which is the true sampling model, presents the highest evidence (-1123.54), followed by $\mathcal{M}_3$ (-1131.04). The lowest marginal likelihood value (-1158.45) is for model $\mathcal{M}_2$. Note that the variability associated with the computation of the model predictive probability is higher than that of the previous simulation study due to the higher complexity of these models.

4.5 Application - Sardine landings in the Mediterranean Sea

Let us recall that, as discussed in Chapter 3, the European sardine (*Sardina pilchardus*) is one of the most commercially important species that exhibited the highest over-exploitation rates in the past 20 years (Coll et al., 2008). The data analysed are derived from the online database www.seaaroundus.org of the research initiative Sea Around Us. However, our focus lies in analysing the tonnes of European sardine caught by countries in the Mediterranean Sea from 1970 to 2014, irrespective of the fishery type, in contrast to the analysis from the previous chapter.

Remember that the countries participating in our study included Albania, Algeria, Bosnia and Herzegovina (B&H), Croatia, France, Greece, Italy, Montenegro, Morocco, Slovenia, Spain, and Turkey. This online database is derived mainly from FAO global fisheries catch statistics, complemented by the statistics of various international and national agencies, and reconstructed data sets. Finally, a logarithmic transformation was applied to the landing data set in order to approach the normality assumption. Figure 4.3 shows the temporal pattern of the logarithm of the tonnes of sardines caught per country.

The amount of fish caught at the beginning of the study is very variable among the different countries. B&H is initially well below the level of the other countries. The temporal evolution of this quantity in most of the countries seems to be rather stable, although a slightly increasing trend can be seen in most countries, especially B&H and




Figure 4.3: Annual log tons of European sardines (*Sardina pilchardus*) caught per country, from 1970 to 2014.

Turkey. Slovenia's behaviour in recent years has been different from that of the others, with a decreasing trend in the number of catches. Observed annual changes in the same country with respect to the previous years are generally not very large. Thus, it would seem reasonable to consider an autoregressive component when modelling the evolution of the sardine landings.

4.5.1 Modelling of sardine fisheries

We are in a balanced scenario, similar to the situation described in Subsection 4.4.1. Let y_{ij} be the logarithm of the total tonnage of sardines caught in country i ($i = 1, \dots, 12$) at measurement j ($j = 1, \dots, 45$) in year t_j . The calendar time serves as the natural time scale for this study, where $t_1 = 0$ corresponds to 1970, the first year of the study, and $t_{45} = 44$ corresponds to year 2014.

We consider three longitudinal models, $\mathcal{M}_1$, $\mathcal{M}_2$, and $\mathcal{M}_3$, to assess the dynamics of the official sardine fishery carried out by country i in the Mediterranean Sea from

1970 to 2014. They were defined as follows

$$\begin{aligned}\mathcal{M}_1 : f(\mathbf{y}_i \mid \boldsymbol{\theta}_1, \phi_1) &= \mathcal{N}(\mathbf{1} \beta_0 + \mathbf{1} b_{0i} + \mathbf{t} b_{1i}, \mathbf{1} \sigma^2), \\ \mathcal{M}_2 : f(\mathbf{y}_i \mid \boldsymbol{\theta}_2, \phi_2) &= \mathcal{N}(\mathbf{1} \beta_0 + \mathbf{t} b_{1i}, \boldsymbol{\Sigma}_{AR}), \\ \mathcal{M}_3 : f(\mathbf{y}_i \mid \boldsymbol{\theta}_3, \phi_3) &= \mathcal{N}(\mathbf{1} \beta_0 + \mathbf{t} b_{1i} + \mathbf{w}_i, \mathbf{1} \sigma^2),\end{aligned}\tag{4.15}$$

where in all models $\mathbf{1}$ is a vector of ones of dimension 45, $\mathbf{t} = (t_1 = 0, \dots, t_j = j - 1, \dots, t_{45} = 44)'$, β_0 is a common intercept and σ^2 a common variance. Individual random effects b_{0i} and random slopes b_{1i} are mutually independent and conditionally normally distributed as $(b_{0i} \mid \sigma_0^2) \sim \mathcal{N}(0, \sigma_0^2)$ and $(b_{1i} \mid \sigma_1^2) \sim \mathcal{N}(0, \sigma_1^2)$. The elements of the variance-covariance matrix in model $\mathcal{M}_2$ are as in (4.2), and the autoregressive element $\mathbf{w}_i$ of model $\mathcal{M}_2$ is defined as in (4.4).

We complete the Bayesian models in (4.15) by specifying a prior distribution for the subsequent parameters and hyperparameters of each model. In all of them we assume prior independence and a vague prior scenario. We consider uniform distributions for the standard deviation parameters $\pi(\sigma) = \pi(\sigma_0) = \pi(\sigma_1) = \pi(\sigma_w) = \mathcal{U}(0, 5)$ and $\pi(\rho) = \mathcal{U}(-1, 1)$ for the autoregressive parameter in models $\mathcal{M}_2$ and $\mathcal{M}_3$. A normal distribution is considered for the common intercept parameter $\pi(\beta_0) = \mathcal{N}(0, 5^2)$.

The estimation of the log-evidence for all of the three models in (4.15) was derived through the power posterior according to the same strategy used in the previous section by means of the Gibbs sampling. We considered 200 temperatures for models $\mathcal{M}_1$ and $\mathcal{M}_2$, 80,000 iterations, and a burn-in of 20,000 for each temperature value τ_r . However, due to the complexity of the model $\mathcal{M}_3$, we increased the number of temperature values to 500, and obtain the discretisation $0 = \tau_0 < \tau_1 < \dots < \tau_{499} = 1$, where $\tau_r = (r/499)^5$ and $r = 0, \dots, 499$, as well as to increase the number of iterations to 200,000 and the burn-in to 50,000. Table 4.3 shows the mean and the standard deviation of ten replicates of the approximate log-evidence of the three models.

The evidence value of the model $\mathcal{M}_1$ is clearly lower than the other two models

Table 4.3: Mean (standard deviation) of the approximate log evidence for each model in (4.15) by means of the power posterior method.

$\mathcal{M}_1$	$\mathcal{M}_2$	$\mathcal{M}_3$
$-515.57(0.29)$	$-191.38(0.14)$	$-193.36(0.79)$

$\mathcal{M}_2$ and $\mathcal{M}_3$, this indicates that the autoregressive term is relevant in modelling sardine fishing. The last two models have similar log-evidence values but if one of them should be selected, this would be $\mathcal{M}_2$, albeit by a very small margin. Note that the approximate Bayes factor of model $\mathcal{M}_2$ to $\mathcal{M}_3$ is 7.25, which gives evidence in favour of $\mathcal{M}_2$, but not very strong. An alternative possibility that could be interesting, but beyond our scope of this memory, would be to deal with models $\mathcal{M}_2$ and $\mathcal{M}_3$ through model averaging procedures (Hoeting et al., 1999).

4.5.2 Longitudinal modelling of European sardine landings

We carry out a Bayesian statistical analysis of the sardine landings by means of model $\mathcal{M}_2$ and discuss some of the relevant outputs derived from the subsequent posterior distribution.

It has been evaluated through Gibbs sampling, running a chain of 250,000 iterations after a burn-in of 50,000, and thinning the chain at every 250th iteration to reduce its autocorrelation. A summary of the posterior outputs is shown in Table 4.4.

Table 4.4: Summary of the approximate posterior distribution of the parameters and hyperparameters in model $\mathcal{M}_2$.

	mean	sd	$q_{0.025}$	$q_{0.975}$
β_0	7.92	0.61	6.48	9.05
σ_0	3.02	0.95	0.68	4.67
σ_1	0.03	0.02	0.00	0.06
σ	0.31	0.01	0.29	0.34
ρ	0.97	0.03	0.91	1

At the beginning of the study, $t = 0$, we notice a substantial common intercept as well as a large heterogeneity among the different countries. The country-specific variability related to the slope is not very large although its magnitude may be relevant

due to the magnitude of the t values. The variability associated with the model is around 0.31 and it is worth noting the high precision of this estimate. Finally, we observe very high and positive values of the correlation coefficient ρ (around 0.97), which would indicate the importance of the autoregressive term in the modelling.




Figure 4.4: Random intercepts vs. random slopes. Approximate posterior mean for the random intercepts and the slope effects by country according to $\mathcal{M}_2$.

The posterior distribution of the random effects associated with each country provides useful information on the individual country patterns. Figure 4.4 shows the posterior mean of the random effect associated with the intercept and the slope of the countries. It can be seen that at the beginning of the time series a group of countries, including Algeria, Croatia and Morocco, had higher landing values than the overall average and kept this superiority over time. However, Greece and Spain and, to a greater extent, Italy and France also started with above-mean levels of fishing but their growth seems to slow down over time. B&H, Turkey, and Slovenia are countries with quite different dynamics from the rest. B&H started with a level of fishing well below the common mean, but increased the landings during the time-series (possibly as a consequence of the end of the Balkan wars) and now it is above the rest of the countries. The same trend can be observed in Turkey, although in the beginning it was in the mean

of the rest of the countries. In contrast, Slovenia's growth has slowed down over the years from an initial state in the mean of the rest of countries.

Finally, we can use the MCMC sample to investigate the autocorrelation function, which depends on the autoregressive parameter ρ . Figure 4.5 shows its approximate posterior mean and 95% credible interval. As expected, the correlation decreases with the increasing of the time lag. However, the uncertainty involved in the posterior distribution is higher as the lag between the two response variables increases.




Figure 4.5: Posterior mean (line) and 95% credible interval (shaded area) of the autocorrelation function according to the posterior distribution of the autoregressive parameter ρ in $\mathcal{M}_2$.

The analysis we presented on sardine landing is an interesting illustration of the potential of Bayesian longitudinal models to analyse the general and individual behaviour of the different experimental units, in this case countries.

4.6 Conclusions

The power posterior distribution, the key concept in the power posterior methodology, allows computation of the marginal likelihoods by extending the Gibbs sampling

process quite naturally, i. e., by doing Gibbs sampling in each of the power posterior distributions, which are as simple to derive as the posterior distribution. This makes the implementation of this method ideal for Bayesian longitudinal models with different types of random effects and different levels of complexity. Some variations in the power posterior algorithm can be implemented using importance sampling to avoid sampling from the posterior distribution (Xie et al., 2011) or from distributions close to the prior (Fan et al., 2011). These improvements may slightly reduce the computational cost of the method. In addition, small changes in the trapezoidal rule for estimating the evidence on the logarithmic scale can be applied in order to reduce the bias of the approximation (Friel, Hurn, and Wyse, 2014).

In the two simulated studies we discussed, the correct model (i. e., that generating the data) is always selected among competing models with different sources of random variation, demonstrating the efficiency of the power posterior method. Moreover, following this methodology, one of the main conclusions in the analysis of the European Sardine landings in the Mediterranean Sea is that the autoregressive term is relevant in its modelling. Actually, the model that includes two sources of random variation (random effects and autoregressive errors) is that with the highest marginal likelihood value.

Chapter 5

A Bayesian hidden Markov model for assessing the hot hand phenomenon in basketball shooting performance

5.1 Introduction

The study of success or failure streaks in sport is a vast issue that has its foundation in the literature of psychology, where it is known as “psychological momentum” (Adler, 1981). Although this topic can be studied in different sports, this chapter focuses on “momentum”, known as the *hot hand* phenomenon in the course of a basketball game.

In basketball, a player or a team is said to have a *hot hand* when their performance in a particular period of time is better than expected or they are on a streak of making consecutive shots. Belief in this phenomenon is evident both in the media and among fans, as well as in many experts from the world of sport. But it also generates some controversy. Some authors, such as Gilovich, Vallone, and Tversky (1985), state that while most fans believe in some sort of correlation among hits and misses in basketball,

the *hot hand* phenomenon is somewhat of an illusion. On the other hand, different studies contradicted Gilovich's results. In particular, Miller and Sanjurjo (2018) recently argued that there was clearly a bias in that analysis. However, evidence can be hard to compare due to the diversity of methods used to assess this behaviour. Some authors made use of hidden Markov models (HMMs) to analyse this behaviour in basketball (Sun, 2004; Wetzels et al., 2016; Sandri, Zuccolotto, and Manisera, 2020) or in other sports such as baseball (Albert and Bennett, 2001). Finally, Ötting et al. (2020) made use of a hidden autoregressive term which indicates the performance level of a player in darts.

This chapter discusses the *hot hand* phenomenon in basketball shooting performance and proposes a general Bayesian longitudinal HMM (BLHMM) to describe that effect. This model structures each game played by the same team as an observational unit and includes both information on the different shots of the same game as well as those in the rest of the games of the season. This is a novel approach that incorporates longitudinal information for baskets made or missed in each game within a general scenario controlled by a HMM that governs the effectiveness of the team's play. The proposed model allows analysing the periods in which the team is in a *hot* or *cold* phase and connecting these phases with the success and failure of the shots in the game. This model could be a powerful tool for assessing the overall performance of a basketball team, especially when it is on a good or bad streak, that could easily be adapted to any other team sport played with a ball.

The structure of the chapter is as follows. Section 2 introduces the BLHMM in terms of a sampling model based on two connected processes, one latent (or hidden) and the other observed, and a prior distribution for all of the unknown elements of the model. Section 3 deals with the posterior distribution for some characteristics of the basketball team performance such as transition probabilities, occupancy and sojourn times, and probabilities of making a basket. Section 4 applies the Bayesian model to a real data set from the Miami Heat team in the season 2005-2006 of the National Basketball Association (NBA). We conclude with a discussion, including potential

directions for future research.

5.2 Bayesian longitudinal hidden Markov modelling

Probabilities in the Bayesian framework are always conditional and will therefore appear as such in their corresponding descriptions. This situation increases the complexity of the notation but clarifies the theoretical framework of Bayesian inference much better.

We propose a BLHMM for analysing the shooting performance of a basketball team in a season or a group of j games as a joint probabilistic model $f(\mathbf{y}, \mathbf{Z}, \boldsymbol{\theta}, \boldsymbol{\phi})$ for the observed process $\mathbf{y}$, the hidden process $\mathbf{Z}$, the random effects $\boldsymbol{\phi}$ associated with both processes, and the parameters and hyperparameters $\boldsymbol{\theta}$. This joint distribution can be factorised as expressed in (1.2)

$$f(\mathbf{y}, \mathbf{Z}, \boldsymbol{\theta}, \boldsymbol{\phi}) = \underbrace{f(\mathbf{y} | \mathbf{Z}, \boldsymbol{\theta}, \boldsymbol{\phi})}_{\text{Observable process}} \underbrace{f(\mathbf{Z} | \boldsymbol{\theta}, \boldsymbol{\phi})}_{\text{Hidden process}} f(\boldsymbol{\phi} | \boldsymbol{\theta}) \pi(\boldsymbol{\theta}).$$

Additionally, we assume a general framework of conditional i.i.d. for game i of the season ($i = 1, \dots, N$), so that

$$f(\mathbf{y}, \mathbf{Z}, \boldsymbol{\theta}, \boldsymbol{\phi}) = \left(\prod_{i=1}^N f(\mathbf{y}_i | \mathbf{Z}_i, \boldsymbol{\theta}, \boldsymbol{\phi}) f(\mathbf{Z}_i | \boldsymbol{\theta}, \boldsymbol{\phi}) \right) f(\boldsymbol{\phi} | \boldsymbol{\theta}) \pi(\boldsymbol{\theta}). \quad (5.1)$$

We describe below the different probabilistic elements included in the product (5.1), with the first two expressed in terms of a generic game i .

The hidden process. This is a hidden Markov chain (HMC) that describes the *hot hand* situation of consecutive shots at the basket of the team in game i . The transition probabilities of the HMC are modelled using logistic mixed regression models.

Let $\{Z_{ij}, j = 1, \dots, n_i\}$ be a HMC which represents the state, *cold* (C) or *hot* (H), of the team in shooting j at the basket in game i , where n_i is the number of shots

in the game i . The transition probability matrix is expressed as

$$P_i = \begin{pmatrix} p_i^{(CC)} & p_i^{(CH)} \\ p_i^{(HC)} & p_i^{(HH)} \end{pmatrix}, \quad (5.2)$$

where $p_i^{(CH)} = P(Z_{ij+1} = H \mid Z_{ij} = C, \boldsymbol{\theta}, \boldsymbol{\phi})$ is the conditional (on $\boldsymbol{\phi}$ and $\boldsymbol{\theta}$) transition probability from the *cold* C to the *hot* state H in the game i , $p_i^{(CC)} = 1 - p_i^{(CH)}$, $p_i^{(HC)} = P(Z_{ij+1} = C \mid Z_{ij} = H, \boldsymbol{\theta}, \boldsymbol{\phi})$, and $p_i^{(HH)} = 1 - p_i^{(HC)}$. Thus, the symbol $\mathbf{Z}_i$ of (5.1) is the random vector $(Z_{i1}, \dots, Z_{in_i})$.

The complete specification of the Markov chain needs to set a probability distribution for the initial state of the chain $\boldsymbol{\delta}_i = (\delta_i^{(C)}, \delta_i^{(H)}) = (P(Z_{i1} = C \mid \boldsymbol{\theta}, \boldsymbol{\phi}), P(Z_{i1} = H \mid \boldsymbol{\theta}, \boldsymbol{\phi}))$. We know that, from a theoretical standpoint, the values of the discrete latent states are natural numbers and not letters. However, for the sake of readability, we chose to represent the values of the variables in the chain with letters, C and H .

We assume a logistic mixed regression model for the transition probabilities as follows:

$$\begin{aligned} \text{logit}(p_i^{(CH)} \mid \boldsymbol{\theta}, \boldsymbol{\phi}) &= \mathbf{X}_i \boldsymbol{\beta}^{(CH)} + b_i^{(CH)}, \\ \text{logit}(p_i^{(HC)} \mid \boldsymbol{\theta}, \boldsymbol{\phi}) &= \mathbf{X}_i \boldsymbol{\beta}^{(HC)} + b_i^{(HC)}, \end{aligned}$$

where $\mathbf{X}_i$ is a vector of baseline covariates for individual i , and $\boldsymbol{\beta}^{(CH)}$ and $\boldsymbol{\beta}^{(HC)}$ are regression coefficient vectors associated with transitions from C to H and from H to C , respectively. The random effects associated with these transition probabilities within the game are $\mathbf{b}_i = (b_i^{(CH)}, b_i^{(HC)})$, which are assumed conditionally jointly normally distributed as $(\mathbf{b}_i \mid \boldsymbol{\theta}) \sim \mathcal{N}(\mathbf{0}, \Sigma_b)$, where Σ_b is a variance-covariance matrix. Note that $\boldsymbol{\beta} = (\boldsymbol{\beta}^{(CH)}, \boldsymbol{\beta}^{(HC)})$ and Σ_b are parameters and hyperparameters, respectively, and, consequently, they are included in the generic $\boldsymbol{\theta}$ vector. The same applies for the random effects b_i 's as elements in $\boldsymbol{\phi}$.

The observed process. A Bernoulli longitudinal model is used to assess the success

or failure of shots at the basket in each game i , with probability depending on the state (*cold* or *hot*) of the team. These state-dependent probabilities are modelled through logistic mixed regression models.

Let the random variable y_{ij} be the success (1) or failure (0) of shot j in game i of the team. Thus, the symbol $\mathbf{y}_i$ of (5.1) is the random vector $(y_{i1}, \dots, y_{in_i})$. Each variable in $\mathbf{y}_i$ is Bernoulli whose probability depends on the latent state Z_{ij} as follows,

$$\begin{aligned} f(y_{ij} \mid \boldsymbol{\theta}, \boldsymbol{\phi}, Z_{ij} = C) &= \mathcal{B}(\gamma_{ij}^{(C)}), \\ f(y_{ij} \mid \boldsymbol{\theta}, \boldsymbol{\phi}, Z_{ij} = H) &= \mathcal{B}(\gamma_{ij}^{(H)}), \end{aligned}$$

where $\gamma_{ij}^{(C)}$ and $\gamma_{ij}^{(H)}$ are the success probability parameters associated with the *cold* and *hot* states, respectively, denoting as $\mathcal{B}$ the Bernoulli distribution. Consider the parameters vectors $\boldsymbol{\gamma}_i^{(C)} = (\gamma_{i1}^{(C)}, \dots, \gamma_{in_i}^{(C)})$ and $\boldsymbol{\gamma}_i^{(H)} = (\gamma_{i1}^{(H)}, \dots, \gamma_{in_i}^{(H)})$. These probabilities can be modelled through a logistic mixed regression model expressed as

$$\begin{aligned} \text{logit}(\boldsymbol{\gamma}_i^{(C)} \mid \boldsymbol{\theta}, \boldsymbol{\phi}) &= \mathcal{X}_i \boldsymbol{\alpha}^{(C)} + \mathcal{Z}_i \mathbf{a}_i^{(C)}, \\ \text{logit}(\boldsymbol{\gamma}_i^{(H)} \mid \boldsymbol{\theta}, \boldsymbol{\phi}) &= \mathcal{X}_i \boldsymbol{\alpha}^{(H)} + \mathcal{Z}_i \mathbf{a}_i^{(H)}, \end{aligned}$$

where $\mathcal{X}_i$ and $\mathcal{Z}_i$ are the design matrices for the fixed $\boldsymbol{\alpha} = (\boldsymbol{\alpha}^{(C)}, \boldsymbol{\alpha}^{(H)})$ and the random effects $\mathbf{a}_i = (\mathbf{a}_i^{(C)}, \mathbf{a}_i^{(H)})$, respectively. Matrix $\mathcal{X}_i$ includes the covariates associated with the success of shooting, for instance the distance to the basket or the type of shot. Random effects $\mathbf{a}_i$ are conditionally independent and assumed to follow a conditional multivariate normal distribution $(\mathbf{a}_i \mid \Sigma_a) \sim \mathcal{N}(\mathbf{0}, \Sigma_a)$. Note that parameters $\boldsymbol{\alpha}$ and hyperparameters in Σ_a are part of $\boldsymbol{\theta}$ (along with $\boldsymbol{\beta}$, Σ_b , and $\boldsymbol{\delta} = (\boldsymbol{\delta}_1, \dots, \boldsymbol{\delta}_N)$), whereas the $\mathbf{a}_i$'s of $\boldsymbol{\phi}$ (along with $\mathbf{b} = (\mathbf{b}_1, \dots, \mathbf{b}_N)$). Therefore, $\boldsymbol{\theta} = (\boldsymbol{\alpha}, \boldsymbol{\beta}, \Sigma_a, \Sigma_b, \boldsymbol{\delta})$ and $\boldsymbol{\phi} = (\mathbf{a}, \mathbf{b})$, where $\mathbf{a} = (\mathbf{a}_1, \dots, \mathbf{a}_N)$.

Finally, as is customary, the Bayesian model is completed by specifying a prior distribution, denoted as $\pi(\boldsymbol{\theta})$, over the parameters and hyperparameters of the model.

5.3 Basketball team performance

Let $\pi(\boldsymbol{\theta}, \boldsymbol{\phi} \mid \mathcal{D})$ be the posterior distribution that contains all information of the behaviour of the model. This information includes direct knowledge of the performance of the team of interest both at a general level, such as the coefficients of the logistic regressions associated with either the hidden transition probabilities or the state dependent probabilities of making a basket, and at a specific level of the different games in the analysed championship. At the specific level, the posterior distribution provides knowledge on the random effects associated with a particular game on transition probabilities or probabilities of making a basket. That information is very valuable but does not provide direct knowledge on some of the indicators of the team performance that may be important to better understand the strengths and weaknesses of the team. This is the case of transition probabilities, probability of making a basket, probability of being in a *cold* or a *hot* state, occupancy times, first-passage times, etc. Since the stochastic behaviour of these quantities depends on $(\boldsymbol{\theta}, \boldsymbol{\phi})$, we can apply the Bayesian paradigm and compute the posterior distribution of each of them from the posterior distribution. We discuss these inferences below.

5.3.1 Transition probabilities

For the sake of space, we will only develop the results of the inference for the transition probabilities of the HMC defined in (5.2) from the *cold* to the *hot* state $p_i^{(CH)}$. The same strategy can be applied, without loss of generality, to the rest of the transition probabilities of the chain. Therefore,

$$P(Z_{i,j+1} = H \mid Z_{ij} = C, \boldsymbol{\theta}, \boldsymbol{\phi}) = \frac{\exp\{\mathbf{X}_i \boldsymbol{\beta}^{(CH)} + b_i^{(CH)}\}}{1 + \exp\{\mathbf{X}_i \boldsymbol{\beta}^{(CH)} + b_i^{(CH)}\}}. \quad (5.3)$$

Since this transition probability depends on $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$, we can approach its posterior distribution,

$$\pi(P(Z_{i,j+1} = H \mid Z_{ij} = C, \boldsymbol{\theta}, \boldsymbol{\phi}) \mid \mathcal{D}), \quad (5.4)$$

from the simulated sample of the posterior $\pi(\boldsymbol{\theta}, \phi \mid \mathcal{D})$ and then compute the relevant characteristics of that distribution (mean, standard deviation, credible intervals, etc). If we also want to evaluate these probabilities in a general way, without distinguishing between games, we could obtain the marginal transition probability

$$P(Z_{i,j+1} = H \mid Z_{ij} = C, \boldsymbol{\theta}) = \int P(Z_{i,j+1} = H \mid Z_{ij} = C, \boldsymbol{\theta}, \phi) f(\phi \mid \boldsymbol{\theta}) d\phi, \quad (5.5)$$

which is independent of the game and only depends on $\boldsymbol{\theta}$. Consequently, its posterior distribution,

$$\pi(P(Z_{i,j+1} = H \mid Z_{ij} = C, \boldsymbol{\theta}) \mid \mathcal{D}), \quad (5.6)$$

could be approximated using a sample from the conditional posterior distribution $\pi(\boldsymbol{\theta} \mid \mathcal{D})$ which is easily obtained through the sample from $\pi(\boldsymbol{\theta}, \phi \mid \mathcal{D})$.

5.3.2 t -step transition probabilities and stationary distribution

The conditional t -step transition probabilities of the chain for game i are defined as

$$\begin{aligned} p_i^{(CH)}(t) &= P(Z_{i,j+t} = H \mid Z_{i,j} = C, \boldsymbol{\theta}, \phi), \quad p_i^{(CC)}(t) = 1 - p_i^{(CH)}(t), \\ p_i^{(HC)}(t) &= P(Z_{i,j+t} = C \mid Z_{i,j} = H, \boldsymbol{\theta}, \phi), \quad p_i^{(HH)}(t) = 1 - p_i^{(CH)}(t). \end{aligned}$$

From the Markov chain theory (Kulkarni, 2016) we know that the matrix that collects these probabilities, $P_i^{(t)}$, can be calculated from the transition probability matrix P_i in (5.2) as

$$\begin{aligned} P_i^{(t)} &= \frac{1}{p_i^{(CH)} + p_i^{(HC)}} \begin{bmatrix} p_i^{(HC)} & p_i^{(CH)} \\ p_i^{(HC)} & p_i^{(CH)} \end{bmatrix} \\ &\quad + \frac{(1 - p_i^{(CH)} - p_i^{(HC)})^t}{p_i^{(CH)} + p_i^{(HC)}} \begin{bmatrix} p_i^{(CH)} & -p_i^{(CH)} \\ -p_i^{(HC)} & p_i^{(HC)} \end{bmatrix}. \end{aligned}$$

Since the HMC for game i has finite state space, is irreducible and aperiodic, its

stationary distribution will be the limit distribution, $\Delta_i = (\Delta_i^{(C)}, \Delta_i^{(H)})$, which turns out to be

$$\Delta_i^{(C)} = \frac{p_i^{(HC)}}{p_i^{(CH)} + p_i^{(HC)}}, \quad \Delta_i^{(H)} = \frac{p_i^{(CH)}}{p_i^{(CH)} + p_i^{(HC)}}. \quad (5.7)$$

Note that the vector Δ_i is also the left eigenvector of the transition probability matrix P_i , associated with the eigenvalue one.

The Bayesian framework allows us to calculate the posterior distribution of each of the t -step transition probabilities of the chain for each specific game as well as for a generic game by means of the posterior distribution of the conditional marginal distribution as explained in Subsection 5.3.1.

5.3.3 Occupancy times

Occupancy times in a chain refers to the expected number of visits the process makes to each state of the chain in a given number of transitions (Kulkarni, 2016). This is an indicator of the frequency with which the chain visits the different states of the process. In particular, for each game i we represent by $m_i^{(CH)}(t) = E(V_i^{(CH)}(t) \mid \theta, \phi)$ the conditional (on (θ, ϕ)) expected value of the number of visits $V_i^{(CH)}(t)$ to state H from state C in the first t transitions of the chain. In the case of our HMC, these conditional expectations are expressed in terms of the transition probabilities as follows

$$\begin{aligned} \begin{bmatrix} m_i^{(CC)}(t) & m_i^{(CH)}(t) \\ m_i^{(HC)}(t) & m_i^{(HH)}(t) \end{bmatrix} &= \frac{t+1}{p_i^{(CH)} + p_i^{(HC)}} \begin{bmatrix} p_i^{(HC)} & p_i^{(CH)} \\ p_i^{(HC)} & p_i^{(CH)} \end{bmatrix} \\ &+ \frac{1 - (p_i^{(HC)} + p_i^{(CH)} - 1)^{(t+1)}}{(p_i^{(CH)} + p_i^{(HC)})^2} \begin{bmatrix} p_i^{(CH)} & -p_i^{(CH)} \\ -p_i^{(HC)} & p_i^{(HC)} \end{bmatrix}. \end{aligned} \quad (5.8)$$

For inferential purposes, we should proceed in the same way as with the transition probabilities. We could calculate the posterior distribution of each occupancy time that would give us information on each game or proceed in a general way by integrating

out the random effects associated with the individual games and obtain posterior information of the number of visits among a given number of transitions in a generic game.

5.3.4 Sojourn times

The sojourn time of the basketball team in state k , $k \in \{C, H\}$, of game i , $\tau_i^{(k)}$, is the number of shots required for the team to leave state k for the first time, i.e.

$$\begin{aligned} P(\tau_i^{(k)} = t \mid \boldsymbol{\theta}, \boldsymbol{\phi}) &= P(Z_{ij+t} \neq k, Z_{it^*} = k, t^* = j+1, \dots, j+t-1 \mid Z_{ij} = k, \boldsymbol{\theta}, \boldsymbol{\phi}) \\ &= [p_i^{(kk)}]^{t-1} (1 - p_i^{(kk)}), \quad t = 0, 1, \dots \end{aligned}$$

Therefore, the conditional distribution of the sojourn time in state C (H) of the game i will be a geometric distribution of parameter $p_i^{(CC)}$ ($p_i^{(HH)}$). Consequently, we can compute the posterior distribution for $\tau_i^{(C)}$ and $\tau_i^{(H)}$

$$\begin{aligned} \pi((\tau_i^{(C)} = t \mid \boldsymbol{\theta}, \boldsymbol{\phi}) \mid \mathcal{D}), \\ \pi((\tau_i^{(H)} = t \mid \boldsymbol{\theta}, \boldsymbol{\phi}) \mid \mathcal{D}), \end{aligned} \tag{5.9}$$

or its marginal posterior distribution integrating out the random effects as above.

5.3.5 Probabilities of making a basket

The procedure for inferring the probability of making the j th shot of game i when the team is in a *hot* or *cold* state is similar to the one we have developed for calculating the posterior distribution of transition probabilities in subsection 5.3.1. Thus, the relevant conditional posterior distributions are

$$\begin{aligned} \pi(P(y_{ij} = 1 \mid Z_{ij} = C, \boldsymbol{\theta}, \boldsymbol{\phi}) \mid \mathcal{D}), \\ \pi(P(y_{ij} = 1 \mid Z_{ij} = H, \boldsymbol{\theta}, \boldsymbol{\phi}) \mid \mathcal{D}). \end{aligned}$$

If interested in a general performance of the team, independently of the game,

we can also marginalise the random effects associated with the games, compute the conditional marginal distributions $P(y_{ij} = 1 \mid Z_{ij} = C, \boldsymbol{\theta}, \mathcal{D})$ and $P(y_{ij} = 1 \mid Z_{ij} = H, \boldsymbol{\theta}, \mathcal{D})$ and its subsequent posterior distributions

$$\begin{aligned} \pi(P(y_{ij} = 1 \mid Z_{ij} = C, \boldsymbol{\theta}) \mid \mathcal{D}), \\ \pi(P(y_{ij} = 1 \mid Z_{ij} = H, \boldsymbol{\theta}) \mid \mathcal{D}). \end{aligned}$$

In addition, we can compute the posterior distribution $\pi(P(y_{ij} = 1 \mid \boldsymbol{\theta}) \mid \mathcal{D})$ of the probability of making the j th shot of game i when the state of the team is unknown. It can be simulated from the expressions

$$\begin{aligned} P(y_{ij} = 1 \mid \boldsymbol{\theta}, \phi) &= P(y_{ij} = 1 \mid Z_{ij} = C, \boldsymbol{\theta}, \phi) P(Z_{ij} = C \mid \boldsymbol{\theta}, \phi) \\ &+ P(y_{ij} = 1 \mid Z_{ij} = H, \boldsymbol{\theta}, \phi) P(Z_{ij} = H \mid \boldsymbol{\theta}, \phi), \end{aligned}$$

where

$$\begin{aligned} P(Z_{ij} = C \mid \boldsymbol{\theta}, \phi) &= P(Z_{ij} = C \mid Z_{i1} = C, \boldsymbol{\theta}, \phi) P(Z_{i1} = C \mid \boldsymbol{\theta}, \phi) \\ &+ P(Z_{ij} = C \mid Z_{i1} = H, \boldsymbol{\theta}, \phi) P(Z_{i1} = H \mid \boldsymbol{\theta}, \phi), \end{aligned}$$

and the simulated sample from the posterior distribution $\pi(\boldsymbol{\theta}, \phi \mid \mathcal{D})$.

5.4 Miami Heat shooting in the 2005-2006 NBA season

On 20 June 2006, the Miami Heat team defeated the Dallas Mavericks in the finals and won their first NBA championship. The analysis of the information play-by-play provided by the NBA (2020) about the team is very interesting, and gives many clues on their successful season. Specifically, we are interested in analysing their shooting performance. We consider all of the *field* and *free throws* of all games of the championship, a total of 11,042 shots in 105 games, 5,922 of which were baskets made. For each of these shots the distance in feet from the shooting position to the basket, the game, and the sequential order per game in which each shot was taken were taken into

account.

Figure 5.1 is a shot chart (Zuccolotto and Manisera, 2020) of the Miami Heat's *field shots* during the 2005-2006 NBA season. It can be clearly seen that most shots are taken just beyond the *three-point line* or just below the basket. On the other hand, as might be expected, most of the shots the team took from further away were missed.




Figure 5.1: Shot chart of the Miami Heat during the 2005-2006 season, showing the location of shots on the court, and indicating the shot results and the exact position of the basket.

5.4.1 Modelling the shooting performance of the Miami Heat team

We apply the Bayesian model presented in Section 5.2. The sampling model is also defined in terms of the two sub-processes, one defined in terms of an HMC that accounts for the *hot* and *cold* states, and the other in terms of an observed Bernoulli longitudinal variable that assesses the success or failure of each shot in relation to the *hot* or *cold* state of the chain.

The transition probabilities of the HMC $\{Z_{ij}, j = 1, \dots, n_i\}$ defined for each game i , $i = 1, \dots, N = 105$, are expressed through the logistic mixed regression

models

$$\begin{aligned}\text{logit}(p_i^{(CH)} | \boldsymbol{\theta}, \boldsymbol{\psi}) &= \beta_{CH} + b_i^{(CH)}, \\ \text{logit}(p_i^{(HC)} | \boldsymbol{\theta}, \boldsymbol{\psi}) &= \beta_{HC} + b_i^{(HC)},\end{aligned}$$

which express each of the two relevant probabilities, $p_i^{(CH)}$ and $p_i^{(HC)}$, in terms of an intercept (indicating the average logit of the transition probabilities), common to all games and a specific random effect, $b_i^{(CH)}$ and $b_i^{(HC)}$, associated with the game. These random effects are mutually independent and conditionally normally distributed as $(b_i^{(CH)} | \sigma_{CH}) \sim \mathcal{N}(0, \sigma_{CH}^2)$ and $(b_i^{(HC)} | \sigma_{HC}) \sim \mathcal{N}(0, \sigma_{HC}^2)$. The initial probability vector $\boldsymbol{\delta} = (\delta_C, \delta_H)$ is also considered.

For the observed part of the sampling model, the probabilities $\gamma_{ij}^{(H)}$ and $\gamma_{ij}^{(C)}$ of making a basket on shot j in game i when the team is respectively in the *hot* or the *cold* state are described as

$$\begin{aligned}\text{logit}(\gamma_{ij}^{(C)}) &= \alpha_C + \alpha_d \text{Dist}_{ij} + \alpha_{FT} I_{FT}(ij) + a_i, \\ \text{logit}(\gamma_{ij}^{(H)}) &= \alpha_H + \alpha_d \text{Dist}_{ij} + \alpha_{FT} I_{FT}(ij) + a_i,\end{aligned}$$

where α_C and α_H are the common intercepts for the probability associated with the *cold* and the *hot* state respectively, α_d and α_{FT} are the regression coefficients associated with the covariates Dist_{ij} and $I_{FT}(ij)$. These covariates respectively describe the distance from the basket in the j th shot of the game i and an indicator variable that is 1 when the j -th shot of game i is a *free throw* and zero otherwise. Random effects a_i are assumed normally distributed, i. e. $(a_i | \sigma_a) \sim \mathcal{N}(0, \sigma_a^2)$, and conditional independent given σ_a , for any i .

In order to complete the specification of the BLHMM model we assume prior independence within a minimally informative prior scenario. In particular, we select a beta distribution for the initial probability $\pi(\delta_C) = \text{Be}(1, 1)$, and uniform distributions for the standard deviation parameters $\pi(\sigma_a) = \pi(\sigma_{CH}) = \pi(\sigma_{HC}) = \mathcal{U}(0, 10)$. For

the two regression coefficients, α_d and α_{FT} , associated with the covariates a wide normal distribution is selected, $\pi(\beta_{CH}) = \pi(\beta_{HC}) = \pi(\alpha_d) = \pi(\alpha_{FT}) = \mathcal{N}(0, 10^2)$. Finally, to avoid the problem of identifiability because of the label switching issue (see among others McCulloch and Tsay, 1994; Frühwirth-Schnatter, 2001; Spezia, 2009), we include the following restriction in the α_C 's and α_H 's prior distributions

$$\pi(\alpha_C) = \pi(\alpha_H) = \mathcal{N}(0, 10^2), \quad \alpha_C \leq \alpha_H.$$

5.4.2 Posterior distribution

The complexity of the BLHMM model makes the posterior distribution analytically intractable. Thus, as in previous chapters, we approximated it by means of MCMC sampling methods via the JAGS software. Three parallel chains were run for 30000 iterations each after burn-ins of 30000 iterations. In addition, based on the estimated autocorrelation in the sample, and in order to reduce it, the chains were also thinned at every 30th iteration. Moreover, the full analysis, performed by an R code (R version 4.0.5), and the data are available as supplementary material at https://github.com/gcalvobayarri/hot_hand_model.git.

Table 5.1: Posterior summaries (mean, standard deviation, 95% credible interval and $\hat{R}$) for the parameters and hyperparameters of the Miami Heat shooting performance *hot hand* model.

Submodel		mean	sd	$q_{0.025}$	$q_{0.975}$	$\hat{R}$
Latent	β_{CH}	-0.49	0.05	-0.58	-0.39	1.001
	β_{HC}	0.38	0.06	0.27	0.49	1.000
	δ_C	0.55	0.06	0.43	0.68	1.000
	σ_{CH}	0.07	0.05	0.00	0.18	1.011
	σ_{HC}	0.10	0.07	0.00	0.25	1.048
Observable	α_C	-0.15	0.07	-0.29	-0.01	0.999
	α_H	12.59	1.19	10.52	14.97	1.014
	α_d	-0.42	0.05	-0.51	-0.33	1.017
	α_{FT}	6.37	0.69	5.16	7.75	1.017
	σ_a	0.15	0.08	0.00	0.31	1.024

The posterior distribution provides useful information on the general performance

of the Miami Heat team in a game. Table 5.1 shows the posterior summary for the parameters and hyperparameters included in the model. Furthermore, we computed a standard convergence diagnostic measure, $\hat{R}$ (Gelman et al., 2014), for each parameter and hyperparameter. It is noteworthy that all $\hat{R}$ values were close to 1, which indicates good convergence of the MCMC algorithm.

First, we focus on the elements of the hidden model. Posterior means of the common intercepts included in the transition probabilities β_{CH} and β_{HC} clearly indicate a higher probability of moving from the *hot* state to the *cold* state than from the *cold* state to the *hot* state. In particular, the posterior mean (5.6) of the probability of switching from *C* to *H* in a generic game is 0.38, and from *H* to *C* is 0.59. Thus, remaining in the *cold* state in one transition is more likely than switching to the *hot*, or than remaining in the *hot* state. This can also be observed in Figure 5.2, which displays the posterior distribution of the transition probabilities for a generic game. Moreover, the expected posterior probability of being in the *cold* state when the game starts is around 0.55 according to the posterior mean of δ_C .

On the other hand, for the observed Bernoulli sub-process the posterior distribution of the common intercepts α_C and α_H shows a large difference of the magnitude of α_H with respect to α_C as seen in their means $E(\alpha_C \mid \mathcal{D}) = -0.15$ and $E(\alpha_H \mid \mathcal{D}) = 12.59$. Following the same procedure as above for the transition probabilities, we have a posterior expected probability of 0.46 of making a shot from a distance of 0 feet to the basket when the team is in the *cold* state, and almost 1 in the case of the *hot* state when the distance to the basket is 0.25 feet.

The two covariates considered in the modelling are relevant. The coefficient parameter associated with the distance to the basket, α_d , is completely negative, $E(\alpha_d \mid \mathcal{D}) = -0.42$, with a small posterior standard deviation, $SD(\alpha_d \mid \mathcal{D})=0.03$, which means that the probability of successful shooting decreases when the distance is further away. In addition, the credible interval of the coefficient associated with the *free throws* has a large positive expected value, $E(\alpha_{FT} \mid \mathcal{D}) = 6.37$. Therefore, making a *free throw* looks easier than making a *field goal* from that distance. Finally, the random




Figure 5.2: Posterior distribution histograms of the transition probabilities (a) $p^{(CC)}$ (in purple) and $p^{(CH)}$ (in yellow), and (b) $p^{(HC)}$ (in purple) and $p^{(HH)}$ (in yellow) integrating the random effects associated with the game.

effects associated with the games included in the success probability are also important since their associated standard deviation σ_a has a posterior mean of 0.15. It can be stated that a relevant portion of the variability can be attributed to the variations in the scoring efficiency of the team in the different games.

The information we can extract from the posterior distribution provides an immense number of possibilities to assess different aspects of the performance of the team. We can discuss some of them right now.

Occupancy times

Occupancy times admit many different posterior outputs. We only focus here on the posterior distribution of occupancy times in the *cold* and the *hot* state in a generic game in which the team attempted $n = 120$ shots. Figure 5.3 shows the posterior distribution of these times and two possibilities of initial state, *cold* and *hot*.

It is interesting to note that the team spends more time in the *cold* state than in the *hot* (with a posterior mean of 74.05 and 46.85, respectively), and that the state in which the team starts playing is practically irrelevant because there is hardly any difference between the two figures.




Figure 5.3: Posterior distribution histograms of the occupancy times in the states *hot* (in yellow) and *cold* (in purple) in a game with $n = 120$ shots that starts (a) in the *cold* and (b) in the *hot* state.

Stationary distribution

Although we could consider a stationary distribution of the chain associated with each game of the season, we focus on the posterior distribution $\pi((\Delta^{(C)}, \Delta^{(H)} | \theta) | \mathcal{D})$ of the stationary or limit distribution of the chain corresponding to a generic game. Table 5.2 shows a summary of the posterior distribution of the stationary distribution for the *cold* and *hot* state.

Table 5.2: Posterior summaries (mean, standard deviation and 95% credible interval) for the stationary distribution of the hidden Markov chain associated with a generic game.

	mean	sd	$q_{0.025}$	$q_{0.975}$
$\Delta^{(C)}$	0.61	0.02	0.57	0.65
$\Delta^{(H)}$	0.39	0.02	0.35	0.43

In summary, the probability of being in the *cold* state for the Miami Heat team in the 2005-2006 season was around 0.61. Thus, the probability of being in the *hot* state it was 0.39. Both distributions have very little variability.

Sojourn times

We consider a *cold* or a *hot* streak when the team stays in the same state for more than three shots. This number of shots is an arbitrary choice that can only be justified in order to illustrate the potential of our modelling to improve the understanding of the

behaviour of the team. Figure 5.4 shows a violin plot of the posterior distribution of the probability for a *cold* and for a *hot* streak in a generic game of the Miami Heat. There, we observe the probability of a *cold* streak (around 0.25) is nearly three times higher than the probability of a *hot* streak (less than 0.1). The much smaller amplitude of the *hot* distribution compared to the *cold* one is also evident.




Figure 5.4: Violin and box plot of the posterior probability distribution for a *cold* and a *hot* streak.

Probability of making a basket

The probability of making a basket depends on the state the team is currently visiting, the distance to the basket at which the shot is made, and whether the shot is a *free throw* or not. Figure 5.5 shows the posterior mean and a 95% credible interval for the probability of making a basket depending on the distance at which the shot (non-*free throw*) is made when the team is in the *cold* state, *hot* state or in the case where the state of the team is unknown.

When the team is in the *cold* state, one can see that for easy shots (i.e. shots close to the basket) there is a probability of around 0.5 of making a basket. However, from a distance of 10 feet, it is almost impossible to succeed. On the other hand, when the team is in the *hot* state, it is very likely to make a successful shot up to 15 feet. Then, from this distance, the probability starts to decrease slowly. Further, for a shot, in which the current state is unknown, the probability of success is also negatively

related to the distance to the basket. Shots closer to the basket have a probability of success around 0.7, whereas for intermediate shots this probability stabilises at around 0.5. Finally, the probability of making a three-point shot drops to 0.4.




Figure 5.5: Posterior means and 95% credible interval for the probabilities of making a basket depending on the distance, in feet, in three different scenarios. Purple curve represents the case when the team is in the *cold* state, the yellow plot is for the *hot* state, and the green plot is the general case when the state of the team is unknown.




Figure 5.6: Posterior mean and 95% credible interval for the probability of making a basket from 25 feet distance in each of the games of the season when the state of the team is unknown.

It is also interesting to visualise the team's performance over the different games of the season. In that sense, Figure 5.6 shows the mean and a 95th credible interval of the probability of hitting a shot from 25 feet, in which the team's state is unknown, in each of the games of the season. We observe a very marked regularity in all the games of the season, very little variability among games and a very stable behaviour in all phases of the season.

5.5 Conclusions

During our analysis of the Miami Heat season, we noticed something more like a *cold hand* instead of a *hot hand*. Based on our results, in any given game, the *cold* streaks seen to have more occurrences than *hot* streaks. Moreover, according to the posterior distribution of the transition probabilities, the two states are clearly differentiated: in fact, the team's performance is very different depending on which state it is in.

Our BLHMM can be employed as a useful tool for analysing the performance of a team or a player. It could even be used to compare the performances of different teams as it provides information that is not possible to obtain with traditional modelling. It might also be used to analyse the performance of an opponent's team during the course of a game.

Finally, for future research, the inclusion of a third intermediate state that is neither *hot* nor *cold* is an interesting possibility. In addition, the exploration of continuous-time Markov chains could be worthwhile to better model this phenomenon.

Chapter 6

Combining Bayesian modelling and optimisation: a wheelchair basketball tool.

6.1 Introduction

Wheelchair basketball is widely recognised as one of the most popular team sports for people with a physical disability. It is regulated by the International Wheelchair Basketball Federation (IWBF). According to its rules, a wheelchair basketball game is played on a standard basketball court with two teams consisting of five players each. Moreover, the sport is based on a Player Classification Points System (PCPS) that rates each player according to his or her playing capacity. The valid player classifications in the IWBF are 1.0, 1.5, 2.0, 2.5, 3.0, 3.5, 4.0 and 4.5. At no time in a game shall a team have players participating whose total points value exceed the 14-point limit (IWBF, 2021).

National leagues may slightly differ from international standards. In the Rollstuhlbasketball bundesliga (RBBL), the German league where teams consist of mixed genders, each female player receives a 1.5 on-court points bonus. It should be noted that there is a limit to the total score on the court, which is 17.5. Additionally, similar

bonuses may also be granted to younger or beginner players.

Choosing the optimal line-up for a game can be challenging in any sport, but it may be particularly difficult for sports that have a classification system for players. As stated by Gil-Agudo, Del Ama-Espinosa, and Crespo-Ruiz (2010), the majority of scientific research on sports for individuals with disabilities focuses on player classification. This study aims to develop a supplementary mathematical tool that provides coaches with the most effective wheelchair basketball team line-ups based on previous performance data.

To achieve this, we propose a novel methodology for team composition in wheelchair basketball, that can be generalised to sports that use PCPS. The proposed methodology involves first computing a Bayesian longitudinal model to analyse the temporal performance of the players. Then, using the knowledge provided by the posterior predictive distribution of this model, we define a linear optimisation problem that aims to maximise the performance of an ideal line-up team. Our proposal accounts for game constraints from the PCPS and the inherent uncertainty that arises from the posterior predictive distribution of the model. It is worth noting that the presence of Bayesian modelling in this proposal allows for the incorporation of uncertainty in the solution.

Finally, we illustrate this methodology by applying our proposal to the Doneck Dolphins Trier team, which participates in the first division of the RBBL. In particular, we consider three different performance measurements by player during the 2022-2023 Regular Season.

The remainder of this chapter is organised as follows: Section 6.2 provides a comprehensive description of the methodology proposed in this study. In Section 6.3, we delve into the practical application of our proposal within an actual team, highlighting its effectiveness. Lastly, in Section 6.4, we present the conclusions drawn from our research and outline key avenues for future research.

6.2 Methodology

6.2.1 Modelling the performance

Let $\mathbf{y}_i$ be the response vector which represents a performance variable measured over a period of time for a given player i . The entire vector of response variables is denoted by $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_N)$, where N is the total number of players studied. Additionally, considering the vector of parameters and hyperparameters $\boldsymbol{\theta}$, and the vector of random effects $\boldsymbol{\phi}$, we can express the longitudinal Bayesian model as in equation (1.4):

$$\begin{aligned} f(\mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\phi}) &= f(\mathbf{y} \mid \boldsymbol{\theta}, \boldsymbol{\phi}) f(\boldsymbol{\phi} \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta}) \\ &= \left[\prod_{i=1}^N f(\mathbf{y}_i \mid \boldsymbol{\theta}, \boldsymbol{\phi}) \right] f(\boldsymbol{\phi} \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta}). \end{aligned}$$

Recall that the individual sizes of $\mathbf{y}_i$ may vary. We denote n_i the number of performance measurements for the player i .

Depending on the nature of the problem, the selection of the appropriate model can vary, i. e., the exact form of $f(\mathbf{y}_i \mid \boldsymbol{\theta}, \boldsymbol{\phi})$ or $f(y_{ij} \mid \boldsymbol{\theta}, \boldsymbol{\phi})$.

6.2.2 Prediction for a new game

The second step of our proposal is to calculate predictions for all of the players studied. In the Bayesian framework, the posterior predictive distribution is used to compute a prediction. In the case we want to predict the punctual performance of all the players, $\mathbf{y}^{(pred)} = (y_1^{(pred)}, \dots, y_N^{(pred)})$ we calculate it by

$$f(\mathbf{y}^{(pred)} \mid \mathcal{D}) = \int f(\mathbf{y}^{(pred)} \mid \mathcal{D}, \boldsymbol{\theta}, \boldsymbol{\phi}) \pi(\boldsymbol{\theta}, \boldsymbol{\phi} \mid \mathcal{D}) d(\boldsymbol{\theta}, \boldsymbol{\phi}), \quad (6.1)$$

where $\mathcal{D}$ represents the observations of $\mathbf{y}$.

Therefore, to compute the posterior predictive distribution, it is necessary to know the posterior distribution of the parameters, hyperparameters and random effects, i. e., $\pi(\boldsymbol{\theta}, \boldsymbol{\phi} \mid \mathcal{D})$.

More specifically, we perform simulations to generate a certain number of predictions S which will be used in the subsequent step. The algorithm executed at iteration s , $s = 1, \dots, S$, comprises two distinct stages:

1. Simulate $\theta^{(s)}$ and $\phi^{(s)}$ from $\pi(\theta, \phi \mid \mathcal{D})$.
2. Generate $\mathbf{y}^{(pred)(s)}$ by simulating from $f(\mathbf{y}^{(pred)} \mid \mathcal{D}, \theta^{(s)}, \phi^{(s)})$,

where $\theta^{(s)}$ and $\phi^{(s)}$ are the s -th simulations from the posterior distribution $\pi(\theta, \phi \mid \mathcal{D})$, and $\mathbf{y}^{(pred)(s)} = (y_1^{(pred)(s)}, \dots, y_N^{(pred)(s)})$ is the s -th simulated vector of performance team predictions.

6.2.3 Selecting the optimal composition

The ultimate goal of this proposal is to select the optimal line-up based on the data collected and analysed for each player. After obtaining S realisations of the performance predictions for the players, the optimisation problem involves maximising the performance of the best line-up team at each iteration s ,

$$\max \sum_{i=1}^N z_i y_i^{(pred)(s)}, \quad (6.2)$$

where z_i indicates the presence of the player i in the team combination, i. e., this variable will be 1 if the player i is included in the line-up team and 0 otherwise.

There are two restrictions in this problem. First, just five players of the same team can play together on the court. The mathematical expression for this condition is the following expression

$$\sum_{i=1}^N z_i = 5. \quad (6.3)$$

The second constraint, which characterises wheelchair basketball, is related to the PCPS. As noted above, the maximum punctuation value allowed on the field by the IWBF is 14. Thus, this general restriction can be express as

$$\sum_{i=1}^N C_i z_i \leq 14, \quad (6.4)$$

where C_i is the “class” of player i , i. e., the functional punctuation of this player.

Finally, it is worth emphasising that solving this linear optimisation problem yields a realistic combination of players that optimises team performance. Moreover, since this problem is solved for each simulated predicted value (S times), we obtain a posterior probability distribution for all possible solutions.

6.3 Doneck Dolphins Trier case

6.3.1 Data

The methodology was applied to a wheelchair basketball team, the Doneck Dolphins Trier, which has been competing in the first division of the RBBL in Germany since 2003. The aim was to determine optimal line-ups for the team based on data from the 2022-2023 season, obtained from the RBBL website (<http://stats.rollstuhlbasketball.de/>) and using only data from players who played more than 40 minutes during the season. Table 6.1 shows general information for the nine players who met this criterion. Three of these players are women. One player has a functional classification of 1, another has 1.5, and yet another has 2, while three players have a functional classification of 3.5, and three more have 4.5, which is the highest possible value indicating a high level of capability in playing basketball.

To carry out the analysis, three different performance ratings were used. These measures are commonly used in official competitions and sports studies to evaluate player performance, although they are primarily focused on offensive performance.

One of the measures we used is the individual player efficiency (EFF), developed by Manley (1989). EFF is computed by adding several positive events and subtracting

Table 6.1: Name, functional classification value and sex of the players of the Doneck Dolphins Trier team who play more than 40 minutes during the 2022-2023 RBBL season.

Player	Classification	Sex	Index (<i>i</i>)
Annabel Breuer	1.5	Woman	1
Correy Rossi	2	Man	2
Dejon Green	3.5	Man	3
Dirk Passivan	4.5	Man	4
Lucas Jung	1	Man	5
Natalie Passivan	4.5	Woman	6
Patrick Dorner	3.5	Man	7
Svenja Erni	3.5	Woman	8
Walter Vlaanderen	4.5	Man	9

some negative events as follows:

$$\text{EFF} = (\text{points} + \text{rebounds} + \text{assists} + \text{steals} + \text{blocks}) - \\ (\text{missed field goals} + \text{missed free throws} + \text{turnovers}).$$

Another measure we used is the Performance Index Rating (PIR), which is employed in the Eurocup and many European leagues in the International Basketball Federation (FIBA) (Torres-Unda et al., 2016). PIR considers more events than EFF and includes fouls drawn, shots rejected, and personal fouls:

$$\text{PIR} = (\text{points} + \text{rebounds} + \text{assists} + \text{steals} + \text{blocks} + \text{fouls drawn}) - \\ (\text{missed field goals} + \text{missed free throws} + \text{turnovers} + \\ \text{shots rejected} + \text{personal fouls}).$$

The last performance metric utilised is the Win Score (Berri and Schmidt, 2010). Win Score replaces missed field goals and missed free throws with field goals attempted and free throws attempted, respectively. In addition, Win Score weights some

events differently than EFF and PIR:

$$\begin{aligned} \text{Win Score} = & (\text{points} + \text{rebounds} + \frac{1}{2}\text{assists} + \text{steals} + \frac{1}{2}\text{blocks}) - \\ & (\text{field goals attempted} + \frac{1}{2}\text{free throws attempted} + \\ & \text{turnovers} + \frac{1}{2}\text{personal fouls}). \end{aligned}$$

As a final step, we divided each player's respective metrics by the number of minutes they played to assess their individual performance in every game.


The observations of EFF, PIR, and Win Score per minute by player in each game of the 2022-2023 season are presented in Figure 6.1. It is evident that Dirk Passiwan ($i = 4$) performed exceptionally well compared to the other players, consistently ranking above them in most games. Additionally, the PIR metric exhibits greater variability both within and between players, while Win Score observations are noticeably less dispersed.

6.3.2 Applying the proposed methodology


As shown in Figure 6.1 and Table 6.1, $N = 9$ players are considered in this study. Moreover, the 2022-2023 RBBL season of the Doneck Dolphins Trier consisted on $n_i = 18$ games. The longitudinal variable y_{ij} represents the performance for the player i at game j , and we will assume that it is distributed following a mixed linear model as

$$\begin{aligned} f(y_{ij} \mid \boldsymbol{\theta}, \boldsymbol{\phi}) = & \mathcal{N}(\beta_0 + b_{0i} + b_{0j} + \beta_W I_W(i) + \beta_C C_i + \beta_H I_H(j) + \\ & (\beta_1 + b_{1i})j, \sigma^2), \end{aligned} \quad (6.5)$$


where the parameter vector contains all parameters and hyperparameters of the model, $\boldsymbol{\theta} = (\beta_0, \beta_W, \beta_C, \beta_H, \beta_1, \sigma)$, and $\boldsymbol{\phi}$ includes all random effects. Individual random effects b_{0i} and random slopes b_{1i} ($i = 1, \dots, 9$) are conditionally independent and normally distributed as $(b_{0i} \mid \sigma_0^2) \sim \mathcal{N}(0, \sigma_0^2)$ and $(b_{1i} \mid \sigma_1^2) \sim \mathcal{N}(0, \sigma_1^2)$. Moreover, the random effect associated with the game b_{0j} ($j = 1, \dots, 18$) is also conditionally



(a) EFF per minute.



(b) PIR per minute.



(c) Win score per minute.

Figure 6.1: Spaghetti plots of the three performance ratings per minute by player for the Doneck Dolphins Trier team during the 2022-2023 RBBL regular season.

normally distributed, $(b_{0j}|\sigma_{0m}^2) \sim \mathcal{N}(0, \sigma_{0m}^2)$. We consider as covariates the gender indicator variable $I_W(i)$, 1 when the player is female or 0 otherwise, the functional classification, C_i , and a dummy variable $I_H(j)$ indicating if the game was played at home (1) or not (0). Furthermore, to complete the Bayesian model, we consider a vague informative prior scenario, and select a wide normal distribution for β_0 , β_1 and for the coefficient parameters associated with the covariates, $\pi(\beta_0) = \pi(\beta_1) = \pi(\beta_W) = \pi(\beta_C) = \pi(\beta_H) = \mathcal{N}(0, 10^2)$. Finally, uniform distributions for the standard deviation parameters $\pi(\sigma) = \pi(\sigma_0) = \pi(\sigma_1) = \pi(\sigma_{0m}) = \mathcal{U}(0, 10)$ are assumed.

In preparation for a potential upcoming game after the completion of 18 games in the season, our objective is to select the optimal team line-up. To accomplish this task, we predict the performance of each player for game number 19 based on the posterior predictive distribution for $\mathbf{y}^{(pred)} = (y_{1\ 19}^{(pred)}, \dots, y_{9\ 19}^{(pred)})$ defined in expression (6.1).

As already mentioned, the RBBL bundesliga has some particularities with respect to the international regulations. To account for the RBBL bundesliga rules, we made some modifications to the class restriction constraint (6.4) in our optimisation problem. In particular, if there is no female player of the team on the basketball court, this constraint will be

$$\sum_{i=1}^N C_i z_i \leq 14.5. \quad (6.6)$$

On the other hand, when there is just one female player in the line-up team, the restriction is

$$\sum_{i=1}^N C_i z_i \leq 16, \quad (6.7)$$

and when there are two or more women

$$\sum_{i=1}^N C_i z_i \leq 17.5. \quad (6.8)$$

6.3.3 Posterior results

We approximate the posterior distribution of θ and ϕ using MCMC sampling methods through the JAGS software. Three parallel chains were executed for 100 000 iterations each, preceded by burn-ins of 300 000 iterations. In addition, to mitigate autocorrelation, the chains were thinned at every 100th iteration based on estimated autocorrelation in the sample. Consequently, we obtained $S = 3000$ performance predictions for each player. Finally, we employed the `lpSolve` package (Berkelaar et al., 2020) to solve the linear optimisation problem. The complete analysis, performed using R code (version 4.0.5), and the data are available as supplementary material at https://github.com/gcalvobayarri/Wheelchair_basketball_lineups.git.


The number of possible line-ups of the team with five players is 126, although not all of them will be valid because they will not meet constraints (6.6), (6.7), and (6.8). In particular, out of the 6 possible line-ups with an all-male team only 2 are valid line-ups; of the 45 potential line-ups with a team of four men and one woman, only 27 are acceptable; of the 60 possible teams with three men and two women only 48 meet the criterion (6.8), and all of the 15 teams with three women and two men meet the above constraint. Therefore, the number of possible line-ups is 92 and we can estimate the posterior probability that alignment L_k , $k = 1, \dots, 92$, is the optimal team through the quotient between the number of optimisation times that have resulted in line-up L_k and the number of optimisation problems solved.

According to the obtained results, the seven team line-ups with the highest probabilities are


$$\begin{aligned}
L_1 &= \{A. Breuer, C. Rossi, D. Passivan, P. Dorner, W. Vlaanderen\}, \\
L_2 &= \{A. Breuer, D. Passivan, P. Dorner, S. Erni, W. Vlaanderen\}, \\
L_3 &= \{A. Breuer, C. Rossi, D. Passivan, N. Passivan, W. Vlaanderen\}, \\
L_4 &= \{A. Breuer, C. Rossi, D. Green, D. Passivan, W. Vlaanderen\}, \\
L_5 &= \{A. Breuer, D. Green, D. Passivan, S. Erni, W. Vlaanderen\}, \\
L_6 &= \{C. Rossi, D. Passivan, L. Jung, S. Erni, W. Vlaanderen\}, \\
L_7 &= \{A. Breuer, C. Rossi, D. Passivan, S. Erni, W. Vlaanderen\}.
\end{aligned}$$

The bar plots of Figure 6.2 illustrates the probabilities of these seven compositions ($L_1, L_2, L_3, L_4, L_5, L_6$, and L_7) which are identified as the optimal line-ups based on EFF, PIR and Win Score metrics. Each bar represents the probability associated with a specific composition, indicating its probability of being the optimal choice according to the analysed performance metrics. The remarkable thing that one can observe is that, based on the three performance metrics, L_1, L_2 , and L_3 emerge as the top three optimal line-up teams, in that order. However, the ranking of L_4, L_5, L_6 , and L_7 differs when considering the Win Score metric in comparison to EFF and PIR. Finally, it is worth noting that Dirk Passivan and Walter Vlaanderen appear in all of the seven line-up teams, followed by Annabel Breuer who is included in six out of seven compositions.


Another valuable piece of information provided by the results is the probability of a player being included in the optimal team. We denote the event of including a specific player in the line-up as I_i , where $i = 1, \dots, 9$ corresponds to the indices of the players listed in Table 6.1. Figure 6.3 shows the probability of selecting each player in the team's best composition, considering the three performance metrics we evaluated. Upon observing this graph, it becomes apparent that three players, namely Dejon Green, Lucas Jung, and Natalie Passivan, have the lowest probabilities of being included in the optimal composition. Conversely, four players stand out above the rest



(a) EFF per minute.



(b) PIR per minute.



(c) Win score per minute.

Figure 6.2: Bar plots showing the estimated probabilities associated with the most optimal line-ups based on the three performance metrics for the seven compositions: L_1 , L_2 , L_3 , L_4 , L_5 , L_6 , and L_7 .

of the team: Annabel Breuer, Correy Rossi, Dirk Passivan, and Walter Vlaanderen.




Figure 6.3: Estimated probabilities $P(I_i \mid \mathcal{D})$ of player i being included in the optimal line-up team based on EFF, PIR and Win Score metrics.

We can also estimate the posterior probability of several players playing together. For example, the posterior probability that both Annabel Breuer and Natalie Passivan are in the line-up that maximises the Win Score metric is given by:

$$P(I_1, I_6 \mid \mathcal{D}) = 0.24.$$

Additionally, if we want to estimate the probability that these two players are included in the team composition, knowing that Correy Rossi is in the line-up, it is calculated as follows:

$$P(I_1, I_6 \mid I_2, \mathcal{D}) = \frac{P(I_1, I_6, I_2 \mid \mathcal{D})}{P(I_2 \mid \mathcal{D})} = \frac{0.23}{0.61} = 0.38.$$

The amount of information we can extract from our analysis could be highly relevant. A notable example may be the following. In Figure 6.3 the four players most likely to be selected for the line-up are identified: Annabel Breuer, Correy Rossi, Dirk Passivan and Walter Vlaanderen. Now, the question arises: which player from the

Doneck Dolphins Trier team would be the ideal complement to complete the composition? This question can be addressed in a simple way as follows:

$$P(I_i | I_1, I_2, I_4, I_9, \mathcal{D}) = \frac{P(I_i, I_1, I_2, I_4, I_9 | \mathcal{D})}{P(I_1, I_2, I_4, I_9 | \mathcal{D})}, \text{ with } i = \{3, 5, 6, 7, 8\},$$

where $P(I_i | I_1, I_2, I_4, I_9, \mathcal{D})$ represents the probability of player i being selected given that players 1, 2, 4, and 9 have already been chosen. Figure 6.4 shows a bar plot with the estimation of these probabilities considering EFF, PIR and Win Score metrics.




Figure 6.4: Estimated probabilities of the inclusion of players Dejon Green, Lucas Jung, Natalie Passivan, Patrick Dorner and Svenja Erni in the line-up team given that players Annabel Breuer, Correy Rossi, Dirk Passivan and Walter Vlaanderen have already been chosen, based on EFF, PIR and Win Score metrics.

Finally, upon examining Figure 6.4, it becomes evident that Patrick Dorner is the player who best completes the group. It is important to note that this team, consisting of $\{\text{Annabel Breuer, Correy Rossi, Dirk Passivan, Patrick Dorner, Walter Vlaanderen}\}$, corresponds to line-up L_1 . Additionally, the case of Natalie Passivan deserves special attention. On one hand, Figure 6.3 indicates that the probability of including this player in the best composition is one of the lowest. However, when considering the information provided by Figure 6.4, this player emerges as the second-best option to complete the most optimal line-up. This demonstrates the valuable information that can be derived from the results obtained through our methodology.

6.4 Conclusions

In this chapter we proposed a mathematical methodology that allows evaluating the best compositions for a wheelchair basketball team by using previous data of several players. Our proposal combines Bayesian longitudinal modelling, which analyses the performance data of players over time, with integer linear optimisation to determine the line-ups that maximise performance. This approach incorporates uncertainty throughout the entire process, resulting in a distribution of solutions rather than a single optimal solution to the optimisation problem.

The example we carried out using the data from a real wheelchair basketball team shows the mathematical and practical potential of our proposal. We clearly identify the most optimal composition of the Doneck Dolphins Trier team, and illustrate how to calculate the possible compatibilities or incompatibilities between players on the court. In our example, we focused exclusively on offensive metrics, but it could be supplemented with other types of metrics, such as those related to defense, as described by Oliver (2004).

To conclude, this methodology has the potential to be immensely valuable for coaches, providing valuable insights into the ideal compositions for wheelchair basketball teams in upcoming games. Furthermore, it can be readily applied to other team sports that cater to individuals with physical disabilities and incorporate a PCPS within their regulations. This includes powerchair football, powerchair hockey, and wheelchair rugby. Moreover, this methodology can also be adapted for other team sports that do not have a PCPS, by replacing the specific restrictions associated with the PCPS rule with realistic constraints present in sports such as football, rugby, basketball, and others.

Chapter 7

Final remarks

7.1 Final comments

This memory provides a summary of the work conducted during the intense period in which this thesis was developed. Undoubtedly, the primary motivation for this research was the many real-world problems encountered in the analysis of longitudinal and repeated measures data across various fields of knowledge.

On the applied side, we can highlight three notable findings. In Chapter 2, we demonstrated the relevance of both gender and age when studying the optimal intra-abdominal space of a patient undergoing laparoscopy. In Chapter 3, a decrease in the quantity of landed tonnes was documented for artisanal and industrial fishing in the Mediterranean Sea during the last three decades. Additionally, a mutually beneficial relationship between these two fishing sectors within each country was assessed. In Chapter 5, we observed something akin to a *cold hand* phenomenon during the 2005-2006 season of the NBA's Miami Heat, despite specifically investigating the opposite phenomenon.

However, the focus on practical problems generated interesting methodological contributions. We produced four methodological contributions in our work. First, in Chapter 3, we developed a joint longitudinal model that connects the responses of the same individual, establishing a linear relationship between the random effects as-

sociated with the individual of the different response variables. Then, in Chapter 4, we proposed a methodology for comparing longitudinal models with random effects and serial correlation terms by calculating their marginal likelihood through the power posterior method. Next, in Chapter 5, we demonstrated the ability of the Bayesian longitudinal hidden Markov models to describe the *hot hand* phenomenon in the shooting performance of a basketball team over a game, via a latent Markov chain with two possible states, i.e. *cold* and *hot*. Finally, in Chapter 6, we introduced an innovative approach employed in the realm of sports, wherein we integrate Bayesian longitudinal modelling techniques for examining the temporal performance data of players, alongside the utilisation of integer linear optimisation. This integrated framework enabled us to identify the optimal line-up configurations that maximised the overall performance.

Regarding the prior distribution of the unknown parameter vector θ , we followed the same strategy along all of the studies presented in this work. In general terms, we split the parameters and hyperparameters within θ into two subgroups: fixed effects (denoted as β 's or α 's throughout the manuscript) and standard deviations (usually represented by σ or similar symbols). Throughout the thesis, we consistently adopted the standard deviation as the parameter to quantify the variability of a normal distribution, rather than using precision or variance. This preference is based on the suggestion by Martínez-Beneito and Botella-Rocamora (2019). Thus, the standard deviation is measured on the same scale of the data, facilitating its interpretation and enhancing its intuitiveness. Moreover, whenever possible, we assumed independence among the various elements of θ , while maintaining a minimally informative prior scenario.

Specifically, the elicited marginal prior distribution for any fixed effect is a normal distribution with a large variance, typically set around 100 as $\mathcal{N}(0, 10^2)$. This choice is vaguely informative although proper, ensuring that the posterior distribution remains proper as well. This property enabled us to perform model selection by evaluating the marginal likelihoods of different models, as in Chapters 3 and 4. In this context, we avoided the use of improper uniform distributions, although widely accepted in the literature, to prevent potential issues related to posterior impropriety.

For the standard deviation parameters, our preferred choice for the prior marginal distributions was always a wide uniform distribution, either $\mathcal{U}(0, 10)$ or $\mathcal{U}(0, 100)$. An alternative prior distribution, particularly for precision, could be the gamma distribution (Gelman, 2006). However, Martínez-Beneito and Botella-Rocamora (2019) argue that the gamma distribution either exhibits undesirable behaviour (relative to its percentiles) or converges to an improper distribution, resulting in an improper posterior distribution.

Finally, we are convinced that transparency in scientific development is a necessary requirement for its continuous growth and improvement. It is worth mentioning the effort we made to ensure that. Whenever possible, all analyses we conducted are reproducible. For this reason, the codes we implemented and the data sets used in our analysis are freely available at <https://github.com/gcalvobayarri>.

7.2 Future research

Lastly, we can highlight some lines of research that we would like to address in the future or that we are already developing. As explained in Chapter 2, our goal was to develop a model that could accurately estimate the ideal space for performing laparoscopy based on patient information. In the second phase of our work, we are improving the previous model and collecting anthropometric measurements from new patients, such as sagittal abdominal diameter and waist-hip ratio. The aim of these enhancements is to identify more accurately the optimal intra-abdominal volume of the patient for conducting the surgery.

Chapter 4 presents a methodology for comparing longitudinal Bayesian models using the power posterior methodology. Specifically, these models include random effects and serial correlation terms. Our next goal is to compare our proposal with other existing approaches. For example, it would be very interesting to compare the results we might obtain through the R package `bridgesampling` with those presented here.

In terms of future investigations, it would also be interesting to consider the inclusion of a third intermediate state that is neither *hot* nor *cold* in the study of Chapter 5, where we developed a Bayesian longitudinal hidden Markov model to assess the *hot hand* phenomenon in shooting performance. Additionally, exploring continuous-time Markov chains could be valuable for a more accurate modelling of this phenomenon.

Finally, extending the methodology developed in Chapter 6 to other sports could be highly interesting. These sports may cater to individuals with disabilities, who would also benefit from a PCPS. However, we posit that the constraint system used in the linear optimisation problem can also be applied to team sports that do not employ a PCPS. For instance, in football, one could consider constraints such as the required number of players, the maximum and minimum number of forwards, and so on.

Appendix A

JAGS codes

A.1 Model Section 2.3 - Logistic growth mixed-effects model

```

model{
  for(i in 1 : N){
    for( j in 1 : individual_sizes[i]){
      y[i, j] ~ dnorm(Mi[i, j], lambda)
      Mi[i, j] <- a[i] / (1 + exp(-c[i] * IAP[i, j]-b[i]))
    }
  }

  ## PRIORS
  # sigma prior
  lambda <- 1 / (sigma * sigma)
  sigma ~ dunif(0,10)

  # a, b, c prior
  for(i in 1 : N){
    a[i] <- beta0[1] + u_a[i] + betaW[1] * sex[i] +
      betaA[1] * age[i]
    b[i] <- beta0[2] + u_b[i] + betaW[2] * sex[i] +
      betaA[2] * age[i]
    c[i] <- beta0[3] + betaW[3] * sex[i]

    # ADP
    ADP[i] <- -(log(5-2*sqrt(6)) + b[i]) / c[i]
  }

  # beta prior
  for(i in 1 : 2){

```

```

    betaA[i] ~ dnorm(0, 0.01)
  }

  for(i in 1 : 3){
    betaW[i] ~ dnorm(0, 0.01)
  }

  beta0[1] ~ dunif(0, 20)
  beta0[2] ~ dnorm(0, 0.01)
  beta0[3] ~ dunif(0, 10)

  for(i in 1 : N){
    u_a[i] ~ dnorm(0, Alambda)
    u_b[i] ~ dnorm(0, Blambda)
  }

  # sigma_a,b,c prior
  Alambda <- 1 / (sigmaA * sigmaA)
  sigmaA ~ dunif(0,10)

  Blambda <- 1 / (sigmaB * sigmaB)
  sigmaB ~ dunif(0,10)
}

```

A.2 Model Section 3.3 - $\mathcal{M}_1$

```

model{
  for (i in 1 : 10){
    for (t in 1 : T){
      # Industrial
      Y1[i, t] ~ dnorm(Mi1[i, t], lambda[1])
      Mi1[i, t] <- beta01 + b0[i, 1] +
        (beta11 + b1[i, 1]) * (t - 1)

      # Artisanal
      Y2[i, t] ~ dnorm(Mi2[i, t], lambda[2])
      Mi2[i, t] <- beta02 + b0[i, 2] +
        (beta12 + b1[i, 2]) * (t - 1)
    }
  }

  for (t in 1 : T){
    # Industrial Albania
    Y1[11, t] ~ dnorm(Mi1[11, t], lambda[1])

```

```

Mil[11, t] <- beta01 + b0Albania +
  (beta11 + b1Albania) * (t - 1)

# Artisanal Bosnia
Y2[11, t] ~ dnorm(Mi2[11, t], lambda[2])
Mi2[11, t] <- beta02 + b0Bosnia +
  (beta12 + b1Bosnia) * (t - 1)
}

## PRIORS
# lambda
lambda[1] <- 1 / (sigma[1] * sigma[1])
sigma[1] ~ dunif(0, 100)

lambda[2] <- 1 / (sigma[2] * sigma[2])
sigma[2] ~ dunif(0, 100)

# beta0, beta1
beta01 ~ dnorm(0, 0.01)
beta02 ~ dnorm(0, 0.01)
beta11 ~ dnorm(0, 0.01)
beta12 ~ dnorm(0, 0.01)

# b0, b1 prior
for (i in 1 : 10){
  b0[i, 1:2] ~ dmnorm(rep(0,2), Omega0[,])
  b1[i, 1:2] ~ dmnorm(rep(0,2), Omega1[,])
}

b0Albania ~ dnorm(0, Omega0[1, 1])
b0Bosnia ~ dnorm(0, Omega0[2, 2])

b1Albania ~ dnorm(0, Omega1[1, 1])
b1Bosnia ~ dnorm(0, Omega1[2, 2])

Omega0 <- inverse(Sigma0)
Omega1 <- inverse(Sigma1)

Sigma0[1, 1] <- sigma0[1] * sigma0[1]
Sigma0[2, 2] <- sigma0[2] * sigma0[2]
Sigma0[1, 2] <- rho0 * sigma0[1] * sigma0[2]
Sigma0[2, 1] <- Sigma0[1, 2]

Sigma1[1, 1] <- sigma1[1] * sigma1[1]
Sigma1[2, 2] <- sigma1[2] * sigma1[2]
Sigma1[1, 2] <- rho1 * sigma1[1] * sigma1[2]
Sigma1[2, 1] <- Sigma1[1, 2]

```

```

rho0 ~ dunif(-1, 1)
rho1 ~ dunif(-1, 1)

sigma0[1] ~ dunif(0, 100)
sigma0[2] ~ dunif(0, 100)

sigma1[1] ~ dunif(0, 100)
sigma1[2] ~ dunif(0, 100)
}

```

A.3 Model Section 3.3 - $\mathcal{M}_2$

```

model{
  for (i in 1 : 10){
    Y1[i,1] ~ dnorm(mu1[i, 1], lambda[1])
    mu1[i,1] <- beta01 + b0[i, 1]

    Y2[i,1] ~ dnorm(mu2[i, 1], lambda[2])
    mu2[i,1] <- beta02 + b0[i, 2]

    for (t in 2 : T){
      # Industrial
      Y1[i, t] ~ dnorm(Mi1[i, t], lambda[1])
      mu1[i, t] <- beta01 + b0[i, 1] +
        (beta11 + b1[i, 1]) * (t - 1)
      Mi1[i, t] <- mu1[i, t] +
        rho[1] * (Y1[i,t-1]-mu1[i,t-1])

      # Artisanal
      Y2[i, t] ~ dnorm(Mi2[i, t], lambda[2])
      mu2[i, t] <- beta02 + b0[i, 2] +
        (beta12 + b1[i, 2]) * (t - 1)
      Mi2[i, t] <- mu2[i, t] +
        rho[2] * (Y2[i,t-1]-mu2[i,t-1])
    }
  }

  Y1[11,1] ~ dnorm(mu1[11, 1], lambda[1]) #Albania
  mu1[11,1] <- beta01 + b0Albania

  Y2[11,1] ~ dnorm(mu2[11, 1], lambda[2]) #Bosnia
  mu2[11,1] <- beta02 + b0Bosnia
}

```

```

for (t in 2 : T){
  # Industrial Albania
  Y1[11, t] ~ dnorm(Mi1[11, t], lambda[1])
  mu1[11, t] <- beta01 + b0Albania +
    (beta11 + b1Albania) * (t - 1)
  Mi1[11, t] <- mu1[11, t] +
    rho[1] * (Y1[11,t-1]-mu1[11,t-1])

  # Artisanal Bosnia
  Y2[11, t] ~ dnorm(Mi2[11, t], lambda[2])
  mu2[11, t] <- beta02 + b0Bosnia +
    (beta12 + b1Bosnia) * (t - 1)
  Mi2[11, t] <- mu2[11, t] +
    rho[2] * (Y2[11,t-1]-mu2[11,t-1])
}

## PRIORS
# rho
rho[1] ~ dunif(-1, 1)
rho[2] ~ dunif(-1, 1)

# lambda
lambda[1] <- 1 / (sigma[1] * sigma[1])
sigma[1] ~ dunif(0, 100)

lambda[2] <- 1 / (sigma[2] * sigma[2])
sigma[2] ~ dunif(0, 100)

# mu prior
beta01 ~ dnorm(0, 0.01)
beta02 ~ dnorm(0, 0.01)
beta11 ~ dnorm(0, 0.01)
beta12 ~ dnorm(0, 0.01)

# b0, b1 prior
for (i in 1: 10){
  b0[i, 1:2] ~ dmnorm(rep(0,2), Omega0[,])
  b1[i, 1:2] ~ dmnorm(rep(0,2), Omega1[,])
}

b0Albania ~ dnorm(0, Omega0[1, 1])
b0Bosnia ~ dnorm(0, Omega0[2, 2])

b1Albania ~ dnorm(0, Omega1[1, 1])
b1Bosnia ~ dnorm(0, Omega1[2, 2])

```

```

Omega0 <- inverse(Sigma0)
Omega1 <- inverse(Sigma1)

Sigma0[1, 1] <- sigma0[1] * sigma0[1]
Sigma0[2, 2] <- sigma0[2] * sigma0[2]
Sigma0[1, 2] <- rho0 * sigma0[1] * sigma0[2]
Sigma0[2, 1] <- Sigma0[1, 2]

Sigma1[1, 1] <- sigma1[1] * sigma1[1]
Sigma1[2, 2] <- sigma1[2] * sigma1[2]
Sigma1[1, 2] <- rho1 * sigma1[1] * sigma1[2]
Sigma1[2, 1] <- Sigma1[1, 2]

rho0 ~ dunif(-1, 1)
rho1 ~ dunif(-1, 1)

sigma0[1] ~ dunif(0, 100)
sigma0[2] ~ dunif(0, 100)

sigma1[1] ~ dunif(0, 100)
sigma1[2] ~ dunif(0, 100)
}

```

A.4 Model Section 3.3 - $\mathcal{M}_3$

```

model{
  for (i in 1: 10){
    for (t in 1: T){
      # Industrial
      Y1[i, t] ~ dnorm(Mi1[i, t], lambda[1])
      Mi1[i, t] <- beta01 + b0[i, 1] +
        (beta11 + b1[i, 1]) * (t - 1) + w1[i, t]

      # Artisanal
      Y2[i, t] ~ dnorm(Mi2[i, t], lambda[2])
      Mi2[i, t] <- beta02 + b0[i, 2] +
        (beta12 + b1[i, 2]) * (t - 1) + w2[i, t]
    }
  }

  for (t in 1: T){
    # Industrial Albania
    Y1[11, t] ~ dnorm(Mi1[11, t], lambda[1])
    Mi1[11, t] <- beta01 + b0Albania +

```

```

(betal1 + b1Albania) * (t - 1) + w1[11, t]

# Artisanal Bosnia
Y2[11, t] ~ dnorm(Mi2[11, t], lambda[2])
Mi2[11, t] <- beta02 + b0Bosnia +
  (betal2 + b1Bosnia) * (t - 1) + w2[11, t]
}

for (i in 1 : 11){
  # Industrial
  w1[i, 1] ~ dnorm(0, lambdau[1])
  for (t in 2: T){
    w1[i, t] ~ dnorm(rho[1] * w1[i, t-1], lambdau[1])
  }

  # Artisanal
  w2[i, 1] ~ dnorm(0, lambdau[2])
  for (t in 2 : T){
    w2[i, t] ~ dnorm(rho[2] * w2[i, t-1], lambdau[2])
  }
}

## PRIORS
# lambda
lambda[1] <- 1 / (sigma[1] * sigma[1])
sigma[1] ~ dunif(0, 100)

lambda[2] <- 1 / (sigma[2] * sigma[2])
sigma[2] ~ dunif(0, 100)

# lambdau
lambdau[1] <- 1 / (sigmau[1] * sigmau[1])
sigmau[1] ~ dunif(0, 100)

lambdau[2] <- 1 / (sigmau[2] * sigmau[2])
sigmau[2] ~ dunif(0, 100)

# rho
rho[1] ~ dunif(-1, 1)
rho[2] ~ dunif(-1, 1)

# beta0, beta1
beta01 ~ dnorm(0, 0.01)
beta02 ~ dnorm(0, 0.01)
beta11 ~ dnorm(0, 0.01)
beta12 ~ dnorm(0, 0.01)

```

```

# b0, b1 prior
for (i in 1: 10){
  b0[i, 1:2] ~ dmnorm(rep(0,2), Omega0[,])
  b1[i, 1:2] ~ dmnorm(rep(0,2), Omega1[,])
}

b0Albania ~ dnorm(0, Omega0[1, 1])
b0Bosnia ~ dnorm(0, Omega0[2, 2])

b1Albania ~ dnorm(0, Omega1[1, 1])
b1Bosnia ~ dnorm(0, Omega1[2, 2])

Omega0 <- inverse(Sigma0)
Omega1 <- inverse(Sigma1)

Sigma0[1, 1] <- sigma0[1] * sigma0[1]
Sigma0[2, 2] <- sigma0[2] * sigma0[2]
Sigma0[1, 2] <- rho0 * sigma0[1] * sigma0[2]
Sigma0[2, 1] <- Sigma0[1, 2]

Sigma1[1, 1] <- sigma1[1] * sigma1[1]
Sigma1[2, 2] <- sigma1[2] * sigma1[2]
Sigma1[1, 2] <- rho1 * sigma1[1] * sigma1[2]
Sigma1[2, 1] <- Sigma1[1, 2]

rho0 ~ dunif(-1, 1)
rho1 ~ dunif(-1, 1)

sigma0[1] ~ dunif(0, 100)
sigma0[2] ~ dunif(0, 100)

sigma1[1] ~ dunif(0, 100)
sigma1[2] ~ dunif(0, 100)
}

```

A.5 Model Section 5.4 - Shooting performance model of the Miami Heat team

```

model{
  for (i in 1 : N){
    for (j in 1: n[i]){
      y[i, j] ~ dbern(gamma[i, j])
    }
  }
}

```

```

        gamma[i, j] <- exp(alpha0[Z[i, j]] +
        alphaft * ft[i, j] + alpha_d * dist[i, j] + a[i]) /
        (1 + exp(alpha0[Z[i, j]] + alphaft * ft[i, j] +
        alpha_d * dist[i, j] + a[i]))
    }
}

## PRIORS
# alpha prior
for (k in 1:2){
  alpha0.p[k] ~ dnorm(0, 0.01)
} # To avoid label switching:
alpha0[1] <- min(alpha0.p[])
alpha0[2] <- max(alpha0.p[])

alpha_d ~ dnorm(0, 0.01)
alphaft ~ dnorm(0, 0.01)

for (i in 1 : N){
  a[i] ~ dnorm(0, tau_a)
}

tau_a <- 1 / (sigmaa * sigmaa)

sigmaa ~ dunif(0, 10)

# Z prior
for (i in 1 : N){
  Z[i, 1] ~ dcat(delta0[])
}

delta0[1] ~ dbeta(1, 1)
delta0[2] <- 1 - delta0[1]

for (i in 1 : N){
  for (j in 2: n[i]){
    Z[i, j] ~ dcat(omega[Z[i, j-1], , i])
  }
}

for(i in 1 : N){
  omega[1, 1, i] <- 1 - pCH[i]
  omega[2, 2, i] <- 1 - pHC[i]
  omega[1, 2, i] <- pCH[i]
  omega[2, 1, i] <- pHC[i]

  pCH[i] <- 1 / (1 + exp(-(betaCH + bCH[i])))

```

```

    pHC[i] <- 1 / (1 + exp(-(betaHC + bHC[i])))

    bCH[i] ~ dnorm(0, tauCH)
    bHC[i] ~ dnorm(0, tauHC)
  }

  tauCH <- 1 / (sigmaCH * sigmaCH)
  sigmaCH ~ dunif(0, 10)
  betaCH ~ dnorm(0, 0.01)

  tauHC <- 1 / (sigmaHC * sigmaHC)
  sigmaHC ~ dunif(0, 10)
  betaHC ~ dnorm(0, 0.01)
}

```

A.6 Model Section 6.3 - Longitudinal performance model of the Doneck Dolphins Trier team

```

model{
  for (i in useful){
    # useful = players with more than 40 minutes
    Y[i] ~ dnorm(mu[i], lambda)
    mu[i] <- beta0 + b0[id_player[i]] + b0m[id_game[i]] +
      betaW * W[i] + betaC * Class[i] + betaH * home[i] +
      (beta1 + b1[id_player[i]]) * (id_game[i] - 1)
    # first match t=0
  }

  ## PRIORS
  # lambda
  lambda <- 1 / (sigma * sigma)
  sigma ~ dunif(0, 10)

  # betas prior
  beta0 ~ dnorm(0, 0.01)
  beta1 ~ dnorm(0, 0.01)
  betaC ~ dnorm(0, 0.01)
  betaW ~ dnorm(0, 0.01)
  betaH ~ dnorm(0, 0.01)

  # b0 prior
  for (i in 1: N){
    b0[i] ~ dnorm(0, lambda0)
  }
}

```

```

}
lambda0 <- 1 / (sigma0 * sigma0)
sigma0 ~ dunif(0, 10)

for (i in 1: (max(id_game))){
  b0m[i] ~ dnorm(0, lambda0m)
}
lambda0m <- 1 / (sigma0m * sigma0m)
sigma0m ~ dunif(0, 10)

# b1 prior
for (i in 1: N){
  b1[i] ~ dnorm(0, lambda1)
}
lambda1 <- 1 / (sigma1 * sigma1)
sigma1 ~ dunif(0, 10)
}

```

Appendix B

Technical details

B.1 Autoregressive variance-covariance matrix

Here, we give some details of the variance-covariance matrix used in model $\mathcal{M}_2$ of Chapter 3. For the sake of simplicity, we assume that the response variable is univariate. The model $\mathcal{M}_2$ can be formulated in an alternative way by conditioning y_{ij} on y_{ij-1} (Chi and Reinsel, 1989):

$$\begin{aligned} (y_{i1} \mid \boldsymbol{\theta}, \boldsymbol{\phi}_i) &= \mu_{i1} + \epsilon_{i1}, \\ (y_{ij} \mid y_{ij-1}, \boldsymbol{\theta}, \boldsymbol{\phi}_i) &= \mu_{ij} + \eta(y_{ij-1} - \mu_{ij-1}) + \epsilon_{ij}, \quad j = 2, \dots, n_i \end{aligned} \quad (\text{B.1})$$

where the measurement errors are normally distributed, denoted as $(\epsilon_{ij} \mid \sigma) \sim \mathcal{N}(0, \sigma^2)$, and $\mu_{ij} = \beta_0 + b_{0i} + t_{ij}(\beta_1 + b_{1i})$ for any measurement j of individual i .

Alternatively, this model can be expressed by means of an autoregressive process of the term $(y_{ij} - \mu_{ij})$ as follows

$$\begin{aligned} (y_{i1} - \mu_{i1}) &= \epsilon_{i1}, \\ (y_{ij} - \mu_{ij}) &= \eta(y_{ij-1} - \mu_{ij-1}) + \epsilon_{ij}, \quad j = 2, \dots, n_i. \end{aligned} \quad (\text{B.2})$$

This autoregressive process can be rewritten as a moving average process:

$$(y_{ij} - \mu_{ij}) = \sum_{k=1}^j \eta^{(t_{ij}-t_{ik})} \epsilon_{ik}. \quad (\text{B.3})$$

Thus, the variance of this term is

$$\text{V}(y_{ij} - \mu_{ij}) = \sum_{k=1}^j \eta^{2(t_{ij}-t_{ik})} \sigma^2, \quad (\text{B.4})$$

which is always finite in our case since the data of the study are balanced; in other words, $\mathbf{t} = (t_{i1}, \dots, t_{in_i}) = (0, 1, \dots, 34)$ for any country i . Additionally, this variance is the summation of a finite geometric series with a ratio value η^2 between 0 and 1.

Moving forward, the calculation of the elements of the variance-covariance matrix in Formula (3.5) is straightforward:

$$\text{Cov}(y_{ij}, y_{il} \mid \boldsymbol{\theta}, \boldsymbol{\phi}_i) = \sigma^2 \sum_{k=1}^{\min(j,l)} \eta^{(t_{ij}-t_{ik})+(t_{il}-t_{ik})}.$$

B.2 Complete posterior conditional distribution associated with the power posteriors

In this Section of the appendix, the complete posterior conditional densities required to compute the power posteriors via Gibbs sampling are listed. The conditional densities of the three most general models discussed in Chapter 4 ($\mathcal{M}_1$, $\mathcal{M}_2$, and $\mathcal{M}_3$ of the second simulation study) are fully described. Variability associated with the normal distributions is expressed in terms of the variance. Note that $\mathbf{w}_{-ij} = (w_{i1}, \dots, w_{ij-1}, w_{ij+1}, \dots, w_{in_i})$.

B.2.1 Model $\mathcal{M}_1$ - Study 2

$$\bullet \pi(\beta_0 \mid \mathbf{y}, \mathbf{b}, \mathbf{w}, \sigma, \tau) = \mathcal{N}\left(\frac{10^2 \tau \sum_{i=1}^N \sum_{j=1}^{n_i} (y_{ij} - (b_{1i} t_{ij} + w_{ij}))}{\sigma^2 + 10^2 \tau \sum_{i=1}^N n_i}, \frac{10^2 \sigma^2}{\sigma^2 + 10^2 \tau \sum_{i=1}^N n_i}\right),$$

- $\pi(b_{1i}|\mathbf{y}, \beta_0, \mathbf{w}_i, \sigma, \sigma_1, \tau) = \mathcal{N}\left(\frac{\sigma_1^2 \tau \sum_{j=1}^{n_i} t_{ij} (y_{ij} - (\beta_0 + w_{ij}))}{\sigma^2 + \sigma_1^2 \tau \sum_{j=1}^{n_i} t_{ij}^2}, \frac{\sigma_1^2 \sigma^2}{\sigma^2 + \sigma_1^2 \tau \sum_{j=1}^{n_i} t_{ij}^2}\right),$
- $\pi(w_{i1}|\mathbf{y}, \beta_0, \mathbf{w}_{-i1}, \rho, \sigma, \sigma_w, \tau) = \mathcal{N}\left(\frac{\sigma_w^2 \tau (y_{i1} - (\beta_0 + b_{1i} t_{i1})) + \sigma^2 \rho w_{i2}}{\sigma^2 + \sigma_w^2 \tau}, \frac{\sigma_w^2 \sigma^2}{\sigma^2 + \sigma_w^2 \tau}\right),$
- $\pi(w_{ij}|\mathbf{y}, \beta_0, \mathbf{w}_{-ij}, \rho, \sigma, \sigma_w, \tau) =$

$$\mathcal{N}\left(\frac{\sigma_w^2 \tau (y_{ij} - (\beta_0 + b_{1i} t_{ij})) + \sigma^2 \rho (w_{ij-1} + w_{ij+1})}{\sigma^2 (1 + \rho^2) + \sigma_w^2 \tau}, \frac{\sigma_w^2 \sigma^2}{\sigma^2 (1 + \rho^2) + \sigma_w^2 \tau}\right),$$
- $\pi(w_{in_i}|\mathbf{y}, \beta_0, \mathbf{w}_{-in_i}, \rho, \sigma, \sigma_w, \tau) = \mathcal{N}\left(\frac{\sigma_w^2 \tau (y_{in_i} - (\beta_0 + b_{1i} t_{in_i})) + \sigma^2 \rho w_{in_i-1}}{\sigma^2 + \sigma_w^2 \tau}, \frac{\sigma_w^2 \sigma^2}{\sigma^2 + \sigma_w^2 \tau}\right),$
- $\pi(\sigma|\mathbf{y}, \beta_0, \mathbf{b}, \mathbf{w}, \tau) \propto \frac{1}{\sigma^\tau \sum_{i=1}^N n_i} \exp\left(\frac{-\tau \sum_{i=1}^N \sum_{j=1}^{n_i} (y_{ij} - (\beta_0 + b_{1i} t_{ij} + w_{ij}))^2}{2\sigma^2}\right) \mathcal{U}(0, 10),$
- $\pi(\sigma_1|\mathbf{y}, \mathbf{b}) \propto \frac{1}{\sigma_1^N} \exp\left(\frac{-\sum_{i=1}^N b_{1i}^2}{2\sigma_1^2}\right) \mathcal{U}(0, 10),$
- $\pi(\sigma_w|\mathbf{y}, \rho, \mathbf{w}) \propto \frac{(1 - \rho^2)^{N/2}}{\sigma_w^{\sum_{i=1}^N n_i}} \exp\left(\frac{-(1 - \rho^2) \sum_{i=1}^N w_{i1}^2}{2\sigma_w^2}\right) \times$

$$\exp\left(\frac{-\sum_{i=1}^N \sum_{j=1}^{n_i} (w_{ij} - \rho w_{ij-1})^2}{2\sigma_w^2}\right) \mathcal{U}(0, 10),$$
- $\pi(\rho|\mathbf{y}, \sigma_w, \mathbf{w}) \propto \frac{(1 - \rho^2)^{N/2}}{\sigma_w^N} \exp\left(\frac{-(1 - \rho^2) \sum_{i=1}^N w_{i1}^2}{2\sigma_w^2}\right) \times$

$$\exp\left(\frac{-\sum_{i=1}^N \sum_{j=1}^{n_i} (w_{ij} - \rho w_{ij-1})^2}{2\sigma_w^2}\right) \mathcal{U}(-1, 1).$$

B.2.2 Model $\mathcal{M}_2$ - Study 2

- $\pi(\beta_0|\mathbf{y}, \mathbf{b}, \sigma, \tau) = \mathcal{N}\left(\frac{10^2 \tau \sum_{i=1}^N \sum_{j=1}^{n_i} (y_{ij} - (b_{0i} + b_{1i} t_{ij}))}{\sigma^2 + 10^2 \tau \sum_{i=1}^N n_i}, \frac{10^2 \sigma^2}{\sigma^2 + 10^2 \tau \sum_{i=1}^N n_i}\right),$
- $\pi(b_{0i}|\mathbf{y}, \beta_0, \mathbf{b}_1, \sigma, \sigma_0, \tau) = \mathcal{N}\left(\frac{\sigma_0^2 \tau \sum_{j=1}^{n_i} (y_{ij} - (\beta_0 + b_{1i} t_{ij}))}{\sigma^2 + \sigma_0^2 \tau n_i}, \frac{\sigma_0^2 \sigma^2}{\sigma^2 + \sigma_0^2 \tau n_i}\right),$
- $\pi(b_{1i}|\mathbf{y}, \beta_0, \mathbf{b}_0, \sigma, \sigma_1, \tau) = \mathcal{N}\left(\frac{\sigma_1^2 \tau \sum_{j=1}^{n_i} t_{ij} (y_{ij} - (\beta_0 + b_{0i}))}{\sigma^2 + \sigma_1^2 \tau \sum_{j=1}^{n_i} t_{ij}^2}, \frac{\sigma_1^2 \sigma^2}{\sigma^2 + \sigma_1^2 \tau \sum_{j=1}^{n_i} t_{ij}^2}\right),$
- $\pi(\sigma|\mathbf{y}, \beta_0, \mathbf{b}, \mathbf{w}, \tau) \propto \frac{1}{\sigma^\tau \sum_{i=1}^N n_i} \exp\left(\frac{-\tau \sum_{i=1}^N \sum_{j=1}^{n_i} (y_{ij} - (\beta_0 + b_{0i} + b_{1i} t_{ij}))^2}{2\sigma^2}\right) \mathcal{U}(0, 10),$
- $\pi(\sigma_0|\mathbf{y}, \mathbf{b}_0) \propto \frac{1}{\sigma_0^N} \exp\left(\frac{-\sum_{i=1}^N b_{0i}^2}{2\sigma_0^2}\right) \mathcal{U}(0, 10),$
- $\pi(\sigma_1|\mathbf{y}, \mathbf{b}_1) \propto \frac{1}{\sigma_1^N} \exp\left(\frac{-\sum_{i=1}^N b_{1i}^2}{2\sigma_1^2}\right) \mathcal{U}(0, 10).$

B.2.3 Model $\mathcal{M}_3$ - Study 2

- $\pi(\beta_0|\mathbf{y}, \rho, \mathbf{b}_1, \sigma, \tau) =$

$$\mathcal{N}\left(\frac{10^2\tau\left[(1-\rho^2)\sum_{i=1}^N(y_{i1}-b_{1i}t_{i1})-(1-\rho)\sum_{i=1}^N\sum_{j=2}^{n_i}(\rho y_{ij-1}-y_{ij}+(t_{ij}-\rho t_{ij-1})b_{1i})\right]}{\sigma^2+10^2\tau[(1-\rho^2)N+(1-\rho)^2\sum_{i=1}^N(n_i-1)]},\right.$$

$$\left.\frac{10^2\sigma^2}{\sigma^2+10^2\tau[(1-\rho^2)N+(1-\rho)^2\sum_{i=1}^N(n_i-1)]}\right),$$
- $\pi(b_{1i}|\mathbf{y}, \beta_0, \sigma, \sigma_1, \tau) =$

$$\mathcal{N}\left(\frac{\sigma_1^2\tau\left[t_{i1}(y_{i1}-\beta_0)-\sum_{j=2}^{n_i}(t_{ij}-\rho t_{ij-1})(\rho y_{ij-1}-y_{ij}+(1-\rho)\beta_0)\right]}{\sigma^2+\sigma_1^2\tau[(1-\rho^2)t_{i1}^2+\sum_{j=2}^{n_i}(t_{ij}-\rho t_{ij-1})^2]},\right.$$

$$\left.\frac{\sigma_1^2\sigma^2}{\sigma^2+\sigma_1^2\tau[(1-\rho^2)t_{i1}^2+\sum_{j=2}^{n_i}(t_{ij}-\rho t_{ij-1})^2]}\right),$$
- $f(\sigma|\mathbf{y}, \beta_0, \rho, \mathbf{b}_1, \tau) \propto \frac{1}{\sigma^\tau \sum_{i=1}^N n_i} \exp\left(-\frac{\tau(1-\rho^2)\sum_{i=1}^N(y_{i1}-(\beta_0+b_{1i}t_{i1}))^2}{2\sigma^2}\right) \times$

$$\exp\left(-\frac{\tau\sum_{i=1}^N\sum_{j=2}^{n_i}(y_{ij}-[(1-\rho)\beta_0+b_{1i}(t_{ij}-\rho t_{ij-1})+\rho y_{ij-1}])^2}{2\sigma^2}\right) \mathbf{U}(0, 10),$$
- $\pi(\sigma_1|\mathbf{y}, \mathbf{b}_1) \propto \frac{1}{\sigma_1^N} \exp\left(-\frac{\sum_{i=1}^N b_{1i}^2}{2\sigma_1^2}\right) \mathbf{U}(0, 10),$
- $\pi(\rho|\mathbf{y}, \beta_0, \mathbf{b}_1, \tau) \propto \frac{(1-\rho^2)^{N/2}}{\sigma^\tau \sum_{i=1}^N n_i} \exp\left(-\frac{\tau(1-\rho^2)\sum_{i=1}^N(y_{i1}-(\beta_0+b_{1i}t_{i1}))^2}{2\sigma^2}\right) \times$

$$\exp\left(-\frac{\tau\sum_{i=1}^N\sum_{j=1}^{n_i}(y_{ij}-[(1-\rho)\beta_0+b_{1i}(t_{ij}-\rho t_{ij-1})+\rho y_{ij-1}])^2}{2\sigma^2}\right) \mathbf{U}(-1, 1).$$

Bibliography

- Adler, P. (1981). *Momentum, a Theory of Social Action: A Theory of Social Action*. Sage Pubns.
- Albert, J. and J. Bennett (2001). “Streakiness (or, The Hot Hand)”. *Curve Ball*. Springer, pp. 111–144.
- Albo-Puigserver, M., J. Navarro, M. Coll, J. Aguzzi, L. Cardona, and R. Sáez-Liante (2015). “Feeding ecology and trophic position of three sympatric demersal chondrichthyans in the northwestern Mediterranean”. *Marine Ecology Progress Series* 524, pp. 255–268.
- Armero, C., A. Forte, H. Perpiñán, M. J. Sanahuja, and S. Agustí (2018). “Bayesian joint modeling for assessing the progression of chronic kidney disease in children”. *Statistical Methods in Medical Research* 27.1, pp. 298–311.
- Bayes, T. (1763). “LII. An essay towards solving a problem in the doctrine of chances. By the late Rev. Mr. Bayes, FRS communicated by Mr. Price, in a letter to John Canton, AMFR S”. *Philosophical transactions of the Royal Society of London* 53, pp. 370–418.
- Berger, J. O. and L. R. Pericchi (1996). “The Intrinsic Bayes Factor for Model Selection and Prediction”. *Journal of the American Statistical Association* 91.433, pp. 109–122.
- Berkelaar, M. et al. (2020). *lpSolve: Interface to ‘Lp-Solve’ v. 5.5 to Solve Linear-Integer Programs*. R package version 5.6.15. URL: <https://CRAN.R-project.org/package=lpSolve>.

- Berri, D. and M. Schmidt (2010). *Stumbling on Wins (Bonus Content Edition): Two Economists Expose the Pitfalls on the Road to Victory in Professional Sports, Portable Documents*. Pearson Education.
- Box, G. E. P. (1980). “Sampling and Bayes’s inference in scientific modelling and robustness (with discussion)”. *Journal of the Royal Statistical Society, Series A* 143, pp. 383–430.
- Brigolin, D., P. Girardi, P. Miller, W. Xu, D. Nachite, M. Zucchetta, and F. Pranovi (2018). “Using remote sensing indicators to investigate the association of landings with fronts: Application to the Alboran Sea (western Mediterranean Sea)”. *Fisheries Oceanography* 27.5, pp. 408–416.
- Brosset, P., J.-M. Fromentin, E. Van Beveren, J. Lloret, V. Marques, G. Basilone, A. Bonanno, P. Carpi, F. Donato, V. Č. Keč, et al. (2017). “Spatio-temporal patterns and environmental controls of small pelagic fish body condition from contrasted Mediterranean areas”. *Progress in oceanography* 151, pp. 149–162.
- Calvo, G., C. Armero, V. Gómez-Rubio, and G. Mazzinari (2021). “Bayesian hierarchical nonlinear modelling of intra-abdominal volume during pneumoperitoneum for laparoscopic surgery”. *SORT* 45.2, pp. 1–20.
- Casella, G. and C. Robert (Jan. 2010). *Introducing Monte Carlo Methods with R (Use R!)* ISBN: 978-1-4419-1582-5. DOI: 10.1007/978-1-4419-1576-4.
- Checkley Jr, D. M., R. G. Asch, and R. R. Rykaczewski (2017). “Climate, anchovy, and sardine”. *Annual Review of Marine Science* 9, pp. 469–493.
- Chen, M.-H., Q.-M. Shao, and J. G. Ibrahim (2012). *Monte Carlo Methods in Bayesian Computation*. New York: Springer Science & Business Media.
- Chi, E. M. and G. C. Reinsel (1989). “Models for longitudinal data with random effects and AR(1) errors”. *Journal of the American Statistical Association* 84.406, pp. 452–459.
- Chib, S. (1995). “Marginal likelihood from the Gibbs output”. *Journal of the American Statistical Association* 90.432, pp. 1313–1321.

- Chib, S. and I. Jeliazkov (2001). “Marginal likelihood from the Metropolis–Hastings output”. *Journal of the American Statistical Association* 96.453, pp. 270–281.
- Coll, M., M. Albo-Puigserver, J. Navarro, I. Palomera, and J. M. Dambacher (2019). “Who is to blame? Plausible pressures on small pelagic fish population changes in the northwestern Mediterranean Sea”. *Marine Ecology Progress Series* 617, pp. 277–294.
- Coll, M., I. Palomera, S. Tudela, and M. Dowd (2008). “Food-web dynamics in the south Catalan Sea ecosystem (NW Mediterranean) for 1978–2003”. *Ecological Modelling* 217.1-2, pp. 95–116.
- Cramer, J. S. (2004). “The early origins of the logit model”. *Studies in History and Philosophy of Biological and Biomedical Sciences* 35.4, pp. 613–626.
- Curtis, S. M., I. Goldin, and E. Evangelou (2018). “Package ‘mcmcplots’”.
- Davidian, M. (2008). “Non-linear mixed-effects model”. *Longitudinal data analysis*. Chapman and Hall/CRC. p, pp. 121–156.
- Denwood, M. J. (2016). “runjags: An R package providing interface utilities, model templates, parallel computing methods and additional distributions for MCMC models in JAGS”. *Journal of statistical software* 71, pp. 1–25.
- Díaz-Cambronero, O., B. Flor Lorente, G. Mazzinari, M. Vila Montañes, N. García Gregorio, D. Robles Hernandez, L. E. Olmedilla Arnal, M. P. Argente Navarro, M. J. Schultz, and Erran (2019). “A multifaceted individualized pneumoperitoneum strategy for laparoscopic colorectal surgery: a multicenter observational feasibility study”. *Surgical endoscopy* 33, pp. 252–260.
- Díaz-Cambronero, O., G. Mazzinari, B. Flor Lorente, N. García Gregorio, D. Robles-Hernandez, L. Olmedilla Arnal, A. Martin de Pablos, M. Schultz, C. Errando, and M. Argente Navarro (2020). “Effect of an individualized versus standard pneumoperitoneum pressurestrategy onpostoperative recovery: a randomized clinical trial in laparoscopic colorectal surgery”. *British Journal of Surgery* 107.12, pp. 1605–1614.

- Diggle, P. J., P. Heagerty, K.-Y. Liang, and S. L. Zeger (2002). *Analysis of longitudinal data*. 2nd ed. Oxford University Press.
- Efron, B. and T. Hastie (2021). *Computer age statistical inference, student edition: algorithms, evidence, and data science*. Vol. 6. Cambridge University Press.
- Fan, Y., R. Wu, M.-H. Chen, L. Kuo, and P. O. Lewis (2011). “Choosing among partition models in Bayesian phylogenetics”. *Molecular biology and evolution* 28.1, pp. 523–532.
- Fang, Z. and R. L. Bailey (2001). “Nonlinear Mixed Effects Modeling for Slash Pine Dominant Height Growth Following Intensive Silvicultural Treatments”. *Forest Science* 47.3, pp. 287–300.
- Forstemann, T., J. Trzewik, J. Holste, B. Batke, M. A. Konerding, T. Wolloscheck, and C. Hartung (2011). “Forces and deformations of the abdominal wall—a mechanical and geometrical approach to the linea alba”. *Journal of Biomechanics* 44.4, pp. 600–606.
- Friel, N., M. Hurn, and J. Wyse (2014). “Improving power posterior estimation of statistical evidence”. *Statistics and Computing* 24, pp. 709–723.
- Friel, N. and A. N. Pettitt (2008). “Marginal likelihood estimation via power posteriors”. *Journal of the Royal Statistical Society Series B* 70.3, pp. 589–607.
- Frühwirth-Schnatter, S. (2001). “Markov chain Monte Carlo estimation of classical and dynamic switching and mixture models”. *Journal of the American Statistical Association* 96.453, pp. 194–209.
- Funatogawa, I. and T. Funatogawa (2018). “Nonlinear Mixed Effects Models, Growth Curves, and Autoregressive Linear Mixed Effects Models”. *Longitudinal Data Analysis*. Ed. by J. R. S. in Statistics. Singapore: Springer, pp. 99–117.
- Gelfand, A. E. (1996). “Model determination using sampling-based methods”. *Markov Chain Monte Carlo in Practice*. Ed. by S. R. W. Gilks and D. Spiegelhalter. Chapman & Hall/CRC Interdisciplinary Statistics, pp. 145–161.

- Gelfand, A. E. and D. K. Dey (1994). “Bayesian model choice: asymptotics and exact calculations”. *Journal of the Royal Statistical Society: Series B (Methodological)* 56.3, pp. 501–514.
- Gelfand, A. E., D. K. Dey, and H. Chang (1992). *Model determination using predictive distributions with implementation via sampling-based methods*. Tech. rep. Stanford Univ CA Dept of Statistics.
- Gelfand, A. E. and A. F. Smith (1990). “Sampling-based approaches to calculating marginal densities”. *Journal of the American statistical association* 85.410, pp. 398–409.
- Gelman, A., J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. Rubin (2014). *Bayesian Data Analysis*. Third. Boca Raton: Chapman and Hall.
- Gelman, A. and X.-L. Meng (1998). “Simulating normalizing constants: from importance sampling to bridge sampling to path sampling”. *Statistical Science* 13.2, pp. 163–185.
- Gelman, A. (2006). “Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper)”.
- Gil-Agudo, A., A. Del Ama-Espinosa, and B. Crespo-Ruiz (2010). “Wheelchair basketball quantification”. *Physical Medicine and Rehabilitation Clinics* 21.1, pp. 141–156.
- Gilovich, T., R. Vallone, and A. Tversky (1985). “The hot hand in basketball: On the misperception of random sequences”. *Cognitive psychology* 17.3, pp. 295–314.
- Gronau, Q. F., H. Singmann, and E.-J. Wagenmakers (2020). “bridgesampling: An R Package for Estimating Normalizing Constants”. *Journal of Statistical Software* 92.10, pp. 1–29.
- Hedeker, D. and R. D. Gibbons (2006). *Longitudinal data analysis*. Vol. 451. John Wiley & Sons.
- Hoeting, J. A., D. Madigan, A. E. Raftery, and C. T. Volinsky (1999). “Bayesian model averaging: a tutorial (with comments by M. Clyde, D. Draper and E. I. George, and a rejoinder by the authors)”. *Statistical Science* 14.4, pp. 382–417.

- Ihaka, R. and R. Gentleman (1996). “R: a language for data analysis and graphics”. *Journal of computational and graphical statistics* 5.3, pp. 299–314.
- IWBF (2021). *Official Wheelchair Basketball Rules*.
- Kass, R. E. and Raftery (1995). “Bayes Factors”. *Journal of the American Statistical Association* 90.2, pp. 773–795.
- Kulkarni, V. G. (2016). *Modeling and analysis of stochastic systems*. Chapman and Hall/CRC.
- Laird, N. M. and J. H. Ware (1982). “Random-Effects for Longitudinal Data”. *Biometrics* 38, pp. 963–974.
- Lartillot, N. and H. Philippe (2006). “Computing Bayes Factors using thermodynamic integration”. *Systematic Biology* 55, pp. 195–207.
- Lee, Y. and J. A. Nelder (2004). “Conditional and marginal models”. *Statistical Science* 19.2, pp. 219–238.
- Lindsey, J. K. (2001). *Nonlinear Models in Medical Statistics*. Oxford: Oxford University Press.
- Lindstrom, M. J. and D. M. Bates (1990). “Nonlinear mixed effects models for repeated measures data”. *Biometrics* 46, pp. 673–687.
- Lunn, D. J., A. Thomas, N. Best, and D. Spiegelhalter (2000). “WinBUGS-a Bayesian modelling framework: concepts, structure, and extensibility”. *Statistics and computing* 10, pp. 325–337.
- Manley, M. (1989). *Basketball Heaven*. Doubleday Books.
- Martínez-Beneito, M. A. and P. Botella-Rocamora (2019). *Disease mapping: from foundations to multidimensional modeling*. CRC Press.
- Mazzinari, G., O. Diaz-Cambronero, J. M. Alonso-Iñigo, N. Garcia-Gregorio, B. Ayas-Montero, J. L. Ibañez, A. Serpa Neto, L. Ball, M. Gama de Abreu, and P. Pelosi (2020). “Intraabdominal pressure targeted positive end-expiratory pressure during laparoscopic surgery: an open-label, nonrandomized, crossover, clinical trial”. *Anesthesiology* 132.4, pp. 667–677.

- Mazzinari, G., O. Diaz-Cambronero, A. Serpa Neto, A. C. Martínez, L. Rovira, M. P. Argente Navarro, M. L. Malbrain, P. Pelosi, M. Gama de Abreu, and M. W. Hollmann (2021). “Modeling intra-abdominal volume and respiratory driving pressure during pneumoperitoneum insufflation—a patient-level data meta-analysis”. *Journal of Applied Physiology* 130.3, pp. 721–728.
- McCulloch, R. E. and R. S. Tsay (1994). “Statistical analysis of economic time series via Markov switching models”. *Journal of Time Series Analysis* 15.5, pp. 523–539.
- Mellon-Duval, C., M. Harmelin-Vivien, L. Métral, V. Loizeau, S. Mortreux, D. Roos, and J. M. Fromentin (2017). “Trophic ecology of the European hake in the Gulf of Lions, northwestern Mediterranean Sea”. *Scientia Marina* 81.1, pp. 7–18.
- Meng, X.-. and W. Wong (1996). “Simulating ratios of normalizing constants via a simple identity: a theoretical exploration”. *Statistica Sinica* 6.4, pp. 831–860.
- Meng, X.-L. and S. Schilling (2002). “Warp Bridge Sampling”. *Journal of Computational and Graphical Statistics* 11, pp. 552–586.
- Miller, J. B. and A. Sanjurjo (2018). “Surprised by the hot hand fallacy? A truth in the law of small numbers”. *Econometrica* 86.6, pp. 2019–2047.
- Mulier, J., B. Dillemans, M. Crombach, C. Missant, and A. Sels (2009). “On the abdominal pressure volume relationship”. *The Internet Journal of Anesthesiology* 21.1, pp. 5221–5231.
- Navarro, J., R. Sáez-Liante, M. Albo-Puigserver, M. Coll, and I. Palomera (2017). “Feeding strategies and ecological roles of three predatory pelagic fish in the western Mediterranean Sea”. *Deep Sea Research Part II: Topical Studies in Oceanography* 140, pp. 9–17.
- NBA (2020). *NBAstuffer*. <https://www.nbastuffer.com/analytics101/playbyplay-data/>. Accessed: 2022-05-03.
- Neal, R. M. (2001). “Annealed importance sampling”. *Statistics and computing* 11.2, pp. 125–139.
- Neudecker, J., S. Sauerland, E. A. Neugebauer, R. Bergamaschi, H. J. Bonjer, A. Cuschieri, K.-H. Fuchs, C. A. Jacobi, F. Jansen, and A.-M. Koivusalo (2006).

- “The EAES clinical practice guidelines on the pneumoperitoneum for laparoscopic surgery”. *EAES Guidelines for Endoscopic Surgery: Twelve Years Evidence-Based Surgery in Europe* 16.7, pp. 39–85.
- Neugebauer, E. A., M. Becker, G. F. Buess, A. Cuschieri, H.-P. Dauben, A. Fingerhut, K. H. Fuchs, B. Habermalz, L. Lantsberg, and M. Morino (2010). “EAES recommendations on methodology of innovation management in endoscopic surgery”. *Surgical Endoscopy* 24.7, pp. 1594–1615.
- Newton, M. A. and Raftery (1994). “Approximate Bayesian inference with the weighted likelihood bootstrap”. *Journal of the Royal Statistical Society: Series B (Methodological)* 56.1, pp. 3–26.
- Ntzoufras, I. (2011). *Bayesian modeling using WinBUGS*. John & Sons: Wiley.
- Obermeyer, Z., B. Powers, C. Vogeli, and S. Mullainathan (2019). “Dissecting racial bias in an algorithm used to manage the health of populations”. *Science* 6464.366, pp. 447–453.
- Oliver, D. (2004). *Basketball on paper: rules and tools for performance analysis*. Potomac Books, Inc.
- Ötting, M., R. Langrock, C. Deutscher, and V. Leos-Barajas (2020). “The hot hand in professional darts”. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 183.2, pp. 565–580.
- Pache, B., M. Hübner, J. Jurt, N. Demartines, and F. Grass (2017). “Minimally invasive surgery and enhanced recovery after surgery: the ideal combination?” *Journal of Surgical Oncology* 116.5, pp. 613–616.
- Peek, M. S., E. Russek-Cohen, D. A. Wait, and I. N. Forseth (2002). “Physiological response curve analysis using nonlinear mixed models”. *Oecologia* 132, pp. 175–180.
- Pennino, M.-G., J. M. Bellido, D. Conesa, M. Coll, and E. Tortosa-Ausina (2017). “The analysis of convergence in ecological indicators: an application to the Mediterranean fisheries”. *Ecological Indicators* 78, pp. 449–457.

- Pinheiro, J. and D. M. Bates (2000). *Mixed-Effects Models in S and S-Plus*. New York: Springer-Verlag.
- Plummer, M. (2003). “JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling”. *Proceedings of the 3rd international workshop on distributed statistical computing*. Vienna, Austria., pp. 1–10.
- (2008). “Penalized loss functions for Bayesian model comparison”. *Biostatistics* 9.3, pp. 523–539.
- R Core Team (2021). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. Vienna, Austria. URL: <https://www.R-project.org/>.
- Raftery, A. E., M. A. Newton, J. M. Satagopan, and P. N. Krivitsky (2007). “Estimating the integrated likelihood via posterior simulation using the harmonic mean identity”. *Bayesian Statistics*. Vol. 8. Oxford University Press, pp. 1–45.
- Rue, H., S. Martino, and N. Chopin (2009). “Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations”. *Journal of the royal statistical society: Series b (statistical methodology)* 71.2, pp. 319–392.
- Sandri, M., P. Zuccolotto, and M. Manisera (2020). “Markov switching modelling of shooting performance variability and teammate interactions in basketball”. *Journal of the Royal Statistical Society: Series C* 69.5, pp. 1337–1356.
- Saraux, C., J.-M. Fromentin, J.-L. Bigot, J.-H. Bourdeix, M. Morfin, D. Roos, E. Van Beveren, and N. Bez (2014). “Spatial structure and distribution of small pelagic fish in the northwestern Mediterranean Sea”. *PloS One* 9.11.
- Senn, S. (2016). “Mastering variation: variance components personalised medicine”. *Statistics in Medicine* 30.7, pp. 966–977.
- Skilling, J. (2006). “Nested sampling for general Bayesian computation”. *Bayesian Analysis* 1.4, pp. 833–859.
- Spezia, L. (2009). “Reversible jump and the label switching problem in hidden Markov models”. *Journal of Statistical Planning and Inference* 139.7, pp. 2305–2315.

- Spiegelhalter, D., A. Thomas, N. Best, and D. Lunn (2014). *OpenBUGS user manual*. Version 3.2.3. 87 pp.
- Spiegelhalter, D. J., N. G. Best, B. P. Carlin, and A. Van Der Linde (2002). “Bayesian measures of model complexity and fit”. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 64.4, pp. 583–639.
- Stergiou, K. I., S. Somarakis, G. Triantafyllou, K. P. Tsiaras, M. Giannoulaki, G. Petihakis, A. Machias, and A. C. Tsikliras (2016). “Trends in productivity and biomass yields in the Mediterranean Sea large marine ecosystem during climate change”. *Environmental Development* 17, pp. 57–74.
- Stigler, S. M. (1986). “Laplace’s 1774 memoir on inverse probability”. *Statistical Science* 1.3, pp. 359–363.
- Sun, Y. (2004). “Detecting the hot hand: An alternative model”. *Proceedings of the Annual Meeting of the Cognitive Science Society*. Vol. 26(26).
- Tanner, M. A. (2012). *Tools for Statistical Inference*. Springer.
- Tierney, L. and J. B. Kadane (1986). “Accurate approximations for posterior moments and marginal densities”. *Journal of the American Statistical Association* 81.393, pp. 82–86.
- Torres-Unda, J., I. Zarrazquin, L. Gravina, J. Zubero, J. Seco, S. M. Gil, J. Gil, and J. Irazusta (2016). “Basketball performance is related to maturity and relative age in elite adolescent players”. *The Journal of Strength & Conditioning Research* 30.5, pp. 1325–1332.
- Vehtari, A. and J. Ojanen (2012). “A survey of Bayesian predictive methods for model assessment, selection and comparison”. *Statistics Surveys* 6, pp. 142–228.
- Verbeke, G. and M. Davidian (2009). “Joint models for longitudinal data: Introduction and overview”. *Longitudinal Data Analysis*. Ed. by G. V. G. Fitzmaurice M. Davidian and G. Molenberghs. Boca Raton: Chapman and Hall/CRC. Chap. 13, pp. 319–326.
- Weiss, R. E. (2005). *Modeling longitudinal data*. Vol. 1. Springer.

- Wetzels, R., D. Tutschkow, C. Dolan, S. van der Sluis, G. Dutilh, and E.-J. Wagenmakers (2016). “A Bayesian test for the hot hand phenomenon”. *Journal of Mathematical Psychology* 72, pp. 200–209.
- Xie, W., P. O. Lewis, Y. Fan, L. Kuo, and M.-H. Chen (2011). “Improving marginal likelihood estimation for Bayesian phylogenetic model selection”. *Systematic biology* 60.2, pp. 150–160.
- Zeller, D. and D. Pauly (2016). “Catch reconstruction: concepts, methods, and data sources”. *Global Atlas of Marine Fisheries: A Critical Appraisal of Catches and Ecosystem Impacts*, pp. 12–33.
- Zuccolotto, P. and M. Manisera (2020). *Basketball data science: with applications in R*. CRC Press.