



HAL
open science

Inference of demographic histories in structured populations with an application to human evolution.

Armando Arredondo Soto

► **To cite this version:**

Armando Arredondo Soto. Inference of demographic histories in structured populations with an application to human evolution.. General Mathematics [math.GM]. INSA de Toulouse, 2021. English. NNT : 2021ISAT0036 . tel-03675193

HAL Id: tel-03675193

<https://theses.hal.science/tel-03675193>

Submitted on 23 May 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



THÈSE

En vue de l'obtention du

DOCTORAT DE L'UNIVERSITÉ DE TOULOUSE

Délivré par : *l'Institut National des Sciences Appliquées de Toulouse (INSA de Toulouse)*

Présentée et soutenue le 07/12/2021 par :

Armando ARREDONDO SOTO

Inférence d'histoires démographiques de populations structurées et application à l'évolution humaine.

Inference of demographic histories in structured populations with an application to human evolution.

JURY

BÉATRICE LAURENT-BONNEAU	Professeure d'Université	Présidente du Jury
ASGER HOBOLTH	Professeur d'Université	Rapporteur
GUILLAUME ACHAZ	Professeur d'Université	Rapporteur
HILDE WILKINSON-HERBOTS	Professeure d'Université	Membre du Jury
LOUNÈS CHIKHI	Directeur de Recherche	Co-directeur
OLIVIER MAZET	Maître de Conférences	Co-directeur

École doctorale et spécialité :

MITT : Domaine Mathématiques : Mathématiques appliquées

Unité de Recherche :

Institut de Mathématiques de Toulouse (UMR 5219)

Directeur(s) de Thèse :

Lounès CHIKHI et Olivier MAZET

Rapporteurs :

Asger HOBOLTH et Guillaume ACHAZ

Abstract

Inferring the demographic history of species is one of the greatest challenges in population genetics. Such histories are often represented as histories of size changes, ignoring population structure. Alternatively, when structure is assumed, they are defined a priori as population trees and not inferred. This thesis aims to contribute methods and tools for reconstructing the demographic history of populations that have been structured into an unknown number of sub-populations for long periods of time.

We present two approaches to demographic inference in the presence of structure. The first is based on the IICR (Inverse Instantaneous Coalescence Rate) which is a function of the demographic model and sampling scheme, and can be estimated for a single diploid individual. The proposed method fits observed IICR curves with exact IICR curves obtained under piecewise stationary symmetrical island models, and infers the number of islands, their common size, and the amount of gene flow in different periods of time. Our application to a set of five human PSMCs yielded demographic histories that are in agreement with previous studies suggesting ancient human structure.

The second approach assumes multiple genetic samples, and is centered on the use of the SFS (Site Frequency Spectrum) as a summary statistic for demographic inference. We focus on the efficient computation of the exact expected SFS under a general model of population structure, and show that, for sample sizes up to 26 haploids, it is possible to achieve good numerical accuracy and performance in symmetrical island models by exploiting the sparsity pattern of the transition matrix for the associated Markov process.

Résumé

L'inférence de l'histoire démographique des espèces représente l'un des plus grands défis de la génétique des populations. Cette histoire est souvent représentée comme une histoire de changements de taille de populations, sous l'hypothèse que la structure des populations peut être ignorée. Alternativement, lorsqu'on suppose qu'il y a une structure, elle est souvent définie a priori comme un arbre de populations et n'est pas inférée. Cette thèse vise à proposer des méthodes et des outils pour reconstruire l'histoire démographique de populations qui ont été structurées en un nombre inconnu de sous-populations pendant de longues périodes de temps.

Nous présentons deux approches de l'inférence démographique en présence de structure. La première est basée sur l'IICR (Inverse Instantaneous Coalescence Rate), fonction du modèle démographique et du schéma d'échantillonnage, qui peut être estimée pour un seul individu diploïde en utilisant la méthode du PSMC (pairwise sequentially Markovian coalescent). La méthode proposée ajuste les courbes IICR observées avec les courbes IICR exactes obtenues dans le cadre de modèles en îles symétriques stationnaires par morceaux, et infère le nombre d'îles, leur taille commune et l'importance du flux génétique à différentes périodes de temps. Notre application à un ensemble de cinq PSMCs humains a fourni des histoires démographiques qui sont en accord avec des études précédentes suggérant une structure humaine ancienne.

La seconde approche se base sur des échantillons génétiques multiples, et est centrée sur l'utilisation du SFS (Site Frequency Spectrum) comme statistique pour l'inférence démographique. Nous nous concentrons sur l'efficacité du calcul du SFS exact attendu dans le cadre d'un modèle général de structure de population, et nous montrons que, pour des tailles d'échantillon allant jusqu'à 26 haploïdes, il est possible d'obtenir une bonne précision numérique et de bonnes performances pour les modèles d'îles symétriques en exploitant la propriété de parcimonie de la matrice de transition pour le processus de Markov associé.

Acknowledgements

I would like to thank the Université Fédérale Toulouse Midi Pyrénées and the Région Occitanie for providing the funding for this project, and extending it in times of duress.

To the esteemed members of the jury, thank you for such a rich and insightful discussion session, and particularly to professors Hobolth and Achaz for their very nice reports and their truly inspired questions.

To my advisors Olivier and Lounès, thank you for all the support, and thank you for believing in me. You have been my role model of professional excellence, humility, and humanity throughout these years, and I will forever cherish the time we had together.

To Simon and Simona, thank you for the many sessions of insightful discussion. This work would not be what it is without your contributions.

To my friends and my family that have been with me every step of the way, I would not be here without you. Here's to more fun projects: cheers!

Contents

Introduction	11
Introduction (version française)	15
List of publications	19
1 Preliminaries	21
1.1 The Wright-Fisher model	21
1.2 The coalescent	22
1.2.1 The discrete-time coalescent	22
1.2.2 The continuous-time approximation	23
1.3 Demographic models	25
1.3.1 Population size change	25
1.3.2 Population structure	27
1.4 Demographic inference	28
1.4.1 A model of mutation	28
1.4.2 The SFS	29
1.4.3 The IICR	30
2 Demographic inference using the IICR	33
2.1 The piecewise stationary n-island coalescent	33
2.1.1 The parameter space	33
2.1.2 Computing and scaling the IICR	34
2.2 Optimization framework: search algorithm and optimality criteria	37
2.2.1 Comparison of optimization parameters	39
2.3 Validation framework and results	41
2.3.1 Sampling the parameter space	41
2.3.2 The three types of target IICRs	43
2.3.3 Quantifying the inference error	45
2.3.4 Validation results using exact target IICRs	46
2.3.5 Validation results using T-sim IICRs	50
2.4 Application to humans	54
2.4.1 Results	55
2.4.2 Seq-Sim validations	59

2.4.3	SFS comparisons	68
2.5	Implementation and use cases	69
3	Multi-sample summary statistics for demographic inference	73
3.1	The structured coalescent and the $IICR_k$	74
3.2	The structured coalescent with ancestry tracking	78
3.2.1	State matrices	79
3.2.2	Transitions between states	81
3.2.3	Grouping and sorting the states	82
3.2.4	Construction of the rate matrix	86
3.3	Computing the expected SFS	89
3.3.1	Numerical solution	90
3.3.2	Error analysis	93
3.3.3	Early termination	97
3.4	Model specialization: Symmetrical n-island	98
3.4.1	The rate matrix, revisited	99
3.4.2	The expected SFS in the n-island model	102
4	Conclusions and future work	113
4.1	The IICR and the inference of structure	113
4.2	A note on human evolution	114
4.3	Computing the exact SFS for structured populations	115
4.4	Structured demographic inference using multi-sample statistics	116
4.5	Closing thoughts	118
	Bibliography	123
A	Additional results of validation using exact simulated IICRs	125
A.1	Unscaled IICR	125
A.2	Scaled IICR	131
B	Additional results of validation using T-sim IICRs	137
B.1	Scenarios with 1 component	137
B.2	Scenarios with 2 components	140
B.3	Scenarios with 3 components	143
B.4	Scenarios with 4 components	146
B.5	Scenarios with 5 components	149

Introduction

Population genetics has seen a major transformation in the last few decades with the wide availability of full-scale genomic sequences. Global initiatives like the 1000 Human Genomes Project (Siva, 2008) has made it possible to have a highly detailed insight into the nature of genetic diversity. One of the great challenges facing the field today is analyzing this rapidly increasing mountain of information in a sensible and efficient manner (Beaumont, 2004; Johri et al., 2020). The development of mathematical models for the evolution of populations plays a central role in this effort, along with the associated statistical and numerical analysis methods. This thesis aims to contribute methods and tools for reconstructing the demographic history of populations that have long been geographically structured into an unknown number of sub-populations.

Reconstructing demographic histories in the context of population genetics is often achieved by fitting the parameters of a model according to the observed sequenced data. Some parameters of interest are past population sizes (Beaumont, 1999; Li and Durbin, 2011; Schiffels and Durbin, 2013; Liu and Fu, 2015; Boitard et al., 2016) or the dates of major demographic events such as population splits or admixture events and their proportions (Gutenkunst et al., 2009; de Barros Damgaard et al., 2018). One classical example of such parameter inference is Watterson's θ_W estimator (Watterson, 1975), which measures the population mutation rate by counting the number of places in the genome where a sample of individuals exhibit differences. A more recent example can be found in Li and Durbin's PSMC method (Li and Durbin, 2011), which estimates the ancestral size history of a population using the full sequence of a single diploid individual, under the operating assumption that such population has remained panmictic over time, i.e., isolated and without internal mating structures.

This assumption of panmixia has been widely prevalent in many studies over the years (Li and Durbin, 2011; Liu and Fu, 2015). Indeed, contending with the presence of structure and migration has the potential to greatly complicate any demographic model. Some studies of human evolution have introduced structure in the form of simplified tree models, where the number of populations is fixed a priori, i.e., the topology is not inferred (other aspects, such as divergence times, are inferred). The branches of these trees often represent large continental regions that are themselves assumed to have remained panmictic over long periods of time (Gutenkunst et al., 2009; Prado-Martinez et al., 2013; Noskova et al., 2019). These panmictic and tree models are useful approximations, and they have proven their utility in building stories of human expansions and population splits. However, if the inferred parameter values or conclusions may fundamentally change depending on the underlying model assumptions, then the resulting stories can be questioned (Wakeley, 1999; Mazet et al., 2016; Scerri et al.,

2019).

The present thesis focuses exclusively on models of structured populations. One of our core objectives in this context is to develop methods of demographic inference where the number of populations can be inferred rather than assumed. The other objective is to understand the extent to which the observed genetic diversity within populations can be attributed to the effects of structure and changes in connectivity rather than to changes in population size. The main theoretical background supporting this approach is the structured coalescent (Herbots, 1994) which provides a Markovian framework for the study of the genealogy of a sample of lineages that are dynamically migrating between a set of islands or demes. In Chapter 1 we present a summary of a few key concepts in population genetics such as the Wright-Fisher model and the associated coalescent process and a few of its extensions, including variable population size, deme-based population structure, and the infinite-sites mutation model. We also introduce the Inverse Instantaneous Coalescence Rate (IICR) function and the site frequency spectrum (SFS), both of which are proposed as summary statistics for inferential approaches.

Chapter 2 describes a method for inferring the number of populations and the changes in connectivity in a piecewise-stationary n -island model. The n -island is the simplest model of demographic structure, where all the islands are assumed of the same size, and connectivity is fully symmetrical in the sense that all islands exchange migrants with all other islands at the same rate. The piecewise-stationarity refers to the fact that this migration rate may change in a discrete manner throughout time. The dates of these events are also inferred. This model has some limitations as it ignores spatial distances and other complexities of real species, but the choice for this first attempt was guided by simplicity and computational feasibility. The method in question uses the IICR of a single diploid sample as input data in order to fit the model parameters. First described in Mazet et al. (2016), the IICR is closely linked with the probability distribution of the coalescence times of one pair of lineages. It is an attractive candidate for demographic inference in our context since it can be always defined independently of the assumed demographic model, and previous simulations (Chikhi et al., 2018; Rodríguez et al., 2018) have shown that it is sensitive to population structure, sampling patterns and fluctuations of migration rates. The IICR can be estimated from real data using methods such as the PSMC.

The process of fitting the data was accomplished using a general purpose meta-heuristic to explore the parameter space and minimize the distance between a candidate IICR and the target (observed) one. The chapter also details the approach and results of our validation efforts, which consisted in simulating IICR curves under different n -island models and comparing the inferred parameters with the known ones. We then applied the method to human genomic data using five published PSMC curves (Prado-Martinez et al., 2013) and compared the inferred histories between individuals and with previously inferred scenarios by Rodríguez et al. (2018) and Noskova et al. (2019). This study was published in Arredondo et al. (2021).

Chapter 3 focuses on establishing a framework for demographic inference of structured populations using multi-sample summary statistics. Having multiple samples provides better

information resolution in the not-so-distant past of the sample, since in typical ancestral trees, most of the total branch length is distributed close to the samples near the bottom of the tree (leaves). The chapter introduces two candidate multi-sample summary statistics: the IICR_k and the SFS. The IICR_k is a natural generalization of the IICR that is formulated using the random distribution of the waiting time until the first coalescence in a sample of size k , and can be estimated for real data using methods such as the MSMC (Schiffels and Durbin, 2013). On the other hand, the SFS is a histogram of allele frequencies that has been widely applied as an efficient way for summarizing the genetic diversity of a sample (Wakeley and Hey, 1997; Griffiths and Tavaré, 1998; Nielsen, 2000; Gutenkunst et al., 2009; Excoffier et al., 2013).

We show in this chapter that the Markovian framework of the structured coalescent can be generalized to capture the necessary information for computing these functions within arbitrary numerical accuracy. Similar approaches have been avoided in the past due to the computational challenges of handling the state spaces of the associated Markov processes, which grow in size almost exponentially as a function of the number of samples. However, we show that by not tracking the historical origin of the lineages (mainly used to compute the joint SFS of the populations), using appropriate model specializations that take advantage of inherent symmetries, and exploiting of the sparsity patterns that appear in the rate matrix under certain state orderings, it is possible to very quickly compute the expected SFS of samples of size up to $k = 26$ in models of symmetrical n -island structure. The approach is general in the sense that it allows for many potential types of structured scenarios, however we exclusively showcase the n -island as a model specialization for the implementation and results, once again due to its mathematical simplicity.

All along this work we have strived to not only establish novel methods for approaching the questions of demographic inference in the context of population structure, but also to provide implementations of these methods that are both intuitively usable and highly performant. The codebases for these projects are available in github.com/arredondos/snif and github.com/arredondos/sisifs.

Many questions remain unanswered. In the final chapter on conclusions and future work, we attempt to paint a broad map of where the results of this work can fit in a larger framework of demographic inference for structured populations by indicating some of the ways in which the presented methods can be generalized.

Introduction (version française)

La génétique des populations a connu une transformation majeure au cours des dernières décennies, grâce à une large disponibilité de séquences génomiques complètes. Des initiatives mondiales comme le projet 1000 génomes humains (Siva, 2008) ont permis d'avoir un aperçu très détaillé de la nature de la diversité génétique. L'un des grands défis auxquels le domaine est confronté aujourd'hui est l'analyse de cette masse d'information en rapide augmentation, d'une manière pertinente et efficace (Beaumont, 2004; Johri et al., 2020). Le développement de modèles mathématiques pour l'évolution des populations ainsi que les méthodes d'analyses statistique et numérique associées jouent un rôle central. Cette thèse a pour but d'apporter des méthodes et des outils de reconstruction de l'histoire démographique de populations qui sont depuis longtemps structurées géographiquement en un nombre inconnu de sous-populations.

La reconstruction de l'histoire démographique dans le contexte de la génétique des populations est souvent réalisée en ajustant les paramètres d'un modèle en fonction des données séquentielles observées. Certains paramètres d'intérêt sont les tailles de population passées (Beaumont, 1999; Li and Durbin, 2011; Schiffels and Durbin, 2013; Liu and Fu, 2015; Boitard et al., 2016) ou les dates d'événements démographiques majeurs tels que les divisions de population ou les événements de mélange et leurs proportions (Gutenkunst et al., 2009; de Barros Damgaard et al., 2018). Un exemple classique de cette inférence de paramètres est l'estimateur θ_W de Watterson (Watterson, 1975), qui mesure le taux de mutation de la population en comptant le nombre d'endroits du génome où un échantillon d'individus présente des différences. Un exemple plus récent se trouve dans la méthode PSMC de Li et Durbin (Li and Durbin, 2011), qui estime l'histoire de la taille ancestrale d'une population en utilisant la séquence complète d'un seul individu diploïde, sous l'hypothèse que cette population est restée panmictique au fil du temps, c'est-à-dire isolée et sans structures internes de reproduction.

Cette hypothèse de panmixie a été largement adoptée dans de nombreuses études au fil des ans (Li and Durbin, 2011; Liu and Fu, 2015). En effet, faire face à la présence de la structure et de la migration complique considérablement tout modèle démographique. Certaines études de l'évolution humaine ont introduit une structure sous la forme de modèles d'arbres simplifiés, où le nombre de populations est fixé a priori, c'est-à-dire que la topologie n'est pas inférée (d'autres aspects, comme les temps de divergence, le sont). Les branches de ces arbres représentent souvent de grandes régions continentales qui sont elles-mêmes supposées être restées panmictiques sur de longues périodes de temps (Gutenkunst et al., 2009; Prado-Martinez et al., 2013; Noskova et al., 2019). Ces modèles panmictiques et arborescents sont des approximations utiles, et ils ont prouvé leur utilité pour construire des histoires d'expansions humaines et de scissions de populations. Cependant, si les valeurs des paramètres déduites ou les conclusions peuvent changer fondamentalement en fonction des

hypothèses du modèle sous-jacent, alors les histoires qui en résultent peuvent être remises en question (Wakeley, 1999; Mazet et al., 2016; Scerri et al., 2019).

La présente thèse se concentre exclusivement sur les modèles de populations structurées. Un de nos objectifs principaux dans ce contexte est de développer des méthodes d'inférence démographique où le nombre de populations peut être déduit plutôt que supposé. L'autre objectif est de comprendre dans quelle mesure la diversité génétique observée au sein des populations peut être attribuée aux effets de la structure et des changements de connectivité plutôt qu'aux changements de taille de la population. Le principal fondement théorique de cette approche est le coalescent structuré (Herbots, 1994) qui fournit un cadre markovien pour l'étude de la généalogie d'un échantillon de lignées qui migrent dynamiquement entre un ensemble d'îles ou de dèmes. Dans le chapitre chap :prelim, nous présentons un résumé de quelques concepts clés en génétique des populations tels que le modèle de Wright-Fisher et le processus de coalescence associé, ainsi que quelques-unes de ses extensions, notamment la taille variable de la population, la structure de la population basée sur les dèmes et le modèle de mutation à sites infinis. Nous présentons également la fonction inverse du taux de coalescence instantané (IICR) et le spectre de fréquence des sites (SFS), qui sont tous deux proposés comme statistiques pour les approches inférentielles.

Le chapitre 2 décrit une méthode permettant de déduire le nombre de populations et les changements de connectivité dans un modèle n -île stationnaire par morceaux. Le modèle n -île est le modèle le plus simple de structure démographique, dans lequel toutes les îles sont supposées avoir la même taille, et la connectivité est totalement symétrique dans le sens où toutes les îles échangent des migrants avec toutes les autres îles au même rythme. La stationnarité par morceaux fait référence au fait que ce taux de migration peut changer de manière discrète au cours du temps. Les dates de ces événements sont également déduites. Ce modèle présente certaines limites car il ignore les distances spatiales et d'autres complexités des espèces réelles, mais le choix pour cette première tentative a été guidé par la simplicité et la faisabilité informatique. La méthode en question utilise l'IICR d'un seul échantillon diploïde comme données d'entrée afin d'ajuster les paramètres du modèle. Décrit pour la première fois dans Mazet et al. (2016), l'IICR est étroitement lié à la distribution de probabilité des temps de coalescence d'une paire de lignées. Il s'agit d'un candidat intéressant pour l'inférence démographique dans notre contexte puisqu'il peut toujours être défini indépendamment du modèle démographique supposé, et des simulations antérieures (Chikhi et al., 2018; Rodríguez et al., 2018) ont montré qu'il est sensible à la structure de la population, aux modèles d'échantillonnage et aux fluctuations des taux de migration. L'IICR peut être estimé à partir de données réelles en utilisant des méthodes telles que le PSMC.

Le processus d'ajustement des données a été réalisé à l'aide d'une méta-heuristique d'usage général pour explorer l'espace des paramètres et minimiser la distance entre un IICR candidat et l'IICR cible (observé). Ce chapitre détaille également l'approche et les résultats de nos efforts de validation, qui ont consisté à simuler des courbes d'IICR sous différents modèles n -îles et à comparer les paramètres déduits avec les paramètres connus. Nous avons ensuite appliqué la méthode à des données génomiques humaines en utilisant cinq courbes PSMC publiées (Prado-Martinez et al., 2013) et comparé les histoires inférées entre les individus et

avec les scénarios inférés précédemment par Rodríguez et al. (2018) et Noskova et al. (2019). Cette étude a été publiée dans Arredondo et al. (2021).

Le chapitre 3 établit un cadre pour l'inférence démographique de populations structurées à l'aide de statistiques multi-échantillons. Le fait d'avoir plusieurs échantillons permet une meilleure résolution des informations dans un passé pas si lointain de l'échantillon, puisque dans les arbres ancestraux typiques, la majeure partie de la longueur totale des branches est localisée à proximité des échantillons situés au bas de l'arbre (feuilles). Ce chapitre présente deux statistiques multi-échantillons candidates : l'IICR_k et le SFS. L'IICR_k est une généralisation naturelle de l'IICR formulé en utilisant la distribution aléatoire du temps d'attente jusqu'à la première coalescence dans un échantillon de taille k , et peut être estimé pour des données réelles en utilisant des méthodes telles que le MSMC (Schiffels and Durbin, 2013). D'autre part, le SFS est un histogramme des fréquences alléliques qui a été largement appliqué car c'est un moyen efficace de résumer la diversité génétique d'un échantillon (Wakeley and Hey, 1997; Griffiths and Tavaré, 1998; Nielsen, 2000; Gutenkunst et al., 2009; Excoffier et al., 2013).

Nous montrons dans ce chapitre que le cadre markovien du coalescent structuré peut être généralisé pour capturer les informations nécessaires au calcul de ces fonctions avec une précision numérique arbitraire. Des approches similaires ont été évitées dans le passé en raison des difficultés de calcul liées à la manipulation des espaces d'état des processus de Markov associés, dont la taille croît de manière presque exponentielle en fonction du nombre d'échantillons. Cependant, nous montrons qu'en ne suivant pas l'origine historique des lignées (principalement utilisée pour calculer le SFS conjoint des populations), en utilisant des spécialisations de modèle appropriées qui tirent profit des symétries inhérentes, et en exploitant les propriétés de parcimonie qui apparaissent dans la matrice de taux sous certains ordres d'état, il est possible de calculer très rapidement le SFS attendu d'échantillons de taille allant jusqu'à $k = 26$ dans des modèles de structure symétrique à n -îles. L'approche est générale dans le sens où elle permet de nombreux types de scénarios structurés. Cependant nous présentons exclusivement le modèle en îles pour la mise en œuvre et les résultats, une fois encore en raison de sa simplicité mathématique.

Tout au long de ce travail, nous nous sommes efforcés non seulement d'établir de nouvelles méthodes pour aborder les questions d'inférence démographique dans le contexte de la structure de la population, mais aussi de fournir des implémentations de ces méthodes qui soient à la fois intuitivement utilisables et hautement performantes. Les programmes sont disponibles dans github.com/arredondos/snif et github.com/arredondos/sisifs.

De nombreuses questions restent sans réponse. Dans le dernier chapitre, consacré aux conclusions et aux travaux futurs, nous tentons de dresser une carte générale de la place des résultats de ce travail dans un cadre plus large d'inférence démographique pour les populations structurées, en indiquant certaines des façons dont les méthodes présentées peuvent être généralisées.





List of publications

The present thesis is based on the following publications and manuscripts:

- Armando Arredondo, Beatriz Mourato, Khoa Nguyen, Simon Boitard, Willy Rodríguez, Camille Noûs, Olivier Mazet and Lounès Chikhi. *Inferring number of populations and changes in connectivity under the n-island model*. 2021. *Heredity*, 126(6):896–912.
- Armando Arredondo, Simon Boitard, Simona Grusea, Josué Corujo, Camille Noûs, Lounès Chikhi and Olivier Mazet. *Computing the exact expected Site Frequency Spectrum for structured populations*. 2022. In preparation.

During my doctoral years, I also had the opportunity to pursue other collaborations. The following publications and manuscripts are the results of these exchanges. These results are not included in the present thesis.

- Willy Rodríguez, Olivier Mazet, Simona Grusea, Armando Arredondo, Josué M. Corujo, Simon Boitard, Lounès Chikhi. *The IICR and the non-stationary structured coalescent: towards demographic inference with arbitrary changes in population structure*. 2018. *Heredity*, 121(6):663–678.
- Simon Boitard, Armando Arredondo, Camille Noûs, Lounès Chikhi, Olivier Mazet. *Heterogeneity in effective size across the genome: effects on the Inverse Instantaneous Coalescence Rate (IICR) and implications for demographic inference under linked selection*. 2021. Submitted. *BioRxiv preprint*: doi.org/10.1101/2021.06.11.448122.
- Helena Teixeira, Jordi Salmons, Armando Arredondo, Beatriz Mourato, Sophie Manzi, Romule Rakotondravony, Olivier Mazet, Lounès Chikhi, Julia Metzger, Ute Radespiel. *Impact of model assumptions on demographic inferences — the case study of two sympatric mouse lemurs in northwestern Madagascar*. 2021. Submitted.
- Beatriz Mourato, Armando Arredondo, Pierre Lacoste, Willy Rodríguez, Khoa Nguyen, Simon Boitard, Olivier Mazet, Lounès Chikhi. *Genomic data and demographic inferences under structured models applied to lemurs*. 2022. In preparation.



Chapter 1.

Preliminaries

In this chapter we briefly summarize some basic elements of population genetics in order to make the document more self-contained, and also to establish a few of the naming conventions and assumptions that will be referenced in later chapters. For the sake of brevity, the material here is discussed quite tersely, and only the theoretical concepts that are most closely related to the themes of the following chapters are covered. For a more complete and thorough exposition of population genetics theory, see Hein et al. (2004); Tavaré (2004); Gillespie (2004).

1.1. The Wright-Fisher model

The Wright-Fisher model in its most simple version describes the evolution of a population of $2N$ haploid individuals (often interpreted as genes) with minimal assumptions. The generations are discrete and non-overlapping, and generation $t + 1$ is created by making each individual a copy of one of the $2N$ from generation t , chosen randomly with probability $1/(2N)$.

The model in this form makes a number of strong assumptions about the population, which evolves in a sequence of discrete time steps by replacing the entire generation with its descendants in every step. Each descendant is an exact copy of its predecessor, so the model does not have a concept of sexual reproduction or gene recombination. Additionally, since each predecessor is selected randomly with uniform probability from the population, there is no concept of demographic structure, fitness or selection in the model.

Typically, the populations we are modeling consist of the genes or alleles of a genome, instead of individual organisms. In the following sections, we focus on a backwards-in-time view of the Wright-Fisher model, known as the coalescent, and relax some of the assumptions made here in favor of having more useful models.




Figure 1.1: An instantiation of a Wright-Fisher process with population size $2N = 10$. Panel (a) shows the natural evolution of the process over the course of 16 generations from top (past) to bottom (present). In panel (b) the demes of each generation have been rearranged such that the tree-like nature of the genealogy is more apparent. Panel (c) highlights the genealogy of a sample of three genes. We see that the first coalescence happens two generations in the past, and that the MRCA of the sample is reached after 9 generations. Panels reproduced from Hein et al. (2004) §1.4.3.

1.2. The coalescent

Given the previously discussed Wright-Fisher model (Figure 1.1), we may be interested in asking: how much time (backwards in generations) until the most recent common ancestor (MRCA) for two genes sampled from the population? What if we sampled k genes? What is the probability distribution of these times? Coalescent theory studies these questions, and provides a framework for modeling the ancestral genealogy of a sample within an evolving population.

1.2.1. The discrete-time coalescent

Given a sample of two genes in a Wright-Fisher population of size $2N$, the probability that they share a common ancestor in the previous generation can be computed by taking the ratio of favorable to possible outcomes as:

$$\frac{2N}{(2N)^2} = \frac{1}{2N},$$

therefore, the probability that they *do not* share a common ancestor in the previous generation is:

$$1 - \frac{1}{2N}.$$

With this it follows that the probability for a sampled pair of lineages to find a common ancestor in exactly g generations in the past is:

$$\left(1 - \frac{1}{2N}\right)^{g-1} \frac{1}{2N}, \quad (1.2.1)$$

since they choose different ancestors in the first $g - 1$ generations, and then choose the same ancestor in the next one. In this case we say that the pair of genes coalesces in g generations. We denote the time until coalescence of two lineages by T_2^N , with a probability distribution $P(T_2^N = g)$ as in (1.2.1), which corresponds to a geometric distribution with parameter $1/(2N)$ and an expected value of $2N$.

Generalizing for a sample of size k lineages, we have that the probability of there being no coalescences in the previous generation, again counting favorable and total outcomes, is:

$$\begin{aligned} P(T_k^N > 1) &= \frac{(2N)(2N-1)\cdots(2N-k+1)}{(2N)^k} \\ &= \left(1 - \frac{1}{2N}\right)\left(1 - \frac{2}{2N}\right)\cdots\left(1 - \frac{k-1}{2N}\right) \\ &= 1^{k-1} + \left(\frac{-1}{2N}\right)1^{k-2} + \left(\frac{-2}{2N}\right)1^{k-2} + \cdots + \left(\frac{-k+1}{2N}\right)1^{k-2} + O\left(\frac{1}{N^2}\right) \\ &= 1 - \frac{1}{2N} \sum_{i=1}^{k-1} i + O\left(\frac{1}{N^2}\right) = 1 - \frac{k(k-1)}{4N} + O\left(\frac{1}{N^2}\right), \end{aligned} \quad (1.2.2)$$

where $O(1/N^2)$ captures all the terms that are divided by N^2 or a higher power. If we introduce the additional assumption that $k \ll N$, these terms are negligible and can be discarded. In terms of the model, this is equivalent to assuming that when the size of a random sample is very small compared to the size of the population, it is extremely unlikely that more than two lineages will have a common ancestor in the previous generation, and thus this possibility can be excluded.

With this approximation we have that the probability of observing the first coalescence in a sample of size k in exactly g generations (becoming then a sample of size $k - 1$) is given by:

$$P(T_k^N = g) \approx \left(1 - \frac{k(k-1)}{4N}\right)^{g-1} \frac{k(k-1)}{4N}. \quad (1.2.3)$$

Consequently, the variable T_k^N can be approximated by a geometric distribution with parameter $k(k-1)/(4N) = \binom{k}{2}/(2N)$. Since the ancestors of the genes are chosen with uniform probability, the times $T_k^N, T_{k-1}^N, \dots, T_2^N$ are independent.

1.2.2. The continuous-time approximation

In the previous sections we measure time in discrete units of time (generations), and the corresponding T_k^N random variables reflect this. However, for large populations sizes and

long-running processes (large N and g), the coalescent can be well approximated by a continuous-time process where time is scaled in proportion to the population size. Using this transformation, the coalescence times become independent of the population size. For instance, if we take $t = g/(2N)$, then the probability that a pair of lineages will not coalesce before g generations is:

$$\begin{aligned} P(T_2^N > g) &= \left(1 - \frac{1}{2N}\right)^g \\ \Leftrightarrow P\left(\frac{T_2^N}{2N} > t\right) &= \left(1 - \frac{1}{2N}\right)^{2Nt} \approx e^{-t} \\ \Leftrightarrow P\left(\frac{T_2^N}{2N} \leq t\right) &\approx 1 - e^{-t}, \end{aligned}$$

therefore the variable $T_2 := T_2^N/(2N)$ can be approximated by an exponential distribution of parameter 1 for large values of N . Using a similar transformation we can see that the scaled times $T_k := T_k^N/(2N)$ can be approximated by exponential distributions of parameter $\binom{k}{2}$:

$$\begin{aligned} P(T_k^N > g) &= \left(1 - \binom{k}{2} \frac{1}{2N}\right)^g \\ \Leftrightarrow P(T_k \leq t) &\approx 1 - e^{-\binom{k}{2}t}. \end{aligned}$$

In order to reduce the notation complexity, the approximation $T_k \approx \exp\left(\binom{k}{2}\right)$ will be implicitly assumed in the following equations. With this we can compute the expected value of several quantities related to the genealogy of the samples. For instance, the expected time until the first coalescence among k samples is:

$$\mathbb{E}(T_k) = \frac{1}{\binom{k}{2}} = \frac{2}{k(k-1)}.$$

Other quantities of interest are the expected time until the most recent common ancestor of the samples, given by the total height of the tree: $T_{\text{MRCA}} = T_k + T_{k-1} + \dots + T_2$:

$$\begin{aligned} \mathbb{E}(T_{\text{MRCA}}) &= \sum_{i=2}^k \mathbb{E}(T_i) \\ &= 2 \sum_{i=2}^k \frac{1}{i(i-1)} = 2 \left(1 - \frac{1}{k}\right); \end{aligned} \tag{1.2.4}$$

and the expected total branch length of the tree, which can be computed by weighing the times T_i according to the number of live lineages:

$$\begin{aligned} \mathbb{E}(BL_k) &= \sum_{i=2}^k i \mathbb{E}(T_i) \\ &= 2 \sum_{i=2}^k \frac{1}{i-1} = 2 \sum_{i=1}^{k-1} \frac{1}{i}. \end{aligned} \tag{1.2.5}$$

We can see from equations (1.2.4) and (1.2.5) that taking large samples has quickly diminishing returns, since the expected tree height is bounded by 2, and the expected total branch length grows on the order of $\log(k)$.

An important benefit of looking at the population evolution through the lens of the coalescent is that simulating gene genealogies becomes computationally feasible, even for very long-running processes and very large population sizes. Indeed, given a sample of k genes in a Wright-Fisher population, we can simulate the exponential time T_k , and then choose a random pair from the $\binom{k}{2}$ possible pairs to form a coalescence. Then replace k with $k - 1$ and repeat until the MRCA of the sample is reached. This approach, with many other added capabilities for more general models, is notably used by the ubiquitous simulation software `ms` (Hudson, 2004).

1.3. Demographic models

In this section we discuss some generalizations to the Wright-Fisher model and the associated coalescent process that allow for more flexibility in the population evolution, and thus more realistic and useful models featuring fluctuating population sizes and population structure.

1.3.1. Population size change

In the previous sections we have been assuming that the evolution of the population is constrained to having a constant size throughout time. We now relax this assumption by introducing a population size function. Returning to the discrete-time framework, we denote by $N(g)$ the number of haploid individuals or sequences in the population g generations before the present. Therefore, $N(0) = 2N$ would be the current population size, $N(1)$ the size of the generation prior, and so on. We note that in this model, the probability for two genes sampled in the present to have a common ancestor in the previous generation is

$$\frac{1}{N(1)}.$$

Continuing this argument g generations in the past we have that the probability that a pair of genes have not yet coalesced is given by:

$$P(T_2^N > g) = \prod_{i=1}^g \left(1 - \frac{1}{N(i)}\right). \quad (1.3.1)$$

We can consider a continuous-time version of $N(g)$ by rounding up the coalescent times $2Nt$ to an integer number of generations. With this we can define $N^c(t) := N(\lceil 2Nt \rceil)$ as well as the continuous-time relative size function $\lambda(t) = N^c(t)/N^c(0)$. We now derive the distribution of the time T_2 for this size changing population. Recall that for $x \in (0; 1)$ we have the inequality

$$x \leq -\ln(1 - x) \leq \frac{x}{1 - x},$$

which after substituting x for $1/N(i)$ and adding for $i = 1, \dots, g$ reduces to

$$\sum_{i=1}^g \frac{1}{N(i)} \leq - \sum_{i=1}^g \ln \left(1 - \frac{1}{N(i)} \right) \leq \sum_{i=1}^g \frac{1}{N(i) - 1}. \quad (1.3.2)$$

We note that with the change of variable $\tau = i/2N$ we can have the transformation

$$\sum_{i=1}^g \frac{1}{N(i)} = \sum_{i=1}^g \frac{1}{N^c(i/2N)} = \int_1^g \frac{di}{N^c(i/2N)} = \int_0^{g/2N} \frac{d\tau}{\lambda(\tau)}.$$

Additionally, we know from (1.3.1) that

$$\sum_{i=1}^g \ln \left(1 - \frac{1}{N(i)} \right) = \ln \prod_{i=1}^g \left(1 - \frac{1}{N(i)} \right) = \ln P(T_2^N > g) = \ln P(T_2 > g/2N)$$

therefore (1.3.2) implies that for large N :

$$\begin{aligned} - \ln P(T_2 > g/2N) &\approx \int_0^{g/2N} \frac{d\tau}{\lambda(\tau)} \\ \Leftrightarrow P(T_2 > t) &\approx \exp \left(- \int_0^t \frac{d\tau}{\lambda(\tau)} \right). \end{aligned} \quad (1.3.3)$$

Equation (1.3.3) indicates that there is a very close relation between the relative size change function $\lambda(t)$ of a varying size Wright-Fisher population and the distribution of the time until the coalescence of a pair of genes sampled from such a population. We return to this idea in §1.4.3.

A similar result can be had for the distribution of T_k , i.e., the time until the first coalescence in a sample of size k . Indeed, by counting favorable and total cases, we have:

$$\begin{aligned} P(T_k^N \geq 1) &= \frac{N(1)(N(1) - 1)(N(1) - 2) \cdots (N(1) - (k - 1))}{N(1)^k} \\ &= \prod_{i=1}^{k-1} \left(1 - \frac{i}{N(1)} \right), \end{aligned}$$

which, after following a series of derivations similar to those of (1.2.2), conduces to:

$$\begin{aligned} P(T_k^N \geq 1) &\approx 1 - \binom{k}{2} \frac{1}{N(1)} \\ \Rightarrow P(T_k^N \geq g) &\approx \prod_{i=1}^g \left(1 - \binom{k}{2} \frac{1}{N(i)} \right) \\ &\approx 1 - \binom{k}{2} \sum_{i=1}^g \frac{1}{N(i)} \\ &= 1 - \binom{k}{2} \int_0^{g/2N} \frac{d\tau}{\lambda(\tau)}, \end{aligned}$$

from which we get:

$$P(T_k > t) \approx \exp\left(-\binom{k}{2} \int_0^t \frac{d\tau}{\lambda(\tau)}\right). \quad (1.3.4)$$

1.3.2. Population structure

Another one of the strong assumptions in the basic Wright-Fisher model relates to the fact that parents are chosen at random with equal probability. In a mating population, this would translate to the assumption that every mating pair is equally as likely to produce offspring. In reality, physical distances often impose an important barrier for finding mates, and geographically close pairs are more likely to produce offspring than those that are much further apart.




Figure 1.2: Examples of models of structure.

We can model this idea in the Wright-Fisher scheme using the finite Herbots's model of the structured coalescent (Herbots, 1994). We assume the existence of n populations or *demes* that behave as haploid Wright-Fisher models of size $N_i = 2s_i N$ genes each, where s_i is the relative deme size and N is large as usual. Migration occurs between demes as in each generation a proportion q_{ij} of lineages migrates from deme i to deme j . We denote by m_i the proportion of the population of deme i that was received from other demes in any given generation, such that $m_i = \sum_{j \neq i} q_{ji} s_j / s_i$. Herbots showed that measuring time in units of $2N$ generations and making N go to infinity in such a way that the number of migrants stays bounded, the model converges to a continuous-time Markov process. In this transformation, q_{ij} goes to zero in such a way that the product $2Nq_{ji} s_j / s_i$ converges and has limit $M_{ij}/2$. Thus we can express the transition rates in the rate matrix Q of the Markov process as functions of n , s_i and M_{ij} .

There are several different types of demographic models. They differ by the configuration of their deme sizes and the patterns of connectivity. Some notable examples (see Figure 1.2):

- The symmetrical islands models, or n -island model for short (Wright, 1931): all islands have the same size, and all the migration rates M_{ij} between any pair of islands i and j are equal, so we use the notation $M = (n-1)M_{ij}$ to denote the migration rate received by any given island.

- Continent-island models: one island (the continent) is significantly larger in size than the rest, which are of the same size. Additionally, migration only occurs from the continent to the islands with rate M_1 . In some cases, migration back from the islands into the continent is modeled with another rate M_2 .
- Stepping-stone models: the islands are spatially arranged, and migration can only occur between neighboring islands, with symmetrical rate M . Many different spatial arrangements are possible, but popular ones are rectangular and linear arrangements.

It is important to note that in general structured scenarios, no closed-form expressions can be obtained for quantities of interest like $\mathbb{E}(T_2)$ or the site frequency spectrum (SFS), for instance. One recurring theme of this dissertation is exploring such properties in special cases of structured scenarios.

1.4. Demographic inference

Demographic inference is one of the central topics of this work. We are interested in extracting information from observed data in the form of estimated model parameters. There are several different coalescent models attempting to describe various aspects of the genealogy and using different sets of assumptions, and there is often a multitude of parameters defining such models. Regardless of these variations, mutations are the source of observed genetic diversity, and any model that attempts to explain real-world data must account for the presence of mutations.

1.4.1. A model of mutation

Here we explore a generalization of the Wright-Fisher model where the assumption that each descendant is exactly identical to its parent is lifted. We introduce genetic types by imposing a mutation process on top of the evolution process, where each time a lineage or gene copies its parent from the previous generation, a mutation is introduced with probability u , and thus the offspring becomes a different genetic type. Alternatively, the new lineage is an exact copy of its ancestor with probability $1 - u$.

Under the lens of the coalescent we may be interested in knowing, for instance, the average number of generations we need to trace back in the genealogical history of a single lineage in order to find a mutation. In each step backwards in time, the probability that the lineage experienced a mutation has a Bernoulli distribution of parameter u , therefore the probability that the lineage has mutated after exactly g backward generations is given by

$$P(T_M^N = g) = u(1 - u)^{g-1},$$

where T_M^N is the time until first mutation of a sampled lineage, with a geometric distribution of parameter u . Measuring time in units of $2N$ generations, we can once again obtain a

convenient continuous-time approximation for large N :

$$\begin{aligned} P(T_M^N \leq g) &= 1 - (1 - u)^g, \\ \Leftrightarrow P(T_M \leq t) &\approx 1 - e^{-t\theta/2}, \end{aligned} \quad (1.4.1)$$

where T_M denotes time in units of $2N$ generations and $\theta := 4Nu$ is the scaled mutation rate. This parameter can be interpreted as the expected number of mutations separating a sample of two lineages, since the expected time until their coalescence is $2N$ generations, and thus $2Nu$ mutations are expected on each branch.

Equation 1.4.1 indicates that T_M can be approximated by an exponential variable of parameter $\theta/2$ for large values of N . A consequence of this is that the number of mutations in a branch of length t in the genealogical tree is Poisson distributed, with intensity $t\theta/2$. Therefore, if we denote this number of mutations M_t , we have:

$$P(M_t = k) = \frac{(t\theta)^k}{k!2^k} e^{-t\theta/2}. \quad (1.4.2)$$

Knowing the distribution of M_t allows us to simulate the mutations efficiently. Indeed, given a simulated genealogical tree, we can randomly place M_t mutations on each branch of length t to obtain a realization of the coalescent process with mutation.

1.4.2. The SFS

We can model the process in which mutations occur in genomic sequences in a simple way by introducing the assumption that the probability for a mutation to occur in the same position twice is null. This assumption is justified in the presence of very long sequences with low mutation rates, and it is known as the infinite sites model of mutation, in which mutations always happen in new locations. Under this assumption, we can assign binary states to the positions along the genome. The two states, which can be coded as 1 and 0, can be interpreted as the position containing a mutation or not, respectively. We can thus represent a genomic sequence as a string of zeroes and ones. If two genes are observed to have different codes in a given position, we say that the position is a segregating or polymorphic site. We could distinguish a mutated position from a non-mutated one, for instance, by comparing with a sampled sequence from an outgroup species.

A convenient way of summarizing the genetic diversity of a sample under these assumptions is to simply inspect the values of the sampled sequences in the segregating sites, given that the sequences match exactly in all other sites. Consider the following example, where $k = 6$ sequences were sampled, and 8 segregating sites were observed among them.

The rows of Table 1 contain the sequence encoding for each of the segregating positions (columns, labeled SNP j for *single-nucleotide polymorphism* number j). For instance, we can see that 4 of the 8 polymorphisms were *singletons*, meaning that a mutation happened in exactly one of the sampled sequences. We can similarly track the number of sites that exhibited exactly i mutations, and compile this information in a vector $\xi = (\xi_1, \dots, \xi_{k-1})$, where ξ_i is the number of segregating sites at which there are exactly i copies of the mutant

	SNP 1	SNP 2	SNP 3	SNP 4	SNP 5	SNP 6	SNP 7	SNP 8
Sample 1	0	1	0	0	0	0	1	0
Sample 2	1	0	1	0	0	0	1	0
Sample 3	0	1	1	0	0	1	0	0
Sample 4	0	0	0	0	1	0	1	1
Sample 5	0	0	1	0	0	0	1	0
Sample 6	0	0	0	1	0	1	1	0
Total	1	2	3	1	1	2	5	1

Table 1.1.: An example for how to compute the observed site frequency spectrum (SFS) in a sample of size $k = 6$ genes. The columns show the segregating sites (SNPs), which are the positions in the genome where a mutation is observed for at least one of the sampled sequences. With increased sequence length, the table grows horizontally, but the SFS is always of length $k - 1$ because this is the maximum possible number of observed mutations (a column full of 1s is not possible since this would correspond to a mutation taking place before the most recent common ancestor of the sample, and as such it would not be visible in the coalescence tree).

type. In the example of Table 1, $\xi = (4, 2, 1, 0, 1)$. Under the assumptions of a Wright-Fisher model with constant population size and scaled mutation rate θ , it can be shown that:

$$\mathbb{E}(\xi_i) = \frac{\theta}{i}, \quad i = 1, \dots, k - 1. \quad (1.4.3)$$

See Tajima (1989); Fu (1995); Griffiths and Tavaré (1998); Hudson (2015) for various proofs of (1.4.3). The vector ξ is known as the Site Frequency Spectrum, and it has been the basis for many estimators and test statistics for analyzing genomic data. Notable examples are the θ_W (Watterson, 1975) and θ_π (Tajima, 1983) estimators of θ ; Tajima's D (Tajima, 1989); Fay and Wu's H (Fay and Wu, 2000); and the fixation index F_{ST} (Wright, 1943).

1.4.3. The IICR

We show in §1.3.1 that the distribution of T_2 times in a size-changing Wright-Fisher population is closely related to the relative size function $\lambda(t)$. Indeed, in (1.3.3) we can solve for $\lambda(t)$ to obtain:

$$\lambda(t) = \frac{P(T_2 > t)}{f_{T_2}}, \quad (1.4.4)$$

where f_{T_2} is the probability density function of the T_2 random variable. There are methods that can estimate the right-hand side of (1.4.4) from real data, most notably the PSMC method from Li and Durbin (2011) estimates T_2 by analyzing the local density of the number of segregating sites along the genome of a single diploid individual. This method has been used to estimate the function $\lambda(t)$ in several applications, but it is notable that the quantity $P(T_2 > t)/f_{T_2}$ can always be defined, regardless of the underlying evolutionary model or its assumptions, therefore it is useful to study it independently of the assumptions of a size-changing Wright-Fisher population. This was first done by Mazet et al. (2016), which referred to it as the Inverse Instantaneous Coalescence Rate, or IICR for short.

It is important to stress that the IICR is not tied to any particular model, panmictic or otherwise. It is defined using the distribution of the coalescence times of a sample of size two as in (1.4.4). It can be approximated to arbitrary numerical precision under many demographic models (Rodríguez et al., 2018); it can also be computed empirically by simulating a sample of coalescent times (Chikhi et al., 2018); and it can be read from full sequence genomic data using the appropriate methods, like the previously cited PSMC.

YOUR GENETIC TEST RESULTS
ARE BACK. APPARENTLY YOU'RE
PART OF AN UNBROKEN LINEAGE
STRETCHING BACK BILLIONS OF
YEARS TO THE EARLY EARTH!



Chapter 2.

Demographic inference using the IICR

To use the IICR as a summary of genomic information we first assume that an IICR curve can be obtained, which we will use as the *target* for demographic inference. With simulated data (sequences or T_2 values) this target curve can be obtained under any pre-defined coalescent model that could be expressed with a simulation tool (e.g., the ms program (Hudson, 2004)). With real genome-wide sequence data, the curve can be estimated with the PSMC method of Li and Durbin (2011). We then try to identify a piecewise stationary n -island model that generates an IICR that is identical or similar to the target IICR (or PSMC curve). The similarity between the two IICR curves is quantified with a distance metric defined below. We use a genetic algorithm to explore the parameter space (number of populations, migration rate within a time component, and timing of these changes assuming a fixed number of components for each independent analysis) and minimize that distance. We compute the IICR under the non-stationary structured coalescent (NSSC) of Rodríguez et al. (2018).

2.1. The piecewise stationary n -island coalescent

2.1.1. The parameter space

We first define the parameter space, as this directly determines the family of demographic histories that we can explore and infer from. The piecewise stationarity refers to the fact that, although migration rate is constant and identical between any pair of islands, this rate may be different between consecutive time periods (components), and there is a fixed number γ that represents the number of demographic events. To say that there are $\gamma \geq 0$ changes of gene flow thus means that there are $c = \gamma + 1$ components or periods of constant gene flow. Likewise, the deme size, which is the same for all islands, may in theory change through time in the general model presented in Rodríguez et al. (2018). In the present study we focus on models with constant population size but we present a more general model where deme sizes can change between components. In this more general case, the parameter space includes the number of islands n , the times t_i for the demographic events, and the values of both the migration rate M_i and the local deme size s_i at each new demographic period. Note that n is inferred but it does not change through time. We thus assume no extinction, no population split and no creation of new populations.

Given a fixed integer γ of demographic events to consider ($\gamma \geq 0$) and a collection of bounds B in the form of:

$$B = ([n_{\min}; n_{\max}], [t_{1\min}; t_{1\max}] \dots [t_{\gamma\min}; t_{\gamma\max}], \\ [M_{0\min}; M_{0\max}] \dots [M_{\gamma\min}; M_{\gamma\max}], [s_{0\min}; s_{0\max}] \dots [s_{\gamma\min}; s_{\gamma\max}]), \quad (2.1.1)$$

we define the parameter space $\Phi_{\gamma,B}$ as:

$$\Phi_{\gamma,B} = \left\{ \varphi = (n, t_1 \dots t_\gamma, M_0 \dots M_\gamma, s_0 \dots s_\gamma) \in \mathbb{N} \times \mathbb{R}^{3\gamma+3}, \text{ s.t. } \forall i: \right. \\ \left. \begin{aligned} 2 \leq n_{\min} \leq n \leq n_{\max}; \quad 0 < t_{i\min} \leq t_i \leq t_{i\max}; \\ 0 < M_{i\min} \leq M_i \leq M_{i\max}; \quad 0 < s_{i\min} \leq s_i \leq s_{i\max} \end{aligned} \right\}. \quad (2.1.2)$$

We define bounds for each variable because we use a constrained optimization algorithm, for which all parameters must be bounded (see §2.2). Also, since we focus on the case where there are no deme size changes, we enforce this by using B , as making $s_{i\min} = s_{i\max} = 1$ for all $0 \leq i \leq \gamma$ effectively fixes all deme sizes to 1, and reduces the number of parameters to $2\gamma + 2$.

2.1.2. Computing and scaling the IICR

Given any demographic scenario from $\Phi_{\gamma,B}$, the associated coalescent process is an instance of the NSSC of Rodríguez et al. (2018). Our main object of interest regarding these scenarios is the IICR. In the supplementary materials we provide a brief overview of how to perform its computation for any given $\varphi \in \Phi_{\gamma,B}$ based on the work of Rodríguez et al. (2018).

The computation of the IICR uses functions that receive the time t in units of $2N$ generations, and return values in units of N generations per coalescence, so these IICR functions are dimensionless in the sense that they operate in a *relative* frame of reference.

In order to compare the IICR with PSMC inferences, we need to re-scale both the time and the IICR values by a reference deme of size N which specifies how many haploid genes correspond to a local deme size of 1 as follows:

$$\text{sIICR}(g) = N \text{IICR}(g/2N);$$

where $\text{sIICR}(g)$ refers to the *scaled* IICR, and $\text{IICR}(t)$ to the *unscaled* (dimensionless) one. Note that we use g for generations as the variable name for sIICR to further stress the difference. The parameter space for the sIICR can be thought of as a simple one-dimensional addition to $\Phi_{\gamma,B}$:

$$\hat{\Phi}_{\gamma,B} = \{ \hat{\varphi} = (N, \varphi) \in \mathbb{R} \times \Phi_{\gamma,B} \text{ such that } 0 < N_{\min} \leq N \leq N_{\max} \}.$$

Considering the c components which constitute the model, we note that during component i , taking $0 \leq i \leq \gamma = c - 1$, the underlying coalescent process X_t is being governed by an

n -island model which has transition rate (see Rodríguez et al. (2018) for details and a more general approach):

$$Q_i = \begin{pmatrix} -M_i - \frac{1}{s_i} & M_i & \frac{1}{s_i} \\ \frac{M_i}{n-1} & -\frac{M_i}{n-1} & 0 \\ 0 & 0 & 0 \end{pmatrix},$$

where the three states in the matrix Q_i are sorted, both by row and column, as:

1. the two sampled lineages are in the same deme (configuration 'same');
2. the two sampled lineages are in different demes (configuration 'diff.');
3. the sampled lineages have coalesced (absorption).

Additionally, the rates in this matrix are active during the time interval $[t_i; t_{i+1})$, where we understand t_0 and $t_{\gamma+1}$ to be 0 and $+\infty$ respectively.

We know from that the probability distribution of T_2 is given by the cumulative effects of the exponential functions e^{tQ_i} (see Hobolth et al. (2019) for an extensive review). More formally, given any $t > 0$, let i be the largest index such that $t_i \leq t$, for which case we have:

$$P(t) = \left(\prod_{k=1}^i e^{(t_k - t_{k-1})Q_{k-1}} \right) e^{(t-t_i)Q_i},$$

$$D(t) = \frac{dP(t)}{dt} = \left(\prod_{k=1}^i e^{(t_k - t_{k-1})Q_{k-1}} \right) Q_i e^{(t-t_i)Q_i}. \quad (2.1.3)$$

Since component (p, q) in matrix $P(t)$ gives the probability $\mathbb{P}(X_t = q \mid X_0 = p)$, then the random variable $T_{2,\text{same}}$ of the time until coalescence of two lineages sampled in the *same*




Figure 2.1: Connectivity and population size graphs. Visual representation of the t_i , M_i and s_i of a demographic model with $c = 3$ components. We refer to the upper part of the figure as a connectivity graph. The lower part represents population size changes of the general model. In the present study the s_i are identical and equal to 1 and will not be represented with connectivity graphs. The number of islands is inferred but constant, and it is not shown in this figure (although it is usually shown in a separate figure or panel).

deme (initial state 1) has distribution $F_{\text{same}}(t) = P(t)_{(1,3)}$ and density $f_{\text{same}}(t) = D(t)_{(1,3)}$, and for the case where we sample in two *different* demes (initial state 2), then $T_{2,\text{diff.}}$ would have its distribution given by $F_{\text{diff.}}(t) = P(t)_{(2,3)}$ and its density by $f_{\text{diff.}}(t) = D(t)_{(2,3)}$.

The factor matrices in (2.1.3) can be computed in several ways in the general case (see Herbots (1994) or Hobolth et al. (2019)), but considering this particular instance of size 3×3 , they may also be computed directly given an arbitrary Δt and rate matrix Q :

$$e^{\Delta t Q} = \begin{pmatrix} F_{11} & F_{12} & 1 - F_{11} - F_{12} \\ F_{21} & F_{22} & 1 - F_{21} - F_{22} \\ 0 & 0 & 1 \end{pmatrix}, \quad (2.1.4)$$

where:

$$\begin{aligned} F_{11} &= \frac{(\delta + \alpha - 2Ms)\text{exp}_2 + (\delta - \alpha + 2Ms)\text{exp}_1}{2\delta}, & \text{exp}_1 &= \exp\left(\frac{\Delta t(\delta - \alpha)}{2s(n-1)}\right), \\ F_{12} &= (n-1)F_{21}, & \text{exp}_2 &= \exp\left(\frac{\Delta t(-\delta - \alpha)}{2s(n-1)}\right), \\ F_{21} &= \frac{Ms(\text{exp}_1 - \text{exp}_2)}{\delta}, & \delta &= \sqrt{\alpha^2 - 4Ms(n-1)}, \\ F_{22} &= \frac{(\delta + \alpha - 2Ms)\text{exp}_1 + (\delta - \alpha + 2Ms)\text{exp}_2}{2\delta}, & \alpha &= Mns + n - 1. \end{aligned} \quad (2.1.5)$$

With this we can efficiently compute either IICR functions:

$$\text{IICR}_{\text{same}}(t) = \frac{1 - F_{\text{same}}(t)}{f_{\text{same}}(t)}, \quad \text{IICR}_{\text{diff.}}(t) = \frac{1 - F_{\text{diff.}}(t)}{f_{\text{diff.}}(t)}. \quad (2.1.6)$$

In contrast with the parameter space for the unscaled IICR $(\Phi_{\gamma,B})$ in which there is a one-to-one correspondence between a parameter tuple and the corresponding IICR curve, there are only $3\gamma + 3$ independent degrees of freedom in $\hat{\Phi}_{\gamma,B}$, even though there are $3\gamma + 4$ parameters. This notion is formalized by the following lemma.

Lemma 1. *Given any $\hat{\varphi} = (N, n, t_1 \dots t_\gamma, M_0 \dots M_\gamma, s_0 \dots s_\gamma) \in \hat{\Phi}_{\gamma,B}$, then the parameter tuple:*

$$\hat{\varphi}_0 = \left(\frac{N}{C}, n, Ct_1 \dots Ct_\gamma, \frac{M_0}{C} \dots \frac{M_\gamma}{C}, Cs_0 \dots Cs_\gamma \right)$$

is such that:

$$\text{sIICR}_{\hat{\varphi}}(g) = \text{sIICR}_{\hat{\varphi}_0}(g) \quad \forall g > 0,$$

where C is any rescaling factor for which the coordinates of $\hat{\varphi}_0$ are within the bounds B .

Proof. Let us denote by $\pi_{n,M,s}(t)$ the factors e^{tQ} that appear in the definition of $P(t)$ (equations 2.1.3 and 2.1.4). It is easy to verify that for any $C > 0$, we have:

$$\pi_{n,M,s}(t) = \pi_{n,\frac{M}{C},Cs}(Ct). \quad (2.1.7)$$

Indeed, from the expressions in (2.1.5) we can see that the parameter M always appears in the factor Ms , and the parameter t always appears in the factor t/s , which are invariant under the transformation $(M, s, t) \mapsto (\frac{M}{C}, Cs, Ct)$.

Next, given any:

$$\begin{aligned}\varphi &= (n, t_1 \dots t_\gamma, M_0 \dots M_\gamma, s_0 \dots s_\gamma), \\ \varphi_0 &= (n, Ct_1 \dots Ct_\gamma, \frac{M_0}{C} \dots \frac{M_\gamma}{C}, Cs_0 \dots Cs_\gamma),\end{aligned}$$

and any $t > 0$, we can write $P(t)$ from (2.1.3) as:

$$\begin{aligned}P_\varphi(t) &= \left(\prod_{k=1}^i \pi_{n, M_{k-1}, s_{k-1}}(t_k - t_{k-1}) \right) \pi_{n, M_i, s_i}(t - t_i) \\ &= \left(\prod_{k=1}^i \pi_{n, \frac{M_{k-1}}{C}, Cs_{k-1}}(Ct_k - Ct_{k-1}) \right) \pi_{n, \frac{M_i}{C}, Cs_i}(Ct - Ct_i) \\ &= P_{\varphi_0}(Ct).\end{aligned}\tag{2.1.8}$$

where i is the largest index such that $t_i \leq t$ and subsequently $Ct_i \leq Ct$. We denote now by $F_\varphi(t)$ any of $F_{\text{same}}(t) = P_\varphi(t)_{(1,3)}$ or $F_{\text{diff.}}(t) = P_\varphi(t)_{(1,2)}$. From (2.1.8) we have $F_\varphi(t) = F_{\varphi_0}(Ct)$. In order to introduce scaling, we consider an arbitrary effective size N and perform the corresponding change of variable $t = g/2N$. Thus, by having $\hat{\varphi} = (N, \varphi)$ and $\hat{\varphi}_0 = (N/C, \varphi_0)$, we get:

$$\begin{aligned}F_\varphi\left(\frac{g}{2N}\right) &= F_{\varphi_0}\left(\frac{Cg}{2N}\right) \\ \Rightarrow \frac{1}{2N} f_\varphi\left(\frac{g}{2N}\right) &= \frac{C}{2N} f_{\varphi_0}\left(\frac{Cg}{2N}\right) \\ \Rightarrow N \frac{1 - F_\varphi(g/2N)}{f_\varphi(g/2N)} &= \frac{N}{C} \frac{1 - F_{\varphi_0}(Cg/2N)}{f_{\varphi_0}(Cg/2N)} \\ \Rightarrow \text{sIICR}_{\hat{\varphi}}(g) &= \text{sIICR}_{\hat{\varphi}_0}(g)\end{aligned}\quad \blacksquare$$

The implication of Lemma 1 is that when trying to infer all the parameters of $\hat{\varphi}$ simultaneously, the only parameter for which we may get an absolute estimate is n , as the rest of them can only be distinguished up to an unknown re-scaling factor C . Note that this un-identifiability issue is different from the one identified in Mazet et al. (2016) regarding the inability to discriminate between panmictic and non-panmictic demographies with a single IICR. However, we stress here that in practice this is not necessarily an issue because it suffices to fix one of the model parameters (for instance, $s_0 = 1$) to be able to uniquely map any sIICR curve to its parameters. In the case of constant size this is even less of an issue, since *all* deme sizes are fixed to $s_i = 1$ and thus no further considerations are necessary.

2.2. Optimization framework: search algorithm and optimality criteria

The search algorithm explores the parameter space and uses an optimality criterion to select the structured scenario that best explains a given target IICR curve, either scaled or unscaled.

We also assume that the underlying coalescence times for these target IICRs have cumulative distribution F_0 and density f_0 .

Given a target IICR_0 and a parameter space $\Phi_{\gamma,B}$, we want to find a parameter tuple φ in $\Phi_{\gamma,B}$ such that the exact IICR curve corresponding to the model defined by φ (denoted by IICR_φ) approximates IICR_0 . We thus want to find the minimal distance:

$$\min_{\varphi \in \Phi_{\gamma,B}} d(\text{IICR}_0, \text{IICR}_\varphi). \quad (2.2.1)$$

Regarding the distance d , a straightforward definition would be:

$$d(\text{IICR}_0, \text{IICR}_\varphi) = \int_0^\infty |\text{IICR}_0(t) - \text{IICR}_\varphi(t)| w(t) dt, \quad (2.2.2)$$

where $w(t)$ is a weight function that should take into account the natural distribution of the information in an IICR. One reasonable solution for w is to take a quantity proportional to the density f_0 of the coalescence times because it ensures that the integral in (2.2.2) is finite, and also because it assigns more weight to the temporal periods where the target IICR_0 is expected to be more accurate and reliable since more coalescences are likely to have happened.

We thus consider the family of weight functions:

$$w(t) = \frac{f_0^\omega(t)}{\|f_0^\omega\|_1}, \quad (2.2.3)$$

where $\|\cdot\|_1$ is the L^1 -norm and $\omega > 0$ is a *weight-shifting* parameter, with the purpose of dampening (if $\omega < 1$) or exaggerating (if $\omega > 1$) the effect of the weight f_0 . Unless otherwise noted, we take $\omega = 1$, which corresponds in practice to giving more weight to recent periods of the IICR in direct proportion to the density f_0 in an n -island model.

In practice, we need to consider that all we know about IICR_0 is a stepwise discretization over a bounded interval of time, so a numerical approximation of the distance (2.2.2) is required. This includes approximating the density f_0 of the underlying T_2 distribution. Given a division of time into \mathcal{I} intervals $\{[\tau_{j-1}; \tau_j]\}$ for $1 \leq j \leq \mathcal{I}$, where $\tau_0 = 0$ and $\tau_{\mathcal{I}} < \infty$, we can consider a discrete representation of IICR_0 in the form of a collection of \mathcal{I} values $\{y_j\}$ such that:

$$\text{IICR}_0(t) = y_j, \quad \forall t \in [\tau_{j-1}; \tau_j), \quad 1 \leq j \leq \mathcal{I}. \quad (2.2.4)$$

We can use this form to compute a numerical approximation for the integral in (2.2.2). For instance, a first degree approximation would yield:

$$d(\text{IICR}_0, \text{IICR}_\varphi) = \sum_{j=1}^{\mathcal{I}} |y_j - \text{IICR}_\varphi(\tau_j)| w(\tau_{j-1}) (\tau_j - \tau_{j-1}).$$

As for the values of $w(\tau_j)$, notice that from (1.4.4) we have the identity:

$$f_0(t) = \exp\left(\int_0^t \frac{-d\tau}{\text{IICR}_0(\tau)}\right) / \text{IICR}_0(t),$$

which, using the representation (2.2.4), can be discretized into:

$$f_0(0) = 1/y_0,$$

$$f_0(\tau_j) = \exp\left(\sum_{k=1}^j \frac{\tau_{k-1} - \tau_k}{y_k}\right) / y_j, \quad 1 \leq j \leq \mathcal{I}.$$

We then have that for any given ω , $w(\tau_j)$ can be expressed as:

$$w(\tau_j) = f_0^\omega(\tau_j) / \sum_{k=1}^{\mathcal{I}} f_0^\omega(\tau_k).$$

An alternative option for the definition of d in (2.2.1) could be one that takes into account the ultimate visual nature of the curve fitting task. Assuming that the points $\{\tau_j\}$ are log-distributed and that they will be used for visualization purposes in a horizontally log-scaled plot like Figure 2.11, then the definition:

$$d(\text{IICR}_0, \text{IICR}_\varphi) = \sum_{j=1}^{\mathcal{I}} |y_j - \text{IICR}_\varphi(\tau_j)|, \quad (2.2.5)$$

captures the perceived visual difference between the plots of the two curves. We distinguish distance (2.2.2) from (2.2.5) by denoting them d_ω and d_{visual} respectively. We keep both definitions because we found that the weighted family of distances generally performs better than the visual distance under certain validation tests, but also that the d_{visual} distance can be used to choose the optimal weight parameter in d_ω (see Figure 2.24).

Regarding the optimization problem (2.2.1) itself, we use the *Differential Evolution* method (Storn and Price, 1997). As a genetic meta-heuristics, this algorithm maintains and evolves (using mutation and recombination parameters) a population of solutions iteratively. As a global optimization algorithm, it features mechanisms for escaping local *optima* of the parameter space. In §2.2.1 we explore the potential effects on the inference results of tuning some of the parameters provided by this algorithms implementation. For our validations, the method runs by using multiple steps of optimization which we refer to as *rounds*. In addition, we stress that the method should be used multiple times on real data sets. We set a maximum number of allowed rounds, as well as a tolerance ε for the distance which controls the minimum number of rounds.

2.2.1. Comparison of optimization parameters

Here we explore the effect of various parameters of the optimization algorithm on the speed of convergence of the inference process.

The most important parameters that affect the search process and convergence criteria are:

- strategy: it can be one of 12 possible values (see Table 1) and controls how each new generation of solutions are computed from the previous one. Default is 'best1bin'.
- maxiter: maximum number of iterations the algorithm will perform before forcing convergence. Default is 5000.
- popsize: number of simultaneous solutions during any given generation. Default is 15.
- tol: relative tolerance for search convergence. The convergence criteria is met when the standard deviation of the solutions within a generation is smaller than tol times the average energy (in our case, distance) within that generation. Default is 10^{-2} .
- mutation: per-generation mutation rate of the solutions. The default behaviour is to draw a random value from a uniform distribution in $[0.5, 1]$ each generation.
- recombination: per-generation recombination rate of the solutions. Default is 0.7.

We selected 10 random simulated scenarios with unscaled IICR and $c = 4$ components from the set of exact IICR validations that had not converged before 500 rounds (these would correspond to scenarios off the diagonals in Figure A.4). For each one of them, and for each of the 12 possible values for the strategy parameter, we attempted another 100 rounds. Table 1 shows the number of rounds it took for each of the 10 scenarios to converge in each case (a value of 100 means that convergence was not reached).

Strategy	#1	#2	#3	#4	#5	#6	#7	#8	#9	#10	Total
best2exp	17	1	100	2	100	22	3	100	57	22	424
best2bin	6	35	100	2	11	93	100	4	33	100	484
currenttobest1exp	2	80	100	5	100	69	12	100	8	26	502
best1exp	3	76	100	21	4	100	10	100	62	100	576
rand1exp	4	100	100	10	100	14	100	100	100	38	666
randtobest1exp	32	100	100	9	100	78	100	30	100	100	749
currenttobest1bin	21	100	100	27	100	100	100	94	63	88	793
best1bin	27	100	100	28	54	100	100	100	100	100	809
rand2exp	40	100	100	100	100	24	100	100	100	100	864
rand1bin	100	100	100	3	100	100	100	100	100	100	903
rand2bin	96	100	100	23	100	100	100	100	100	100	919
randtobest1bin	92	100	100	100	100	100	100	100	100	100	992

Table 2.1.: Results of varying the strategy parameter of the differential evolution algorithm on the speed of convergence of 10 difficult demographic scenarios with $c = 4$ components.

As can be seen, strategy 'best2exp' was best with a combined number of 424 rounds for the 10 scenarios.

Next, using this optimal strategy parameter, we tried one alternative value for the rest of the optimization parameters at a time, again allowing a maximum of 100 rounds. The alternative values were as follows:

- maxiter was changed from 5000 to 10000.
- popsize was changed from 15 to 50.
- mutation was changed from random sample in $[0.5, 1]$ to random sample in $[0.5, 1.7]$.

- recombination was changed from 0.7 to 0.9.

The results, shown in Table 2, suggest that these parameters should be left at their default values.

Parameter	no. 1	no. 2	no. 3	no. 4	no. 5	no. 6	no. 7	no. 8	no. 9	no. 10	Total
popsize	100	100	100	100	100	100	100	100	100	100	1,000
recombination	100	100	100	100	100	100	100	100	100	100	1,000
mutation	100	100	100	100	100	100	100	100	100	100	1,000
max-iter	100	100	100	100	100	100	100	100	100	100	1,000

Table 2.2.: Results of varying the popsize, recombination, mutation and max-iter parameters of the differential evolution algorithm on the speed of convergence of 10 difficult demographic scenarios with $c = 4$ components.

2.3. Validation framework and results

We applied the SNIF inferential method to target IICRs generated under piecewise stationary n -island models of increasing complexity (i.e., number of components) and with known parameter values (N, n, t_i, M_i) and then compared the inferred parameter values to those actually used.

In what follows we present various ways of generating random demographic scenarios and then computing appropriate IICR curves from them for use in validation.




Figure 2.2: Flowchart of the validation procedures. Starting from a parameter space $\hat{\Phi}_{\gamma, B}$ we use one of two sampling methods (§2.3.1) to generate a demographic history $\hat{\varphi}$ defined (for the scaled case) by the parameters (N, n, t_i, M_i) . We then compute the IICR of that demographic history using one of three methods (§2.3.2) to obtain the target IICR. After that, we run the inference algorithm on this target IICR curve (using wider bounds than those in B) to obtain an estimated (or inferred) demographic history $\hat{\varphi}_0 = (N^0, n^0, t_i^0, M_i^0)$, which we then compare to the known $\hat{\varphi}$ in order to assess the accuracy of the inference methodology (§2.3.4 and §2.3.5).

2.3.1. Sampling the parameter space

Given a parameter space $\Phi_{\gamma, B}$ (we only discuss the unscaled case here for brevity, but the same principles apply to a scaled parameter space), we sample demographic scenarios from

which we compute the corresponding IICRs. We used two sampling strategies which we call continuous and discrete sampling.

Continuous sampling. Assuming that we want to realise L independent tests, this sampling strategy consists in using uniform or log-uniform distributions for each of the $3\gamma + 3$ random variables:

$$\begin{aligned}
 n &\sim U\{n_{\min}, n_{\min} + 1, \dots, n_{\max}\}, \\
 t_1 &\sim LU_{10}[t_{1\min}; t_{1\max}], \dots, t_\gamma \sim LU_{10}[t_{\gamma\min}; t_{\gamma\max}], \\
 M_0 &\sim U[M_{0\min}; M_{0\max}], \dots, M_\gamma \sim U[M_{\gamma\min}; M_{\gamma\max}], \\
 s_0 &\sim U[s_{0\min}; s_{0\max}], \dots, s_\gamma \sim U[s_{\gamma\min}; s_{\gamma\max}],
 \end{aligned} \tag{2.3.1}$$

where U denotes a uniform distribution (discrete in the case of n and continuous for the rest) and LU_{10} denotes a log-uniform distribution of base 10. This distribution is used for sampling the times of changes in a logarithmic space in order to take into account the natural distribution of information in an IICR.

After sorting the times t_i , we can define the L sampled scenarios by constructing, for $1 \leq j \leq L$, the tuple $(n^j, t_1^j \dots t_\gamma^j, M_0^j \dots M_\gamma^j, s_0^j \dots s_\gamma^j)$. This sampling strategy makes it very unlikely to sample exactly the same parameter values twice or to sample exactly the same M_i values in two consecutive components. However, it sometimes produces demographic scenarios in which consecutive t_i and/or M_i values may be close to each other, and thus difficult to distinguish. This makes it thus harder on our inferential framework compared to cases where we would chose contrasted scenarios with clearly separated events with major changes in migration rates. In other words, our inferential method was sometimes inferring parameters in the case of extremely difficult scenarios as we show below.

In §2.3.4 we show the results obtained using this sampling method with $L = 400$ scenarios. The bounds for sampling and inferring are shown in Table 2.3. In particular, we note that in practice we disallow deme size changes by fixing the bounds of the s_i to 1, which consequently reduces the parameter space to just $2\gamma + 2$ parameters.

	simulation		inference	
	min	max	min	max
n	2	40	2	50
t_i	0.1	50	0.01	100
M_i	0.1	50	0.05	60
s_i	1	1	1	1
N	10^3	10^3	10^2	10^4

Table 2.3.: Parameter bounds used during the generation of demographic scenarios for validation.

Discrete sampling. Here we sampled $L = 100$ independent scenarios from the same parameter space, but using the following set of predefined values:

$$\begin{aligned} n &\in \mathbf{n} = \{2, 5, 10, 15, 20\}, \\ t_i &\in \mathbf{t} = \{1/10, 1/2, 1, 2, 5, 10, 20, 50\} \quad \forall i, \\ M_i &\in \mathbf{M} = \{1/10, 1/5, 1/2, 1, 2, 5, 10, 20, 50\} \quad \forall i, \\ s_i &\in \mathbf{s} = \{1\} \quad \forall i; \end{aligned}$$

The inference process was, however, done within the continuous space. For instance, under this validation scheme (see §2.3.5) we only simulated data with 2, 5, 10, 15, 20 islands but the inference process always allowed n to take any value between 2 and 50. The choice of the L independent simulated data sets was done using the following procedure. We first considered the following cartesian product of dimension $2\gamma + 2$:

$$K = \mathbf{n} \times \mathbf{t}^\gamma \times \mathbf{M}^{\gamma+1} \times \{1\}.$$

and then uniformly drew L tuples from K without replacement. We then sorted the sampled event times obtaining thus a set of L demographic scenarios. We drew randomly (without replacement) from the set K , rejecting scenarios with identical M_i values in two consecutive components, until we reached L accepted scenarios.

2.3.2. The three types of target IICRs

We explored three different types of target IICRs (see Figure 2.3) that could be obtained given a scenario $\varphi \in \Phi_{\gamma,B}$. All IICRs were discretized so as to be comparable to PSMC plots (see equation 2.2.4).

Validating SNIF on PSMC plots across the parameter space described above would be extremely time consuming as it would require simulating genomes and then running the PSMC method (or other related methods) on these genomes before applying our approach. We thus only ran the PSMC method in the case of the scenarios inferred for the human data so as to integrate the uncertainty due to the PSMC inferential process. The issue of uncertainty is crucial but our aim is not to test the PSMC or other inferential methods. To clarify this we explain below the different types of IICR that could be computed given a scenario $\varphi \in \Phi_{\gamma,B}$.

Exact IICR We can compute the IICR for any n -island model at any time value t , but to generate input data we need a discretization as in (2.2.4), so considering that we take a log-distributed sample of size l in the interval $[t_{\min}, t_{\max}]$, we end up with a suitable IICR_0 . Note that even though this IICR has been discretized, its values are exact within machine precision, so it is still an artificial product compared to real data.

For the validations using the exact IICR in §2.3.4 we chose for the distance tolerance between a target and an inferred IICR a value of $\varepsilon = 10^{-10}$ for the unscaled IICRs and an equivalent value of $\varepsilon = 10^{-7}$ for the scaled IICR (since the simulated N was always 1000). It should be noted that this value of ε is quite small even for double-precision floating-point

arithmetic, and thus is only a reasonable choice for validation using exact IICRs (i.e., those where the distance could theoretically be zero).

T-sim IICR The T-sim IICR is obtained by simulating a finite collection of T_2 realizations using `ms` and then building an empirical IICR as in Mazet et al. (2016), using the Kaplan–Meier estimator (Kaplan and Meier, 1958), with log-distributed times. We stress that `ms` scales time in units of $4N$ generations whereas our models use a scale of $2N$ generations (see the example in Figure 2.3), so this must be kept in mind when writing `ms` commands.

Seq-sim IICR We simulate genomic sequences with `ms` and then apply the PSMC method for obtaining the $IICR_0$ to be used by the inference method. Since simulating genomes and performing PSMC analyses is significantly more time consuming than the other two methods, we limited ourselves to validating the Seq-sim IICR for the human PSMC based scenarios that we obtain after performing the demographic inference described in §2.4. The results of this step are shown in §2.4.2 of the Supplementary Materials.




Figure 2.3: The three ways of obtaining a target IICR. The panels represent different types of discretised target IICRs for the same two-components demographic history with $N = 10^3$, $n = 4$, $t_1 = 7$, $(M_i) = (0.5, 1.0)$ and $(s_i) = (1.0, 1.0)$. **(a)** Exact IICR, computed as per section 2.1.2 and discretized to 32 intervals. In gray we show a smoother discretization and keep it in the other panels for reference. This is obtained using the approach of Rodríguez et al. (2018). **(b)** T-sim IICR, obtained by simulating 2×10^4 T_2 samples using the `ms` command `ms 2 20000 -T -L -I 4 2 0 0 0.5 -eM 3.5 1 -eN 3.5 1` and later scaling the IICR using the value of N . This is the approach used in Chikhi et al. (2018) **(c)** Seq-sim IICR, obtained by running PSMC on a genome simulated using the `ms` command `ms 2 100 -t 100 -r 20 2000000 -p 8 -I 4 2 0 0 0.5 -eM 3.5 1 -eN 3.5 1` and later scaling with mutation rate $\mu = 1.25 \times 10^{-8}$. This is obtained by running the PSMC method of Li and Durbin (2011)

2.3.3. Quantifying the inference error

In order to quantify the inference error incurred during the continuous-sampling validation phase, we measured the normalised root-mean-square deviation (nRMSD) between the simulated and inferred parameter values for each parameter of the demographic models ($c = 1$ to $c = 6$ components). These values can be seen in the lower-right corner of every sub-panel in figures A.1 to A.6 and A.8 to A.13. They are also summarised in Figure 2.4. We note that the number of islands n and the effective size N (in scaled scenarios) is very well inferred regardless of the number of components c . On the other hand, the inference accuracy of the connectivity rates M_i does get gradually worse when increasing the number of parameters (see section Validation results using exact target IICRs in the main text).




Figure 2.4: Normalised root-mean-square deviation of the inferred parameters during validation. The top part of each sub-panel corresponds to the parameters of unscaled scenarios, and the bottom part to scaled scenarios. The number of components c of the scenarios is in the horizontal axis in every case.

In an effort to quantify the component misidentification phenomenon, we computed the correlation between 100 randomly sampled pairwise values of simulated and estimated parameters (in this case, the simulated M_2 parameter in unscaled scenarios of $c = 5$ components with the inferred M_j parameters of the same scenarios). The key insight underlying this test is that when a parameter is badly estimated, it may be due to the fact that the method estimated the value from the component that is either just before or just after it. Conversely, the value of M from other components should be non-correlated, since values are taken at random within the range of allowed values. Now, we do not know if the method used the component just before or just after, as this may change from simulation to simulation. To solve this problem we computed the correlation (r^2) by using either both values or just the value that was closest. We did exactly the same for the values taken from other components that are not expected to be correlated. We find indeed that when no correlation is expected there is no correlation, but when we take the best value the correlation increases, and it increases much more when it is a neighbouring value. Additionally, we found that this effect is significantly amplified (r^2 increased from 0.37 to 0.82) when we exclude from our sample the tests where the inferred rates were within 10% of the simulated values,

further indicating that this effect is present mostly when there is a large mismatch between the simulated and inferred M values. These results are displayed in Figure 2.5.




Figure 2.5: Parameter misidentification. We measure the r^2 correlation between M_2 and \hat{M}_j in a sample of 100 scenarios with $c = 5$ components. In ‘random global’ we draw M_j from components 0, 1, 3, and 4. In ‘random far’ we draw M_j from components 0 and 4. In ‘best far’ we draw M_j from components 0 and 4 again, but always keep the best fitting value of the two. Finally, in ‘best near’ we draw M_j from components 1 and 3 and always keep the best fitting value. ‘Filtered at 10%’ indicates that the scenarios where the inferred value was within 10% of the simulated value were excluded.

2.3.4. Validation results using exact target IICRs

A first set of figures (Figures 2.6, A.1 to A.6, and A.8 to A.13) represents the simulated and inferred parameter values on the horizontal and vertical axes, respectively, using the continuous sampling strategy. As explained in §2.3.1, the range of possible values in the inference process was always wider than the range used for the simulated values (see Table 2.3 for the exact values). We quantified the inference error for each parameter by computing the Normalized Root-Mean-Square Deviation (nRMSD). This value is displayed in the lower-right corner of each panel of the previously mentioned figures, and summarized for all parameters in Figure 2.4. For reference, we also highlight the $y = x$ line, indicating what would be a perfect inference, and the region corresponding to 10% of relative error around this line (50% of relative error in the case of the t_i parameters). A summary of how many tests fall within this margin of error (and others) is displayed in figures A.7 and A.14. Altogether, these figures always show the data points near the $y = x$, hence suggesting that the inferred parameter was identical or very close to the simulated parameters. This is particularly obvious for all the parameters corresponding to scenarios with up to four components, where the nRMSD values stayed below 0.5 (the case of the t_i parameters is exceptional, since the exponential distribution of its range disproportionately affects the error measures). For instance, Figure 2.6a shows the results for a model with three components, in which there is a nearly perfect match (nRMSD close to or below 0.1 for the non- t parameters) between the simulated and inferred values for the model parameters. For five- and six-component scenarios the results are still nearly perfect for most of the simulated scenarios but we observe an increasing number of cases (i.e., simulated scenarios) where the parameters are poorly estimated, with

the exception of n , M_0 (and N for scenarios with scaled IICRs) which are almost always well estimated also in such cases (nRMSD values consistently close to or below 0.1). In particular, we can appreciate a gradual degradation of the correspondence between simulated and inferred migration values when the number of components c increases, as the nRMSD monotonically increases to over 0.5 for $c = 6$. These cases can be identified by the dots scattered in the different panels. They start to appear in scenarios with three components, but their number grows with the number of components.




Figure 2.6: Scatter plots of simulated and inferred parameters. Panel (a) corresponds to scenarios with $c = 3$ components, and (b) to scenarios with $c = 6$ components. The different sub-panels represent the simulated (horizontal axis) versus inferred (vertical axis) parameter values for all the parameters (or a representative selection of parameters in the case of panel (b)) of $L = 400$ unscaled simulated scenarios.

These poorly estimated parameters are surprising given the near perfect estimation obtained for most parameter combinations. This is particularly striking because these dots do not seem to be distributed in any clear area of the parameter space. We see at least two possible and non-exclusive interpretations for this result. One is that the search algorithm had not yet *converged* when the maximum number of rounds was reached.

The maximum number of rounds was set to 500 in all simulations because we had found that less than 50 rounds were more than enough in the first tests carried out with one or two components. The search algorithm might however need more than 500 rounds to reach the optimal solution for scenarios of increasing complexity. We thus asked whether the maximum the number of rounds had been reached in the scenarios analysed and whether the proportion

of scenarios with 500 rounds increased with the number of components. We found indeed that the proportion of simulations for which that maximum was reached increased with the number of components. For instance, all five- and six-component scenarios stopped their parameter search at 500 rounds, hence suggesting that at least some had not yet reached an optimum solution. For the cases with one- and two-component scenarios, all 800 independent simulations reached convergence in less than 150 rounds (see Figure 2.7). Again, the choice of the tolerance ϵ plays a role in these results, and selecting larger tolerances will tend to produce earlier convergence in general, but not necessarily better results.




Figure 2.7: Number of performed rounds during validations with exact IICRs. The top panel shows the number of rounds (up to a maximum of 500) that the method required before converging to a scenario with a distance value smaller than the tolerance $\epsilon = 10^{-10}$ for the unscaled case. The bottom panel shows the same information for scaled scenarios and a tolerance of $\epsilon = 10^{-7}$. We represent in different colors the curves corresponding to scenarios with different simulated (and inferred) number of components, ranging from $c = 2$ up to $c = 6$; the same data-set represented in figures A.1 to A.6 (top panel) and figures A.8 to A.13 (bottom panel). Higher number of rounds is not correlated with worse fit when the maximum number of rounds is not reached; and when it is reached it is not necessarily an indication of bad fit either, although all instances of bad fit stem from inferences that reached the maximum allowed number of rounds.

As a test we randomly identified a couple of scenarios with six components that had bad estimates and re-ran the algorithm with 5000 rounds. We found that the distance value consistently decreases with more rounds (see Figure 2.8), but the inferred parameter values may not converge to the simulated ones because with more components there is a higher probability that two consecutive simulated M_i values are very close, thus making the corresponding event time challenging to infer.

The second possible reason for the poorly estimated parameters in Figure 2.6 may be related




Figure 2.8: Convergence of the inferred IICR with the progression of the inference rounds. We measured the best distance achieved ($\omega = 1$) as a function of the completed rounds for up to 5000 rounds. The two curves correspond to two $c = 6$ component scenarios that did not previously converge ($\varepsilon = 10^{-10}$) in 500 rounds (marked with a vertical dashed line). We note that there is a clear trend for the distance between the IICR curves to decrease with more inference rounds, but there is also diminishing returns to performing a very large number of rounds, since the issue of component misidentification may be an insurmountable difficulty for some scenarios (see main text for a more detailed discussion).

to the fact that some simulated components may have a short duration that do not leave a significant mark on the IICR curve, thus leading them to be “skipped”. We refer to this issue as component misidentification, which could lead to a particular estimated parameter to be plotted in the wrong panel. For instance, the method may miss the first change in migration rate at t_1 and identify the second change in migration at t_2 . In such a case the method will assign the inferred t_2 value to the set of inferred t_1 values and plot it in the t_1 panel. This wrongly assigned t_2 value will thus appear away from the diagonal in the t_1 panel even if it was well-estimated. Such misassignment cases for one parameter will also have consequences for the M_i plots, and thus will generate several misassignments across panels. They are also expected to increase in frequency as the number of components increases and as the t_i values become closer to each other. This phenomenon can be observed clearly in the right panels of Figure 2.9. We also present an attempt at quantifying it for the case of $c = 5$ in Figure 2.5. One way to mitigate the effect of this misassignment issue in the analysis of the results is to visualize the simulated and inferred scenarios using what we call a connectivity graph. This connectivity graph represents the times at which migration changes against the values of the migration rates. Such connectivity graphs are featured in the next section.

2.3.5. Validation results using T-sim IICRs

The connectivity graphs and IICR plots obtained from simulated T_2 values show that again the scenarios are generally very well reconstructed (Figure 2.9 and Figures B.1 to B.15).

In Figure 2.9 the connectivity graphs obtained for all the scenarios simulated with three and four components show that the inferred times at which migration rates changed (green vertical lines) and the inferred migration rates (green horizontal lines) are generally overlapping close to the simulated values (dotted vertical and horizontal gray lines). In the right panels of this figure, we show a subset of the inferred migration histogram (in red). Namely, we show the distributions of the migration values that were inferred for components with a simulated migration value of $M_i = 10$ for panel (a) and $M_i = 1$ for panel (b). This allows us to better appreciate the variance of the inferred migration values in the context of the simulated ones, as well as the component misassignment effect mentioned earlier. Indeed, we note here that the incorrectly inferred migration values are clustered around other simulated values, indicating a mismatch in a particular component assignment which does not affect the rest of the inferred demographic history (we present a quantification of this effect for a limited case in Figure 2.5).

For example, consider the right sub-panel of (a). We see that most repetitions correctly inferred a value close to $M = 10$ for the components with that simulated migration rate. However, there were cases where a given component i was simulated with a migration rate of $M_i = 10$, but it was missed entirely (maybe because it did not generate a very different IICR or because it had a short duration), and thus the inferred migration value for component i ultimately reflected either M_{i-1} or M_{i+1} . In panel (b) we can observe the same effect with higher intensity because with more components it is more likely for them to be misaligned or misidentified during inference. See Figure 2.5 for a quantification of this effect on scenarios of $c = 5$ components.

These connectivity graphs (and the one obtained for five components shown in Figure B.13) also show that there are regions of the parameter space where the green lines are more widely distributed. For instance, in the recent past of Figure 2.9b ($t_i < 10^{-3}$ generations) when the simulated M_i value was 0.1 or 0.2 the inferred values seem to vary between 0.05 and 0.3, suggesting that the method identifies periods with low migration rates but that the exact value is difficult to estimate properly, at least in the recent times. These graphs however summarize extremely different scenarios, including scenarios in which consecutive M_i values may be similar. We thus stress that the quality of the inference is dependent on the timing of the changes in migration rates and on the size of the change in M_i values.

Figure 2.10 shows the results for four different scenarios. In each of the four panels, we represented the inferred and target IICR plots, connectivity graphs, N (the size of each the islands) and n (the number of islands) for the corresponding model. Panels (a) and (b) correspond to three- and four-component scenarios, whereas panels (c) and (d) show the results for two five-component scenarios, one for which we obtained very good estimates and one for which the estimates were poorer. In panels (b) and (c) the inferred and simulated M_i and t_i values are on top of each other as can be seen in the connectivity graphs. Similarly, N and n are also well estimated. Here, the IICR plots also overlap, although this does not always

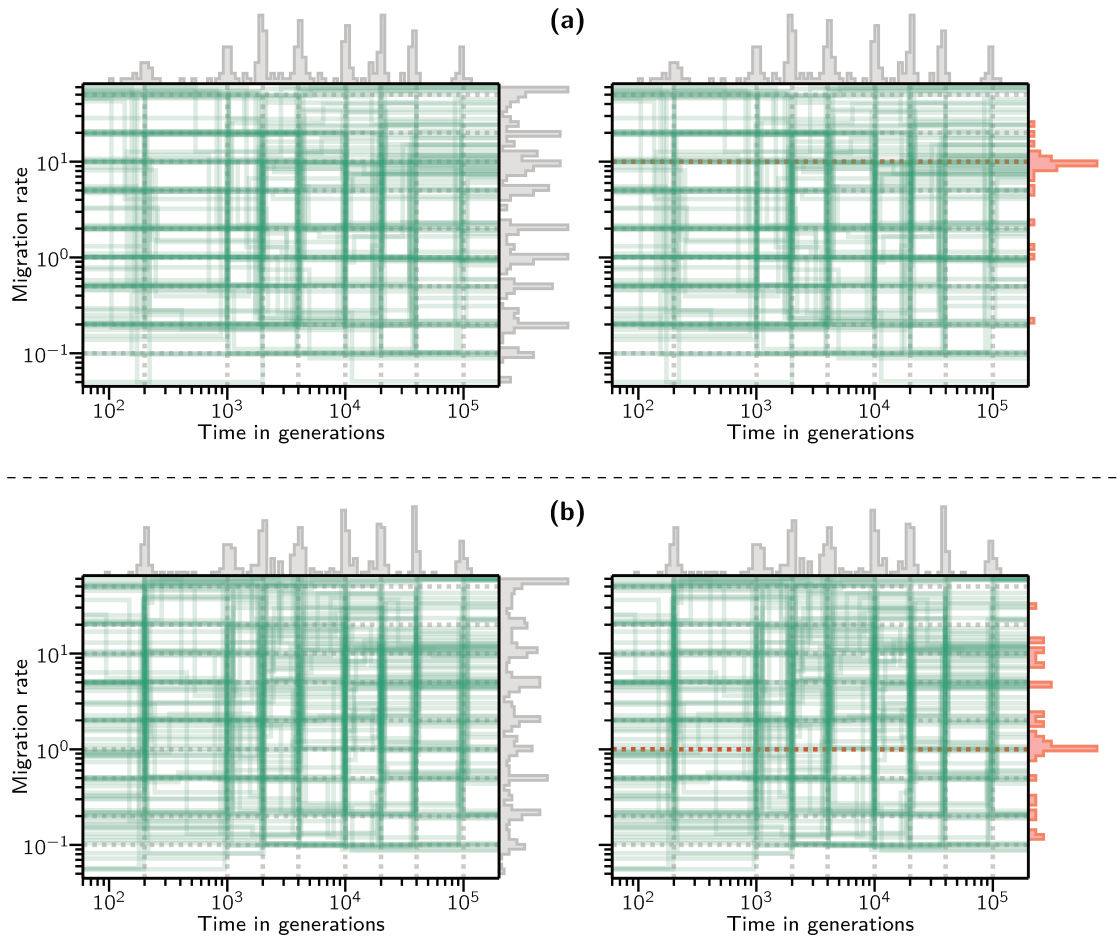


Figure 2.9: Connectivity graphs of 100 independently inferred histories obtained by sampling for each scenario from the values indicated by the dotted lines. **(a)** scenarios with $c = 3$ components. **(b)** scenarios with $c = 4$ components. The right sub-panels show a side histogram with only the inferred migration rates for those components with a specific simulated migration rate (10 for **(a)** and 1 for **(b)**).

guarantee perfect parameter estimation, as is the case in panels (a) and (d). Interestingly, in panel (a) the first change in migration rate (at $t_1 = 200$ generations) is estimated at around 900 generations due to the stochasticity of the IICR plot. This appears to generate some variance in the estimates of N and n but the connectivity graph shows the same trend (increasing connectivity) as in the simulations. In the case of panel (d) we can see that the method had some difficulty in estimating several of the changes in M_i values. This is not surprising as some of the randomly simulated changes do not seem to lead to major changes in the IICR curves. This generates again some variance in the N and n estimates. We also observe a significant variance in the connectivity graph even if several runs overlap nearly perfectly with the simulated connectivity graph.

Altogether the validation tests and figures above suggest that our framework is able to infer changes in connectivity under the n -island model, and that some scenarios can be extremely well inferred whereas others may be more difficult depending on their effect on the IICR plots.

We also observe that for real data it may be helpful to run the analyses for a varying number of rounds, since too few rounds may negatively affect the quality of the fit. Also, once a scenario has been inferred, it is advisable, as an additional validation step, to simulate data under the inferred scenario and use SNIF to re-infer the parameters. This is what we do with the real human data in the next section.




Figure 2.10: Simulated and inferred IICR plots, connectivity graphs, N and n . The four panels correspond to four different scenarios. **(a)** a $c = 3$ components scenario **(b)** a $c = 4$ components scenario **(c)** a $c = 5$ components scenario **(d)** a $c = 5$ components scenario. The left part of each panel represents the target and inferred IICRs (top), and the connectivity graphs (down). The right half of each panel shows the simulated and inferred values for n (top) and N (down). In each IICR graph, the ragged blue line represents the target IICR whereas the red lines represent 10 independently inferred IICRs. The vertical blue and red lines are located at the simulated and inferred values of the event times t_i , respectively. In the connectivity graphs, the black and green lines represent the simulated and inferred connectivity scenarios, respectively. The simulated n and N values are represented by black circles whereas the inferred values for the corresponding parameters are represented by red and green full circles and by grey histogram bars.

2.4. Application to humans

We applied our method to the human genomes published in the great apes study by Prado-Martinez et al. (2013). Namely, we used the PSMC files of five sampled humans identified as Dai, French, Karitiana, Sardinian and Yoruba. Figure 2.11 shows the IICR curves of these human samples, scaled using a mutation rate of $\mu = 1.25 \times 10^{-8}$ and a generation time of 25 years (§2.4).




Figure 2.11: Five human PSMC plots. All five PSMC curves were obtained from the study of Prado-Martinez et al. (2013). The scaling was done as recommended by Li and Durbin (2011) using a value of $\mu = 1.25 \times 10^{-8}$ and a generation time of 25 years. No further processing was necessary since *.psmc* files contain the final result of the effective size or IICR estimate.

For each human PSMC curve we performed demographic inference independently within the following bounds:

$$\begin{aligned}
 n &\in \{2, 3, \dots, 100\}; \\
 t_i &\in [4 \times 10^2, 4 \times 10^5] \quad \forall i; \\
 M_i &\in [1/20, 20] \quad \forall i; \\
 s_i &= 1 \quad \forall i; \\
 N &\in [10^2, 10^4].
 \end{aligned} \tag{2.4.1}$$

The bounds for the t_i are specified in generations, so given a generation time of 25 years, we effectively allowed for the inference of demographic events between 10 thousand and 10 million years ago. Regarding the number of components, we choose $c \in \{2, 3, 4, 5\}$; i.e., between one and four demographic events in agreement with Mazet et al. (2016) who suggested that a minimum of three events were necessary to explain the two humps, and in agreement with our validation simulations which suggest that inference above five components is difficult.

Some of the analyzed PSMC curves exhibit an increase in effective size in the recent past, which could be due to a genuine population growth as noted by Mazet et al. (2016). Given that we choose to specifically rule out changes in deme sizes, we account for this fact by

running every inference a second time, ignoring this period of possible recent expansion. This is accomplished using an option that allows to limit the interval where the distance function is computed. In this case, we restricted both this range and the bounds for the t_i to be between 50 thousand and 10 million years ago, thus ignoring any population size change that may have happened in the last 50,000 years. Note that this option is also useful to ignore very ancient sections of the PSMC plots which may be difficult to trust.

Since real human PSMCs are unlikely to have been generated by an n -island model, the default value of ω used for simulated data may not be the most appropriate, and we thus performed inferences with $\omega \in \{1, 0.5, 0.2\}$. Decreasing values of ω give increasing weight to the most ancient part of the PSMC (see the weighted distances (2.2.3)). The resulting inferred demographic scenarios are shown in section §2.4.1.

To validate the inference process using PSMC outputs, we generated 10 Seq-sim IICRs corresponding to the inferred demographic scenarios for the French, Karitiana and Yoruba individuals. We exclude the Dai and Sardinian populations from this analysis because their corresponding inferred histories are similar to the other three (see Figures 2.18 through 2.24). For each one of the selected scenarios we simulated $nreps = 30$ chromosomes of length $L = 10^8$ base pairs, using the effective size N inferred by SNIF, a per-base per-generation mutation rate of $\mu = 1.25 \times 10^{-8}$ (see (Scally and Durbin, 2012) and references therein). We kept for consistency the scaled recombination rate of $\rho = \theta/5$ as in Li and Durbin (2011), and we ran the `ms` command with $\theta = 4\mu LN$ using:

```
ms 2 nreps -t  $\theta$  -r  $\rho$  L -p 8 -I . . .
```

where the rest of the command follows according to the inferred demography (see Figure 2.3 for a reference). After that we prepared a `.psmcfa` file as input for PSMC, always choosing a bin size of $s = 100$. Then we ran the PSMC with the command:

```
psmc -N25 -t15 -r5 -p "4+25*2+4+6" . . .
```

following Li and Durbin (2011) on human data. We then scaled the resulting curve using the information in the generated `.psmc` file and used these PSMC curves as targets to determine whether we could indeed infer the parameters used for such complex scenarios.

We also applied SNIF to genomic data simulated under the scenarios used to describe recent human evolutionary history by Gutenkunst et al. (2009) and Noskova et al. (2019). Here, we thus ask the following two questions: if human evolution were indeed closer to such splitting models, would our method infer again an n -island model with similar parameters to those inferred from the humans PSMCs? additionally, do these models generate IICR plots that are similar to the human PSMCs? The results of these validations are shown in §2.4.2 of the Supplementary Materials.

2.4.1. Results

Figures 2.12 and 2.13 show the results of using SNIF on the human data.

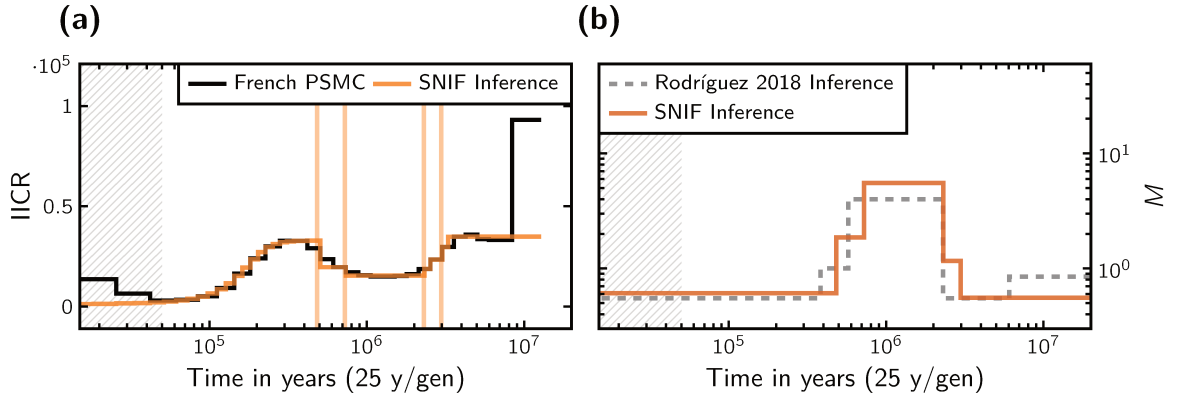


Figure 2.12: Results of performing demographic inference on the French PSMC curve. Panel (a) shows the IICR plot inferred for $c = 5$ components and a weight parameter of $\omega = 0.2$. The vertical lines represent the inferred times of the demographic events. Panel (b) shows the connectivity graph for the same inferred scenario. As a reference point, the connectivity graph of the scenario proposed in Rodríguez et al. (2018) is also shown. The vertical axis in panel (b) represent migration rates (M).

In Figure 2.12, panel (a) shows the PSMC curve of the French individual (scaled with a mutation rate of $\mu = 1.25 \times 10^{-8}$ and a generation time of 25 years) together with the best fitting IICR for the model with $c = 5$ and $\omega = 0.2$. Panel (b) shows the connectivity graphs of the same inferred demographic scenario. We note that the connectivity pattern consisting of a period of relatively high connectivity between roughly 500 kya and 2 Mya agrees with previous results published in Rodríguez et al. (2018). Note that this study used a mutation rate of $\mu = 2.5 \times 10^{-8}$ and not 1.25×10^{-8} . The absolute timing of events and deme sizes are thus different (Rodríguez et al., 2021).

The full set of results related to the inference of human demographics can be found in Figures 2.18 to 2.24, which were placed in the Supplementary Materials for the sake of brevity. The most striking feature of this extended set of plots is the sensitivity of the fit to the value of the weight-shifting parameter ω . Smaller values allow the optimizer to distribute more of the demographic events towards the ancient past and thus allows this region to be better fitted by the inferred IICR. This functionality (together with being able to ignore certain parts of the plots for the computation of the distance function) can be used to make explicit the knowledge (or beliefs) of the researcher regarding the accuracy of the PSMC curve. We notice that the Yoruba individual cannot be well fitted in the recent past for any value of ω , even outside of the designated period of recent population expansion.

Figure 2.13 shows in panel (a) the number of demes n and in panel (b) the reference size N that were inferred from each of the five fitted human PSMCs. Of note here is that all individuals except the Yoruba show a consistent value for these inferred parameters across both number of components and value of ω . The larger variance of the estimated values for the Yoruba individual suggests that a symmetrical island model may not be enough to explain the patterns of diversity in all five sampled human IICRs.

Figures 2.14, 2.15, and 2.16 show the results of applying our method to the IICR curves associated with the demographic model for human expansion published by Noskova et al.




Figure 2.13: (a) Inferred number of islands n and (b) inferred reference sizes N for each human population and each used combination of the weight parameter ω and number of components c (only 4 and 5 are shown here). The bars with the darker color, marked 'restricted' in the legend, correspond to inferences realized with the option of ignoring recent population expansion.




Figure 2.14: Application of our inference method to a tree-like human demographic scenario with three modern populations. (a) IICR plots showing the resulting IICR curve of the European population under this model and the inferred IICR curve obtained with our method (where the recent period of human expansion was ignored) for $c = 5$ components and a weight parameter of $\omega = 0.25$. For reference purposes, we also show the real PSMC curve of the French individual. The grey vertical lines indicate the inferred event times in the C3PO model, and the colored vertical lines the inferred event times by SNIF (b) Connectivity graph of the inferred scenario. For reference, we show the inferred event times in the C3PO model as grey vertical lines.

(2019), which we will refer to as the Classical 3-Populations model—C3PO for short. The C3PO model is a tree-like model with three modern populations that exchange gene flow asymmetrically. It is based on the model of Gutenkunst et al. (2009) and has the same structure but with a higher likelihood and thus can be seen as an improved model with a better

fit to the data. The model stipulates the existence of an ancestral population that experienced an increase in size around 275 thousand years ago (Kya), and the a splitting event at about 150 Kya. This split resulted in two populations that exchanged gene flow asymmetrically: a large one that eventually became the modern African population, and a smaller one ancestral to the modern Eurasian population, whatever this terminology may mean. This ancestral lineage split about 22 Kya into the precursors of the European and Asian populations, which at this point began an exponential increase in size that continued to present day. During this period, all three lineages continued to exchange gene flow asymmetrically. The times for these resize and splitting events are represented as dotted vertical gray lines in Figure 2.14.

It is clear that the nature of this model does not lend itself to be exactly modelled by a symmetrical n -island model, but the piecewise-stationarity of our family of models should still be able to pick up the main demographic events. For example, from an n -island perspective, a merger or joining of two populations (going backwards in time) may be represented by an increase in gene flow, although this effect may be confounded by the actual changes in both the sizes of the populations and migration rates taking place during these events. Also of note is the fact that the first merger event is not visible to our method because it marks the start of the recent population expansion and is thus excluded from the distance computation.

As can be seen in panel (a) of Figure 2.14, these IICRs do not exhibit any major features past the 300 Kya mark, so they do not agree with the human PSMCs of Figure 2.11 (of which the representative ones are again shown in Figure 2.14 in dashed trace for reference), and they also do not generate significant events in the inferred demographic histories. Particularly, varying the value of the weight-shifting parameter ω did not make a great effect in this set of inferences (which is in contrast with the results shown in Figure 2.12). This inferred demographic history can be roughly summarized from panel (b) as having a period of relative high gene flow followed by a sharp decrease near the 300 Kya mark, which can be very clearly attributed to the size increase of the ancestral population in the C3PO model.

The inferred number of demes and their relative sizes for each population can be observed in Figure 2.15. The numbers for the African population is in sharp contrast with the other two populations. We can also observe that for the three populations there is more variance (compared to the results from Figure 2.13) in the inferred values of n and N across the different values of c and ω . This may indicate a weaker link to an underlying n -island model.

In general, there is little agreement between the demographic histories inferred by our method from the PSMC data and the simulated IICRs from the C3PO model. This is expected because of how the two models have fundamentally incompatible structures, not only regarding the island versus tree aspect, but also due to the size changes in the C3PO model that affect the IICR potentially as much as gene flow does. However, we do identify the approximate timings of the two visible demographic events when using $c = 5$ components and the more recent-weighted value of $\omega = 1$. These results also serve as additional validation that our method will not return the same parameter values regardless of the source of the data. They also suggest that the C3PO model is unlikely to be a good model to understand questions about ancient human structure and evolution.

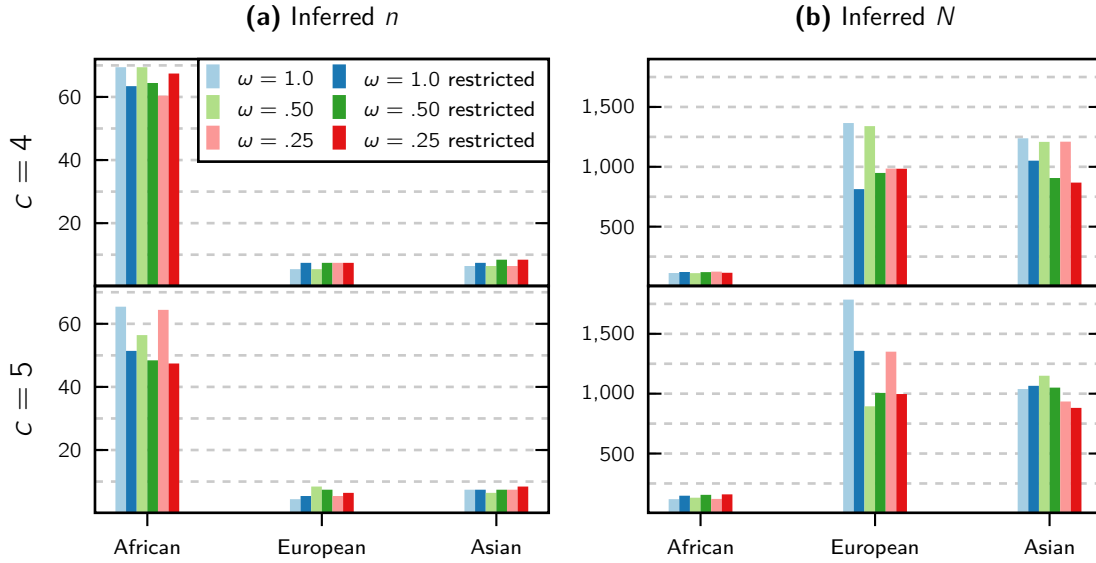


Figure 2.15: Application of our inference method to a generally-accepted tree-like human demographic scenario with three modern populations. **(a)** Inferred number of islands for each modern population. **(b)** Inferred local size of each island. Shown here are the scenarios with 4 and 5 components c , and all three values of the weight-shifting parameter ω . The bars with the darker color, marked 'restricted' in the legend, correspond to inferences realized with the option of ignoring recent population expansion.

2.4.2. Seq-Sim validations

Figure 2.17 shows the results of the validations using seq-sim IICRs. For the three chosen representative human populations (French, Karitiana and Yoruba), we selected the SNIF inferences obtained with the parameter values $c = 5$ components and $\omega = 0.2$. The selection of these values as the preferred ones was supported by the fact that the d_{visual} distance (see Figure 2.24) was best for them. For each of these inferred models we generated two independent genomic sequences of length 3×10^9 base pairs and applied the PSMC method to obtain two independently estimated target IICRs. These seq-sim IICRs are shown in blue on the left panels of Figure 2.17. The connectivity graphs associated with these inferred models are represented in the middle panels by the red curves. The corresponding inferred values of n and N are marked as the black reference circles in the right panels of the figure.

After obtaining the seq-sim IICRs, we applied again our inference method. The goal is to validate whether it would be able to consistently infer the same parameter values of the demographic history regardless of the origin of the source IICR curve. To this end we performed 10 independent inferences from each of the two target IICRs. The inferred IICRs are superimposed on the left panels of Figure 2.17 (transparent red curves, 20 per population). The inferred connectivity graphs are shown in the middle panels (transparent green curves) and the inferred values of n and N are presented in the right panels. We observe an agreement with the previously inferred histories, which suggests that *if* the real history of human evolution were close to piecewise n -island models like those used in this work, our method would be able to infer the parameters properly, and they would be similar to those

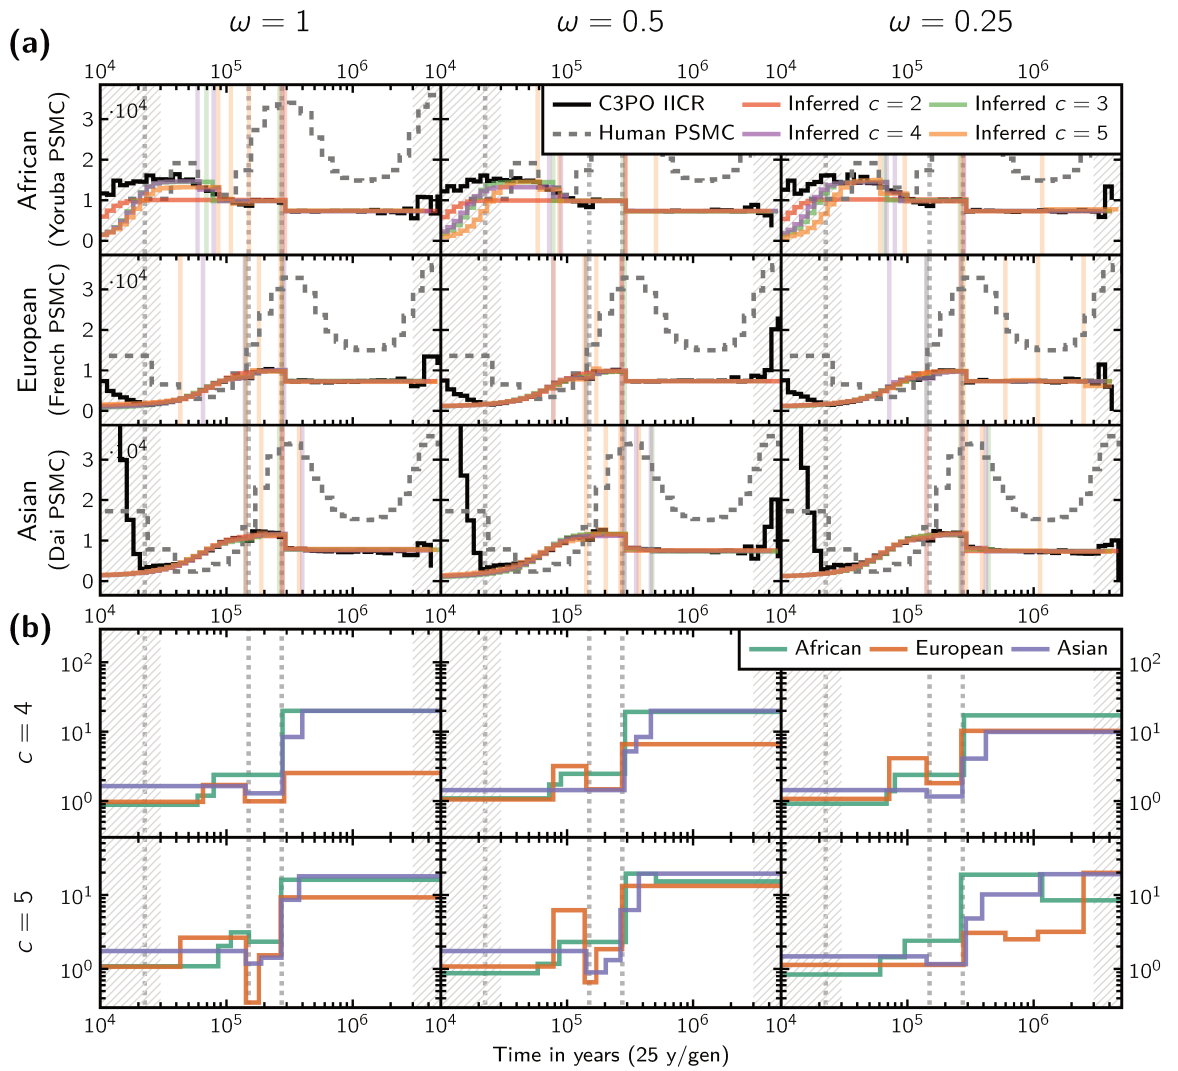


Figure 2.16: Application of our inference method to a tree-like human demographic scenario with three modern populations. **(a)** IICR plots showing the resulting IICR curve of the European population under this model and the inferred IICR curve obtained with our method (where the recent period of human expansion was ignored) for $c = 5$ components and a weight parameter of $\omega = 0.25$. For reference purposes, we also show the real PSMC curve of the French individual. The grey vertical lines indicate the inferred event times in the C3PO model, and the colored vertical lines the inferred event times by SNIF **(b)** Connectivity graph of the inferred scenario. For reference, we show the inferred event times in the C3PO model as grey vertical lines.

shown in Figures 2.12 and 2.13.

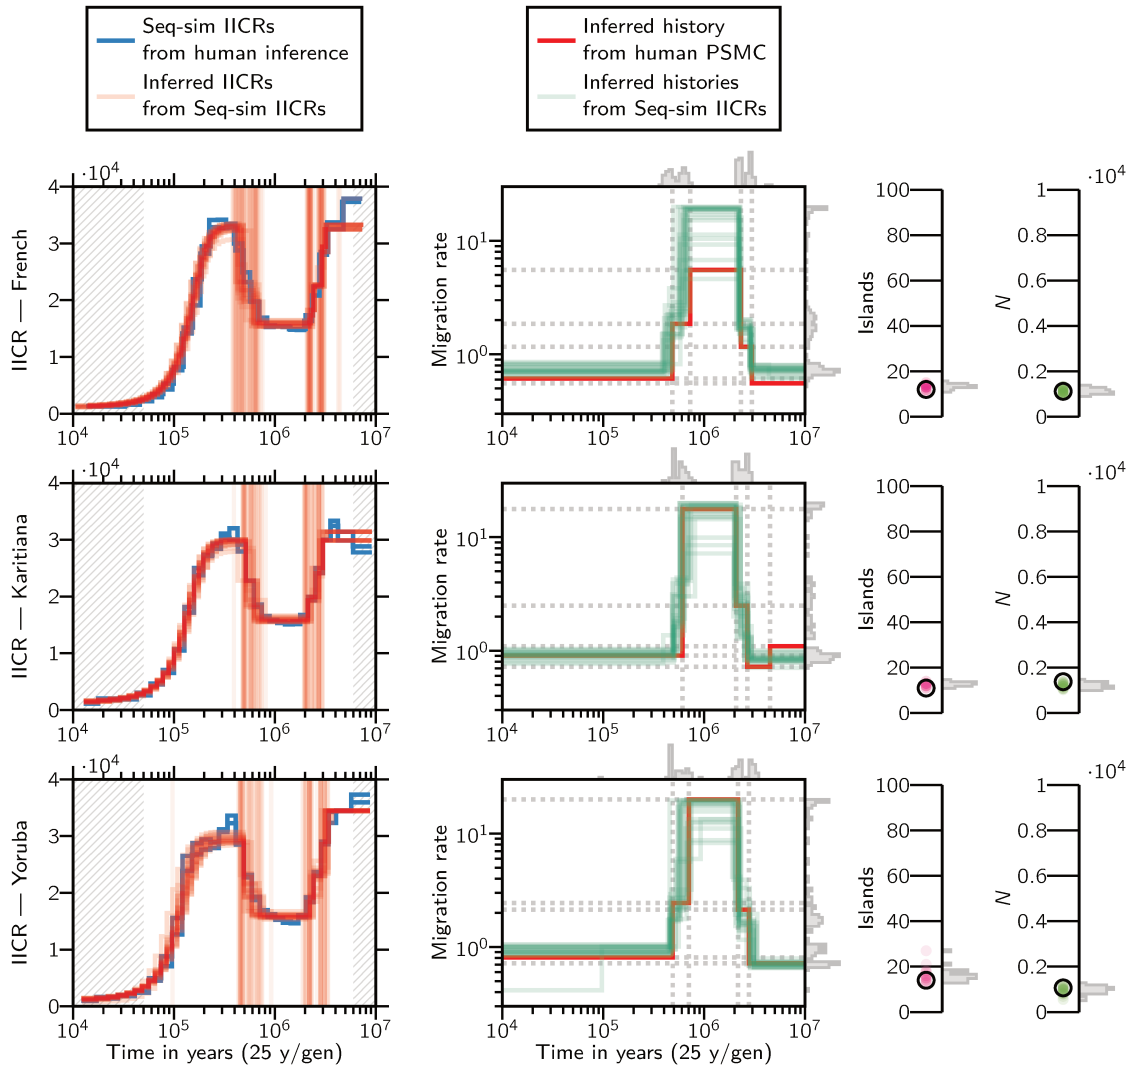


Figure 2.17: Results of the validation using Seq-sim IICRs. The left panels show two independently simulated seq-sim IICRs (obtained using the demographic scenario inferred with $c = 5$ and $\omega = 0.2$ for each of the indicated human population) alongside 10 independent inferred IICRs. The rest of the panels show the connectivity graphs, number of islands, and local deme size of these seq-sim IICRs and their corresponding inferences.




Figure 2.18: IICRs of the inferred human scenarios.




Figure 2.19: IICRs of the inferred human scenarios with restricted inference range.




Figure 2.20: Connectivity graphs of the inferred human scenarios.

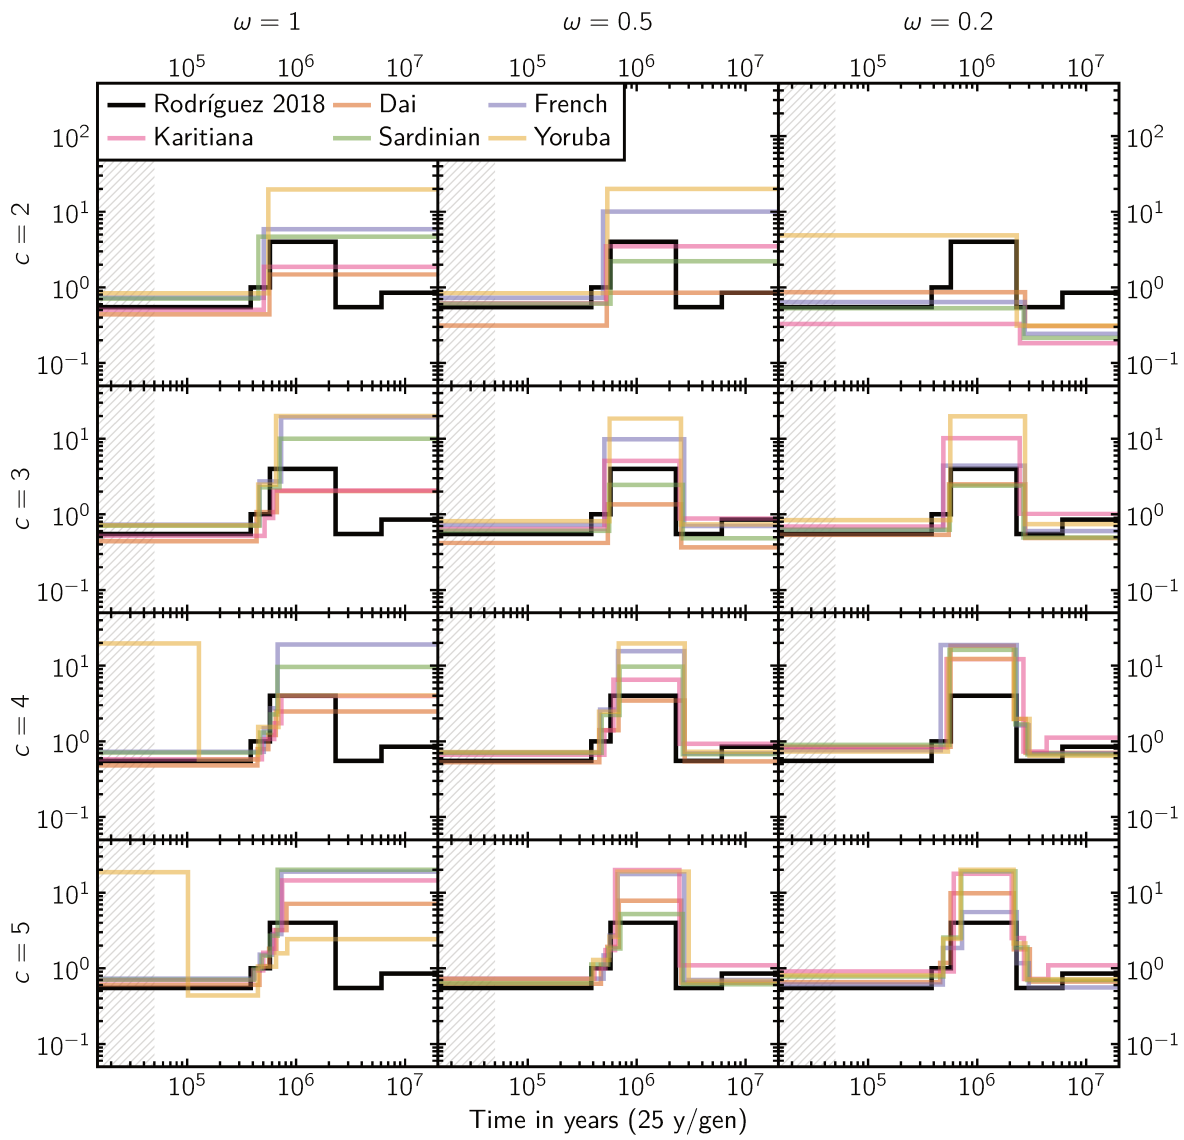


Figure 2.21: Connectivity graphs of the inferred human scenarios with restricted inference range.




Figure 2.22: Inferred number of islands for the human populations.




Figure 2.23: Inferred reference sizes for the human populations.




Figure 2.24: Visual distance (d_{visual}) of the best fitting scenario as a function of ω for various number of components.

2.4.3. SFS comparisons

In this section we aim to make a basic comparison of the Site Frequency Spectrum (SFS) corresponding to the demographic scenarios inferred by SNIF from human PSMC plots to the folded SFS of a sample of 108 humans from the Yoruba “population” (Lapierre et al., 2017), and to the SFS produced by the GADMA method (Noskova et al., 2019) for similar samples from the three populations of the C3PO model. This comparison is intentionally limited in scope because a full SFS study requires a separate study.

We simulated the SFS under two extreme models and an in-between one: one with constant size (i.e., the model inferred by SNIF from the Yoruba PSMC), another one with a recent

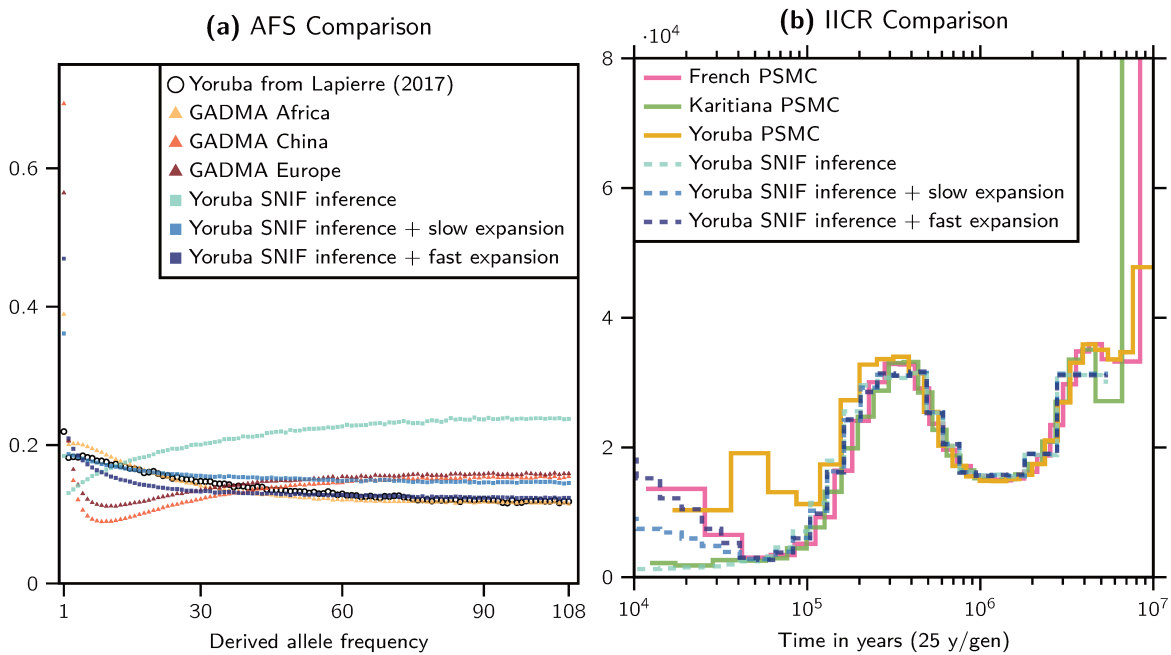


Figure 2.25: Comparison of **(a)** the folded SFSs and **(b)** the IICRs of various real, inferred and simulated human populations. The SFSs in panel **(a)** are transformed according to Lapierre et al. (2017) in order to facilitate the visual comparison. The transformation consists in multiplying the i th frequency by i before folding (the details can be found in Lapierre et al. (2017) pages 440–441).

exponential growth in the last 50,000 years (fast expansion), and another where the population growth was half as fast (slow expansion). The expansion parameters were chosen such that in the case of the fast expansion, the IICR matches in the recent past the human PSMCs that clearly exhibit a signal compatible with recent population expansion (French, for instance). The slow expansion is half as fast in the sense that in the present, the effective size is half as big. In Figure 2.25 panel (b) we plotted the corresponding three PSMCs to show how they look like and in panel (a) predicted the three corresponding SFSs which we compared with the SFSs from Lapierre et al. (2017) and Noskova et al. (2019). We find that this simple change in the recent history is enough to make a significant change of the predicted SFS. We do not try to fit any of the SFSs (Yoruba, European and Chinese). This is just a “proof of concept” simulation which suggests that existing SFSs could be easily fitted with a structured model similar to those inferred by SNIF, but in which we would allow for a recent population size change which would incorporate geography.

2.5. Implementation and use cases

Our method is implemented in a program named SNIF (Structured Non-stationary Inference Framework) which can be found at github.com/arredondos/snif. The method produces parameter estimates and connectivity graphs in order to determine whether there is consistency across individuals or species. We focused on models in which the population size was

maintained constant. The method can in principle infer changes in population size but this should be part of an extension of the present work. At this stage we stress that more work is needed before changes in connectivity and population size can be estimated together with confidence.

The intended use case for our method consists of several inference cycles, each preceded by adjustments to the many available parameters which include: the bounds B for the estimated variables, the number of components c , the weight distance parameter ω , the time interval where the distance is to be computed, the number of inference rounds along with the tolerance ε , some parameters of the search algorithm, and other minor options. This cycle emerges naturally from the fact that the inference process itself runs fairly quickly (a few seconds per round), so it is feasible to prepare a script that generates inferences under several combinations of these parameters and later do a general assessment of what are the most plausible scenarios for the data depending on the visual fit, the consistency of the inferred demographic histories, the distribution of the distances, etc. The SNIF software already includes a number of auxiliary scripts that may be used in this later analysis stage, including for instance the automatic generation of figures similar to the sub-panels of figures 2.12 and 2.13 in the main text, where the target and inferred IICRs can be compared, and the nature of the best fitting scenarios can be understood using the connectivity graphs or the n and N plots.



As for the number of components c , specifying a higher value will generally (but not always) result in a better fit and a lower distance, however it also incurs in longer analysis times and diminishing returns on the new information present in the resulting inferred histories. These aspects can be balanced by increasing the value of c up to the point where the inferred demographic histories start converging (for instance, as more components are added, the connectivity graphs stop revealing new major events and any additional new degrees of freedom are only used to refine already existing details).

As part of the user inference cycles mentioned above, the bounds system can be configured to infer parameters for more strict or specialized symmetrical island models. For instance, it is possible to fix the number of inferred islands to be exactly 3 by setting $n_{\min} = n_{\max} = 3$ in the inference bounds B . Likewise, given an independent estimate of the reference size N for the data, this information may be used to process the corresponding scaled target IICR as an unscaled one using matching inference bounds for the parameter N .

An important point concerning scaling is that the supplied value of the mutation rate μ must be accurately specified when inferring demographic parameters from a PSMC curve. This value is not used during the inference, but it is used in order to properly scale the PSMC curve and convert it into a target IICR, therefore it is important to provide the same value of μ that was used during the PSMC analysis. Otherwise, the only parameter that can be correctly inferred is n , since the rest of the parameters would only be accurate up to a scaling factor. This follows from a similar logic to that of Lemma 1.

We also suggest to use the hand-fitting Python script developed in Chikhi et al. (2018) as a complement to the automated inference process proposed in this work. You may want to compare the results obtained this way with the output of SNIF. Doing this will help make

sense of the data, or help set the bounds for the many SNIF parameters.



Chapter 3.

Multi-sample summary statistics for demographic inference

In this chapter we explore two different multi-sample summary statistics, with the goal of using them in a future demographic inference application, similar to that of Chapter 2 with the IICR summary statistic for a sample of size two haploids. The two summary statistics in question are the $IICR_k$ of an n -island model, and the expected site frequency spectrum (SFS) of a general structured model.

The $IICR_k$ of an n -island model is the most natural generalisation of the IICR. On the other hand, the SFS is a widely used statistic that succinctly summarizes genetic diversity of a demographic sample (Wakeley and Hey, 1997; Griffiths and Tavaré, 1998; Nielsen, 2000; Gutenkunst et al., 2009; Excoffier et al., 2013). For both cases, the methods we present in this chapter are centered on building and analyzing a Markov process that captures all necessary information to compute the given summary statistic of k sampled haploid lineages in a structured model.

Similar Markovian approaches have been used previously in the context of the coalescent (Hobolth et al., 2019). Notably, Herbots (1994) with the study of the structured coalescent; Rodríguez et al. (2018) with an extension to piece-wise stationarity models for studying the Inverse Instantaneous Coalescence Rate (IICR); and Kern and Hey (2017) with the exact computation of the expected Joint AFS for the two-populations isolation-with-migration (IM) demographic model.

The main challenge with the exact approach is that the Markov state space grows very quickly (almost exponentially) with the size of the sample, which poses a hard limit on applicability due to computational intractability. One of the aspects we explore in this chapter to mitigate this effect is the concept of *model specialization*, which is a way to compress the state space by exploiting symmetries in the studied models. The n -island model is best suitable for this kind of analysis due to its high degree of symmetry and mathematical simplicity.

The majority of this chapter sections are dedicated to the study of the SFS summary statistic. The SFS is a distribution $\xi = (\xi_i)_{i=1}^{k-1}$ where ξ_j indicates the *proportion* of observed segregating sites with exactly i derived alleles. The sample can be of any size (although large values of k quickly become computationally prohibitive) and the demographic structure can be specified by an arbitrary migration matrix and deme-sizes vector. In §3.2 we proceed to enumerate the state space for the continuous-time Markov process that will be studied. We

represent the states as $k \times n$ matrices, where k is the haploid sample size, n is the number of demes, and an entry z with coordinates (i, j) indicates that, in the corresponding state, there are z lineages in deme j , each ancestral to i lineages from the sample (see Figure 3.1). This state representation captures the necessary information to exactly compute the expected SFS of the sample under any arbitrary initial sampling configuration π^0 . Indeed, our main result (presented in §3.3) expresses the expected SFS as a conceptually simple function of the data structures derived from our representation.

During §3.2, we present algorithms that will collectively allow us to enumerate the state space and computing the corresponding transition rates. Special attention is placed to the ordering of the states, since this allows the rate matrix to have a special block structure (Figure 3.2) that can be exploited to achieve a computationally efficient method for obtaining the expected times spend under each state. This method is developed in §3.3.

In §3.4 we present a specialization of the demographic model to the symmetrical n -island model. Figure 3.3 shows how, by taking advantage of the symmetrical nature of this model (and updating the proposed algorithms accordingly), the state space can be significantly compressed, thus making the use of larger sample sizes computationally feasible. The figure also shows a comparison of the size of the state space in Kern and Hey (2017) for computing the joint SFS of a two-island isolation-with-migration model.

3.1. The structured coalescent and the ICR_k

In the classical structured coalescent, the evolution of the populations can be modeled with a continuous-time Markov chain in which the states are the vectors $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_n)$ indicating how many lineages there are in every island.

Initially, the sum $\sum_i \alpha_i$ is k , the total number of sampled lineages. The distribution of the α_i varies across the islands as migrations and coalescences occur. Rodríguez et al. (2018) goes into the details of how to build such models for several examples of demographic structures. In all cases, once we have enumerated the states and computed the transition matrix Q , then the exponential function $P(t) = \exp(tQ)$ contains information about the state of the Markov chain at time t . Specifically, we know that $\exp(tQ)[i, j] = P(X(t) = \alpha_j | X(0) = \alpha_i)$. These probabilities are useful for computing functions like the IICR, since the IICR can be obtained given the distribution of T_2 . A special case for the n -island model and a sample size of $k = 2$ was used in Chapter 2.

In the presence of a larger sample ($k > 2$) we may be interested in computing its associated ICR_k function. We know from (1.3.4) how, in a size changing population, the relative size function $\lambda(t)$ is related to the distribution of T_k :

$$\lambda(t) = \binom{k}{2} \frac{P(T_k > t)}{f_{T_k}(t)}. \quad (3.1.1)$$

But just like in the case for the IICR, the right hand side of (3.1.1) is not constrained to the model of panmixia with size change, or to any other demographic model for that

matter, and as such it can be computed exactly given the distribution of T_k , and can be estimated from real demographic data using methods like MSMC (Schiffels and Durbin, 2013) and MAGIC (Weissman and Hallatschek, 2017). Given the similarity of this situation to the un-identifiability problem described in Mazet et al. (2016), it is only fitting that we call this function the $IICR_k$:

$$IICR_k(t) = \frac{1 - F_{t_k}(t)}{f_{T_k}(t)}. \quad (3.1.2)$$

It is important to stress that the distribution of T_k (and therefore, the $IICR_k$) depends on the initial state of the Markov process, i.e., the sampling vector. Let call this vector π^0 , and assume that it is a probability distribution over the possible initial states of the Markov process.

We can now pose the question of what are the possible states in a structured model with n island and k lineages: “*In how many different ways can k lineages migrate and coalesce in n demes?*” The answer depends on whether we can *distinguish*, or tell apart from each other, the individual lineages and/or demes. If we represent each state as an n -tuple $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_n)$, where α_i denotes the set of lineages in island i , then we are *distinguishing* the islands by their position in the tuple. If additionally we *distinguish* the lineages by naming them $\ell_1, \ell_2, \dots, \ell_k$, then we could represent the following four distinct states for $n = 4$ and $k = 3$:

$$\begin{aligned} &(\{\ell_1, \ell_2\}, \{\}, \{\ell_3\}, \{\}), & (\{\ell_1, \ell_3\}, \{\}, \{\ell_2\}, \{\}), \\ &(\{\}, \{\ell_3\}, \{\ell_1, \ell_2\}, \{\}), & (\{\}, \{\ell_2\}, \{\ell_1, \ell_3\}, \{\}). \end{aligned} \quad (3.1.3)$$

Now, if we were to assume that the k lineages are indistinguishable, then there would be no need to use the names ℓ_j , and we could encode in each component x_i simply the number of lineages it contains. The above states reduce thus to the following:

$$\begin{aligned} &(2, 0, 1, 0), \\ &(0, 1, 2, 0). \end{aligned} \quad (3.1.4)$$

Alternatively, if we decide we can indeed distinguish the lineages but not the individual islands, then we would need to keep the ℓ_j , but also introduce a convention on how to arrange the demes. One way to do so is to consider only non-increasing tuples, i.e., encode each state as $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_n)$ where $|\alpha_i| \geq |\alpha_{i+1}|$ for all valid i . Under these assumptions, the states in (3.1.3) reduce to:

$$(\{\ell_1, \ell_2\}, \{\ell_3\}, \{\}, \{\}), \quad (\{\ell_1, \ell_3\}, \{\ell_2\}, \{\}, \{\}). \quad (3.1.5)$$

Finally, if neither islands nor lineages are distinguishable, then all of the cases considered (and more) are included in the following single, terse, state:

$$(2, 1, 0, 0).$$

So, how many states are there in each case?

Ignoring coalescence events, we proceed to count in how many ways can *exactly* k lineages be distributed in n demes. See Stanley (2011, §1.9) for more details:

	distinguishable n islands	indistinguishable n islands
distinguishable k lineages	n^k	$\left\{ \begin{matrix} k \\ 1 \end{matrix} \right\} + \left\{ \begin{matrix} k \\ 2 \end{matrix} \right\} + \cdots + \left\{ \begin{matrix} k \\ n \end{matrix} \right\}$
indistinguishable k lineages	$\left(\begin{matrix} n \\ k \end{matrix} \right)$	$p_1(k) + p_2(k) + \cdots + p_n(k)$

Here, the symbol $\left\{ \begin{matrix} k \\ i \end{matrix} \right\}$ denotes a *Stirling number of the second kind*, which counts the number of ways to partition a set of exactly k labelled objects (the lineages) into exactly i non-empty unlabeled subsets (the demes). These numbers can be computed like this:

$$\left\{ \begin{matrix} k \\ i \end{matrix} \right\} = \frac{1}{i!} \sum_{j=0}^i (-1)^{i-j} \binom{i}{j} j^k.$$

More on Stirling numbers in L Graham (1994, §6.1).

The symbol $\left(\begin{matrix} n \\ k \end{matrix} \right)$ denotes a *multiset coefficient*. It counts the number of multisets of k elements which are taken from a set of n elements (multisets are sets that allow repeated elements). It can be easily shown that:

$$\left(\begin{matrix} n \\ k \end{matrix} \right) = \binom{n+k-1}{k}.$$

More on multisets in Stanley (2011, §1.2).

Finally, $p_i(k)$ is a *restricted partition number*. It counts in how many ways can the number k be expressed as the sum of exactly i non-null numbers.

In general, an integer partition of k is a multiset of positive integers that sum to k . For example, there are 7 possible partitions of $k = 5$, which can be represented as:

$$\{5\}, \{4, 1\}, \{3, 2\}, \{3, 1, 1\}, \{2, 2, 1\}, \{2, 1, 1, 1\}, \{1, 1, 1, 1, 1\}. \quad (3.1.6)$$

The order of the terms in the multiset is unspecified, so $5 = 2 + 3$ is the same partition as $5 = 3 + 2$. An alternative way of representing these partitions is by using *product notation*,

where we specify the number of times each part is repeated, from 1 to k which is the largest possible part. The partitions in (3.1.6) would be re-written as:

$$(00001), (10010), (01100), (20100), (12000), (31000), (50000). \quad (3.1.7)$$

We denote by $P(k)$, or simply P for simplicity, the total number of integer partitions of k . There is no known closed-form expression for P , but Hardy and Ramanujan (1918) provided the asymptotic approximation

$$P \sim \frac{e^{\pi\sqrt{2k/3}}}{4k\sqrt{3}}.$$

Algorithm P generates all partitions of k in standard notation (3.1.6).

Algorithm P: For evaluating routine **Partitions**(k), which visits, in reverse-lexicographic order, all integer partitions $(a_1, a_2, \dots, a_\omega)$ such that $\sum_{i=1}^{\omega} a_i = k$.

```

1  $\omega \leftarrow 1; a \leftarrow (0)_{1 \times k}$ 
2 while True:
3    $a_\omega \leftarrow k; q \leftarrow \omega - \text{id}_{k-1}$ 
4   while True:
5     visit  $(a_1, a_2, \dots, a_\omega)$ 
6     if  $a_q \neq 2$ :
7       break
8      $a_q \leftarrow 1; q \leftarrow q - 1; \omega \leftarrow \omega + 1; a_\omega \leftarrow 1$ 
9   if  $q = 0$ :
10    exit
11   $x \leftarrow a_q - 1; a_q \leftarrow x; k \leftarrow \omega - q - 1; \omega \leftarrow q + 1$ 
12  while  $k > x$ :
13     $a_q \leftarrow x; \omega \leftarrow \omega + 1; k \leftarrow k - x$ 

```

See Knuth (2005, §7.2.1.4) for the derivation and detailed analysis of this algorithm.

Integer partitions can be restricted in a multitude of ways. In this chapter, we are interested in restricting the number of parts. For instance, $5 = 2 + 3$ is a partition in two parts, and $5 = 2 + 1 + 1 + 1$ is a partition in four parts. The symbols $p_j(k)$ thus count the possible number of restricted partitions of k into j parts.

We may choose from these distinguishable/indistinguishable configurations depending on our needs. For instance, if the sizes of the islands are different in the general case, then the islands are inherently distinguishable. It is important to keep in mind that the choice of configuration will affect the size of the state space and the transition rates of the system. In any case, we denote the total number of states in the Markov chain by E .

Returning to the IICR_k , in order to compute this function we must be able to evaluate the state of the continuous-time Markov system until the moment of first coalescence.

The migration events after the first coalescence need not be taken into account, nor any coalescences after the first, since these events do not contribute to the distribution of T_k . For that reason, any coalescence event that happens in the initial states leads to the only absorbing state of the system. This produces a rate matrix Q with the structure:

$$Q = \begin{pmatrix} Q^* & b \\ \mathbb{0} & \mathbb{0} \end{pmatrix}_{E \times E},$$

where Q^* is of size $(E - 1) \times (E - 1)$, and b is of size $(E - 1) \times 1$. We are interested in computing $P(t) = \exp(tQ)$. The coefficient $P(t)[i, E - 1]$ contains the probability that, given an initial sampling in state i , the system has reached the absorbing state by time t , i.e., $P(T_k \geq t)$.

There are many ways to compute matrix exponentials. See for instance the extensive review in Moler and Van Loan (2003). An important problem of these methods, however, is that they are based on matrix factorizations, which have a computational complexity of $O(E^3)$, with E being the size of the matrix. Additionally, factorization methods generally produce dense factors even when the matrix Q is sparse, which also poses a large challenge in memory requirements. As a result, these methods scale poorly with the size of the problem (number of sampled lineages or number of islands in the model).

An alternative approach can be had with the computation of only the action of the exponential matrix. Indeed, we note that for our application, we are only interested in the last column of $\exp(tQ)$. We can express this by saying that we want to compute the vector $(\pi^0)^\top \exp(tQ)$ without having to compute $\exp(tQ)$ first. There are a number of methods that have focused on this in the literature (see for instance Al-Mohy and Higham (2011)), and with an asymptotic time complexity of $O(E^2)$ they provide a great improvement over the factorization methods mentioned before. However, they still have some drawbacks: these methods are based on the truncated computation of the Taylor expansion of $\exp(tQ)$, and a significant part of their computational cost is spent towards finding the ideal cutoff point that minimizes the numerical error. Moreover, they do not take advantage of the sparse nature of the matrix Q .

Recent developments by Sherlock (2021) provide an improvement over these methods. Sherlock proposes an algorithm that computes the action of the exponential while taking into account the properties of Q by virtue of it being specifically the rate matrix of a Markov process. This algorithm achieves a computational cost of $O(E\rho)$, where ρ is the maximum diagonal element of Q in absolute value.

An implementation of Sherlock's algorithm with an application to demographic inference using the IICR_k will be the subject of a future study.

3.2. The structured coalescent with ancestry tracking

We proceed to describe the state space for the continuous-time Markov process that models the migrations and ancestry of a sampled population. We do this by introducing a generalization of the classical structured coalescent process.

The state space of the structured coalescent keeps track of what lineages are in which demes. For instance, an initial state of $\alpha = (0, 2, 1, 0)$ encodes that the model has $n = 4$ demes and $k = 3$ sampled lineages, with two of them located in deme 2 and the other in deme 3. We note that the islands are distinguishable by their index in the tuple, whereas the individual lineages are not (it is not possible to differentiate between the two lineages in deme 2 for instance).

This information is sufficient in order to compute the expected coalescence times $\mathbb{E}(T_k), \mathbb{E}(T_{k-1}), \dots, \mathbb{E}(T_2)$ associated to the Markov process for any structured model. However, it is not enough information to compute a site frequency spectrum, since we would also need information about which lineages are ancestral to which. More precisely, in order to construct an SFS, we need to know at every point in time how many descendants each lineage has in order to tell, given a mutation in an ancient lineage, how many copies it will have in the current sample. We achieve this by representing the states as matrices instead of tuples. We call these matrices *state matrices*.

3.2.1. State matrices

A state matrix has k rows and n columns, and the coefficient in row i and column j indicates how many lineages there are in deme j that are ancestral to i lineages from the sample. The example in Figure 3.1 shows an one possible demographic history for $k = 4$ sampled lineages and $n = 3$ islands, with the corresponding state matrices.

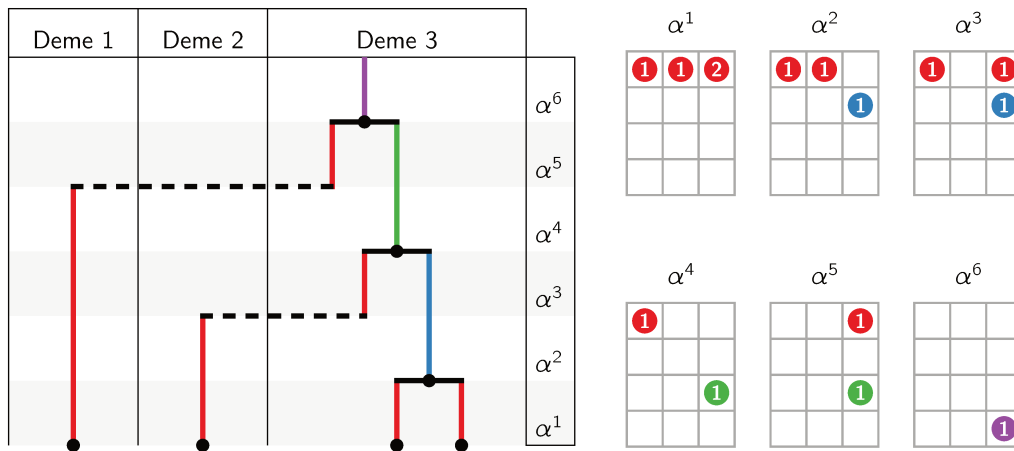


Figure 3.1: Example of a demographic history and the corresponding state matrices. In this case there are $n = 3$ demes in the model, and an initial sample of $k = 4$ lineages. Looking at state α^4 for example, there is one lineage in deme 1 (column $j = 1$), which is ancestral to one lineage (row $i = 1$) from the sample. There is also one lineage in deme 3 (column $j = 3$), which is ancestral to three lineages (row $i = 3$) from the sample. Thus, any mutation to this (green) lineage in α^4 will be present in three of the four sampled lineages, and will thus contribute to coefficient ξ_3 of the SFS. The non-zero coefficients of the state matrices and their corresponding lineages in the coalescence tree have been color-coded in this figure according to the number of ancestral lineages or *weight* for ease of reference. The remaining null coefficients of the matrix are omitted in order to avoid visual clutter. The migration rates, deme sizes and coalescence times of the model are not represented.

Given our description of state matrix for a model with n demes and k samples, there is a relationship between the number of sampled lineages and the coefficients α_{ij} of any state matrix α , which is given by the invariant:

$$\sum_{\substack{1 \leq i \leq k \\ 1 \leq j \leq n}} i \alpha_{ij} = k, \quad \forall \alpha. \quad (3.2.1)$$

We refer to the factor i in (3.2.1) as *the weight* of the lineages that are in the row i . It represents the fact that each lineage in row i is ancestral to i lineages from the sample. The number of lineages that are present in any given state α is denoted $|\alpha|$, so:

$$\sum_{\substack{1 \leq i \leq k \\ 1 \leq j \leq n}} \alpha_{ij} = |\alpha|,$$

and we call r_α the number of coalescences that have taken place up to state α , such that

$$k = |\alpha| + r_\alpha, \quad \forall \alpha.$$

We note that the state matrix associated with any state α is empty for rows below $r_\alpha + 1$. Indeed, if we suppose $\alpha_{ij} \geq 1$ for some $i \geq r_\alpha + 2$, then the weight of that lineage is at least $r_\alpha + 2$, and since the remaining $|\alpha| - 1$ lineages have weight at least 1, then the total weight of the state is:

$$k = \sum_{\substack{1 \leq i \leq k \\ 1 \leq j \leq n}} i \alpha_{ij} \geq (r_\alpha + 2) + (|\alpha| - 1) = k + 1,$$

which is absurd.

We denote by \mathcal{E} the set of all possible states α , and by E the total number of states:

$$\mathcal{E} = \left\{ \alpha \in \mathbb{Z}_+^{k \times n} : \sum_{i=1}^k \sum_{j=1}^n i \alpha_{ij} = k \right\}, \quad |\mathcal{E}| = E.$$

The set \mathcal{E} is clearly defined by the sample size k and number of demes n , but in order to avoid notation clutter we keep this dependency implicit.

Another important notion related to the state matrices are the contribution vectors. The contribution vector of state α , denoted ψ_α , is defined as the sum by rows of the states matrices:

$$\psi_\alpha = \left(\sum_{1 \leq j \leq n} \alpha_{1,j}, \sum_{1 \leq j \leq n} \alpha_{2,j}, \dots, \sum_{1 \leq j \leq n} \alpha_{k,j} \right).$$

Returning to the example from Figure 3.1, the contribution vectors are:

$$\begin{aligned} \psi_{\alpha^1} &= (4, 0, 0, 0); & \psi_{\alpha^2} &= \psi_{\alpha^3} = (2, 1, 0, 0); \\ \psi_{\alpha^4} &= \psi_{\alpha^5} = (1, 0, 1, 0); & \psi_{\alpha^6} &= (0, 0, 0, 1). \end{aligned}$$

In general, for $i < k$, the coefficient $\psi_\alpha(i)$ represents the contribution to $\mathbb{E}(\xi_i)$ for each unit of time the process spends in state α on average (see §3.3). We denote by \mathcal{P} the set of all possible contribution vectors, and by P its total number of elements:

$$\mathcal{P} = \left\{ \psi \in \mathbb{Z}_+^k : \sum_{i=1}^k i \psi(i) = k \right\}, \quad |\mathcal{P}| = P. \quad (3.2.2)$$

In this case, \mathcal{P} depends only on the number of samples k (in fact $\mathcal{P} = \mathcal{E}$ when $n = 1$), but once again we keep this dependency implicit.

3.2.2. Transitions between states

From this point onwards we use the notation e_{ij} to indicate a matrix of size $k \times n$ with null coefficients except for the one in row i and column j , which has a value of 1.

There are two possible types of state transitions: migration events and coalescence events. During a migration event, one lineage in deme j migrates to deme $g \neq j$ with rate m_{jg} . We collect these rates in the matrix

$$\mathcal{M} = (m_{jg}) \text{ with } 1 \leq j, g \leq n; \quad j \neq g.$$

We convene, for the sake of completeness, that the diagonal elements of \mathcal{M} are null.

Given two states α and β , and a weight of $i \in \{1, \dots, k-1\}$, the process can transition from α to β via migration event if for some $j, g \in \{1, \dots, n\}$

$$\beta = \alpha - e_{ij} + e_{ig}, \quad (3.2.3)$$

and this transition happens with rate m_{jg} for every lineage that could potentially migrate, of which there are α_{ij} . Algorithm M generates all possible destination states β given any departure state α .

Algorithm M: For evaluating routine **MigrationDestinations**(\mathcal{M}, α), which given the migration rates \mathcal{M} and a state matrix α of size $k \times n$, visits all the tuples (β, m) where β is a state satisfying (3.2.3) and m (which may be null) is the corresponding transition rate.

```

1 for each  $i$  in  $\{1, 2, \dots, k-1\}$ :
2   for each  $j$  in  $\{1, 2, \dots, n\}$ :
3     if  $\alpha_{ij} \geq 1$ :
4       for each  $g$  in  $\{1, 2, \dots, j-1, j+1, \dots, n\}$ :
5          $\beta \leftarrow \alpha - e_{ij} + e_{ig}$ 
6          $m \leftarrow \alpha_{ij} m_{jg}$ 
7         visit  $(\beta, m)$ 

```

We note that Algorithm M visits all *potential* migration destinations, regardless of whether a transition would be possible given the migration rates \mathcal{M} . This behaviour is desired because, in addition to generating the migration rates, we use the algorithm to build the state space \mathcal{E} by repeatedly simulating all possible migrations (see §3.2.4 and Algorithm A).

A coalescence happens when two lineages in the same deme j are replaced with their immediate common ancestor, also in deme j . This coalescence happens with a rate that is the inverse of the deme size. We collect the deme sizes in the vector

$$\mathcal{S} = (s_j)_{j=1}^n.$$

Given two states α and β , and two values $i, h \in \{1, \dots, k-1\}$, the process can transition from α to β via coalescence event if for some $j \in \{1, \dots, n\}$

$$\beta = \alpha - e_{ij} - e_{hj} + e_{i+h,j}, \quad (3.2.4)$$

and this transition happens with rate $1/s_j$ for every pair of lineages that can coalesce, of which there are $\binom{\alpha_{ij}}{2}$ if $h = i$, and $\alpha_{ij} \alpha_{hj}$ otherwise. Algorithm C generates all possible states to which the process can transition by a coalescence from α .

Algorithm C: For evaluating routine **CoalescenceDestinations**(\mathcal{S}, α), which given the deme sizes and a state matrix α , visits all the tuples (β, c) where β is a state satisfying (3.2.4) and $c \neq 0$ is the corresponding coalescence rate.

```

1 for each  $j$  in  $\{1, 2, \dots, n\}$ :
2   for each  $i$  in  $\{1, 2, \dots, k-1\}$ :
3     if  $\alpha_{ij} \geq 1$ :
4       if  $\alpha_{ij} \geq 2$ :
5          $\beta \leftarrow \alpha - 2e_{ij} + e_{2i,j}$ 
6          $c \leftarrow \alpha_{ij}(\alpha_{ij} - 1)/(2s_i)$ 
7         visit  $(\beta, c)$ 
8       for each  $h$  in  $\{i+1, \dots, k-1\}$ :
9         if  $\alpha_{hj} \geq 1$ :
10         $\beta \leftarrow \alpha - e_{ij} - e_{hj} + e_{i+h,j}$ 
11         $c \leftarrow \alpha_{ij} \alpha_{hj}/s_j$ 
12        visit  $(\beta, c)$ 

```

3.2.3. Grouping and sorting the states

We now focus on defining the rate matrix of the Markov process. This is a matrix Q of size $E \times E$, where coefficient (ℓ_0, ℓ_1) contains the rate of transition from state α^{ℓ_0} to state α^{ℓ_1} when $\ell_0 \neq \ell_1$; and the diagonal elements are defined such that $\sum_{\ell=1}^E Q(\ell_0, \ell) = 0$. To define

such a matrix, we first need to impart an order \prec on the set of all states \mathcal{E} . The focus of this section will therefore be the construction of a bijective map

$$\begin{aligned} \mathcal{L} : \mathcal{E} &\rightarrow \{1, \dots, E\} \\ \alpha^\ell &\mapsto \ell \end{aligned} \quad (3.2.5)$$

such that $1 \leq \ell_0 < \ell_1 \leq E \Leftrightarrow \alpha^{\ell_0} \prec \alpha^{\ell_1}$.

We start with the partial order given by the natural transition of the process from states with more living lineages to states with fewer living lineages. Indeed, migration events do not change the number of living lineages, and a coalescence events decrease them by one, therefore $|\alpha|$ (or equivalently, r_α) decreases (increases) during any given instantiation of the process. We convene then that:

$$r_\alpha < r_\beta \Rightarrow \alpha \prec \beta. \quad (3.2.6)$$

In the example from Figure 3.1, since $r_{\alpha^1} = 0$ and $r_{\alpha^3} = 1$, we have $\alpha^1 \prec \alpha^3$.

By grouping together the states that share the same number of live lineages, we get a partition of the state space into k disjoint parts (one part for each possible value of r_α from 0 to $k - 1$). We denote these parts by \mathcal{E}_i , each with E_i elements, such that $\mathcal{E} = \cup_i \mathcal{E}_i$ and $E = \sum_i E_i$:

$$\mathcal{E}_i = \{\alpha \in \mathcal{E} : r_\alpha = i - 1\}, \quad |\mathcal{E}_i| = E_i, \quad 1 \leq i \leq k.$$

The partition $\{\mathcal{E}_i\}$ can be understood as the quotient set \mathcal{E}/R , where R is the equivalence relation

$$\alpha R \beta \Leftrightarrow r_\alpha = r_\beta \Leftrightarrow |\alpha| = |\beta|; \quad \alpha, \beta \in \mathcal{E}. \quad (3.2.7)$$

This partial order imposes a $k \times k$ block structure on Q , where block (i, h) has dimensions $E_i \times E_h$, and contains the rates of transitioning from the states in \mathcal{E}_i to the states in \mathcal{E}_h .

$$Q = \begin{pmatrix} [Q_{11}]_{E_1 \times E_1} & [Q_{12}]_{E_1 \times E_2} & \cdots & [Q_{1k}]_{E_1 \times E_k} \\ [Q_{21}]_{E_2 \times E_1} & [Q_{22}]_{E_2 \times E_2} & \cdots & [Q_{2k}]_{E_2 \times E_k} \\ \vdots & \vdots & \ddots & \vdots \\ [Q_{k1}]_{E_k \times E_1} & [Q_{k2}]_{E_k \times E_2} & \cdots & [Q_{kk}]_{E_k \times E_k} \end{pmatrix}_{E \times E}. \quad (3.2.8)$$

We note that if the process is at a non-absorbing state $\alpha \in \mathcal{E}_i$, then it can only transition to another state in \mathcal{E}_i via migration, or to a state in \mathcal{E}_{i+1} via coalescence, therefore all the coefficients of block $[Q_{ih}]$ are null when $h \neq i$ and $h \neq i + 1$. Matrix Q is thus block bi-diagonal, with blocks $A_i = [Q_{ii}]$ for $1 \leq i \leq k$ in the diagonal, and blocks $B_i = [Q_{i, i+1}]$ for $1 \leq i \leq k - 1$ in the upper diagonal.

$$Q = \begin{pmatrix} A_1 & B_1 & & & \\ & A_2 & B_2 & & \\ & & \ddots & \ddots & \\ & & & A_{k-1} & B_{k-1} \\ & & & & A_k \end{pmatrix}_{E \times E}.$$

We call A_i the migration blocks and B_i the coalescence blocks of Q . We further note that migration block A_k is null, since there are no more transitions once the process arrives at any of the stationary states of \mathcal{E}_k . An alternative representation of Q is thus possible:

$$Q = \begin{pmatrix} Q^* & B \\ \mathbb{0} & \mathbb{0} \end{pmatrix}, \quad (3.2.9)$$

where

$$Q^* = \begin{pmatrix} A_1 & B_1 & & & & \\ & A_2 & B_2 & & & \\ & & \ddots & \ddots & & \\ & & & A_{k-2} & B_{k-2} & \\ & & & & A_{k-1} & \\ & & & & & \end{pmatrix}_{E^* \times E^*} \quad \text{and} \quad B = \begin{pmatrix} \mathbb{0} \\ \mathbb{0} \\ \vdots \\ \mathbb{0} \\ B_{k-1} \end{pmatrix}_{E^* \times E_k}.$$

Here, Q^* is known as the sub-intensity matrix and contains the rates of transition between the transient states $\mathcal{E}^* = \mathcal{E} - \mathcal{E}_k$. E^* is the number of transient states $E - E_k$.

Next we focus on defining the order \prec for the states within the sets \mathcal{E}_i . We begin by defining another equivalence relation C between states of \mathcal{E} . We say that two states are equivalent under C if they have the same contribution vector:

$$\alpha C \beta \Leftrightarrow \psi_\alpha = \psi_\beta; \quad \alpha, \beta \in \mathcal{E}. \quad (3.2.10)$$

This relation partitions the state space into disjoint sets identified by their common contribution vector. We can thus think of each different contribution vector as the representative of an equivalence class of states, or a *macrostate*. In order to count how many different macrostates there are in a process with sample size k , we need to briefly return to integer partitions. It is clear from their product representation (3.1.7) that the set of all integer partitions of k is the set \mathcal{P} of all possible state contribution vectors defined in (3.2.2). There are thus P possible macrostates in the process. The number of parts in macrostate $\psi_\alpha \in \mathcal{P}$ is given by $\sum_{i=1}^k \psi_\alpha(i) = |\alpha|$. We group the macrostates by their number of parts following a similar procedure to how we group the states by their number of live lineages in (3.2.7). Overloading the name R to define a relation on \mathcal{P} , we say that:

$$\psi_\alpha R \psi_\beta \Leftrightarrow |\alpha| = |\beta| \Leftrightarrow r_\alpha = r_\beta; \quad \psi_\alpha, \psi_\beta \in \mathcal{P}. \quad (3.2.11)$$

This partitions \mathcal{P} into k disjoint sets \mathcal{P}_i , each with P_i macrostates, such that $\mathcal{P} = \cup_i \mathcal{P}_i$ and $P = \sum_i P_i$:

$$\mathcal{P}_i = \{\psi_\alpha \in \mathcal{P} : r_\alpha = i - 1\}, \quad |\mathcal{P}_i| = P_i, \quad 1 \leq i \leq k.$$

We sort the macrostates of \mathcal{P}_i in reverse lexicographic order of standard notation (3.1.6) as does Algorithm P, so for $k = 5$ we have $(4, 1) \prec (3, 2)$ in \mathcal{P}_2 , and $(3, 1, 1) \prec (2, 2, 1)$ in \mathcal{P}_3 . This is a total order in \mathcal{P}_i , so for each $i \in \{1, \dots, k\}$ we will assume the existence of a

bijjective map \mathcal{L}_i defined on \mathcal{P}_i such that $\mathcal{L}_i(\psi) \in \{1, \dots, P_i\}$, and for $\psi, \varphi \in \mathcal{P}_i$, we have $\psi \prec \varphi \Leftrightarrow \mathcal{L}_i(\psi) < \mathcal{L}_i(\varphi)$.

Since $\psi_\alpha \in \mathcal{P}_i \Leftrightarrow r_\alpha = i - 1 \Leftrightarrow \alpha \in \mathcal{E}_i$, this order in \mathcal{P}_i motivates a natural way to introduce a partial order in the elements of \mathcal{E}_i . We convene that for $\alpha, \beta \in \mathcal{E}_i$,

$$\psi_\alpha \prec \psi_\beta \Rightarrow \alpha \prec \beta. \quad (3.2.12)$$

We can group together the uncomparable states of \mathcal{E}_i by returning to the equivalence relation C defined in (3.2.10) such that $\alpha C \beta \Leftrightarrow \psi_\alpha = \psi_\beta$ for $\alpha, \beta \in \mathcal{E}_i$. There are P_i equivalence classes in the quotient set \mathcal{E}_i/C , which we will denote \mathcal{E}_{ip} , each with E_{ip} states. We now have $\mathcal{E}_i = \cup_{p=1}^{P_i} \mathcal{E}_{ip}$, and $E_i = \sum_{p=1}^{P_i} E_{ip}$.

$$\mathcal{E}_{ip} = \{\alpha \in \mathcal{E}_i : \mathcal{L}_i(\psi_\alpha) = p\}, \quad |\mathcal{E}_{ip}| = E_{ip}, \quad 1 \leq p \leq P_i.$$

As was the case with Q in (3.2.8), this partial order on \mathcal{E}_i imposes a $P_i \times P_i$ block structure on migration block A_i , and a $P_i \times P_{i+1}$ block structure on coalescence block B_i . For $p, q \leq P_i$, sub-block (p, q) of A_i has dimensions $E_{ip} \times E_{iq}$, and for $q' \leq P_{i+1}$, sub-block (p, q') of B_i has dimensions $E_{ip} \times E_{i+1, q'}$.

Upon closer inspection, we see that the contribution vectors of the states in \mathcal{E}_i are invariant under migration. Indeed, if $\beta = \alpha - e_{ij} + e_{iq}$, then the sum of row i of both states remains unchanged, and thus $\psi_\alpha = \psi_\beta$. This implies that sub-block (p, q) of migration block A_i has all null entries except when $p = q$, and thus A_i is $P_i \times P_i$ block-diagonal. We denote the diagonal sub-blocks of A_i as $A_{ip} \in \mathbb{R}^{E_{ip} \times E_{ip}}$ for $1 \leq p \leq P_i$. There is no such simplification for coalescence block B_i , but we group all its sub-blocks in each block-column as $B_{iq} \in \mathbb{R}^{E_i \times E_{i+1, q}}$ with $1 \leq q \leq P_{i+1}$ for later convenience:

$$A_i = \begin{pmatrix} A_{i1} & & & \\ & A_{i2} & & \\ & & \dots & \\ & & & A_{iP_i} \end{pmatrix}_{E_i \times E_i}; \quad B_i = \begin{pmatrix} B_{i1} & B_{i2} & \dots & B_{iP_{i+1}} \end{pmatrix}_{E_i \times E_{i+1}}. \quad (3.2.13)$$

To make \prec a total order, we need to define it for the states within each of the P sub-spaces \mathcal{E}_{ip} . There is no further benefit in doing so in terms of the structure of the rate matrix Q in the general case, but it is still a necessary step for building it. We thus convene that the order of the states in \mathcal{E}_{ip} is the one obtained by applying algorithm M to state $\alpha = (\mathcal{L}_i^{-1}(p)^\top, 0_{k \times 1}, \dots, 0_{k \times 1})$ and all its subsequent migration destinations repeatedly (this procedure is detailed in algorithm A). Therefore, for every $i = 1, \dots, k$ and $p = 1, \dots, P_i$ we assume the existence of a bijjective map \mathcal{L}_{ip} defined on \mathcal{E}_{ip} such that $\mathcal{L}_{ip}(\alpha) \in \{1, \dots, E_{ip}\}$ and $\alpha \prec \beta \Leftrightarrow \mathcal{L}_{ip}(\alpha) < \mathcal{L}_{ip}(\beta)$. With this construction, we can finally map every state $\alpha \in \mathcal{E}$ to an index $\ell \in \{1, \dots, E\}$, and the definition of the global map \mathcal{L} (3.2.5) would be

$$\mathcal{L}(\alpha) = \sigma_{ip} + \mathcal{L}_{ip}(\alpha), \quad \forall \alpha \in \mathcal{E} \quad (3.2.14)$$

where:

$$\begin{aligned}\sigma_{ip} &= \sigma_i + E_{i1} + E_{i2} + \dots + E_{ip-1}; \\ \sigma_i &= E_1 + E_2 + \dots + E_{i-1}; \\ p &= \mathcal{L}_i(\psi_\alpha); \\ i &= r_\alpha + 1.\end{aligned}$$

Figure 3.2 presents a schematic summary of some of the objects introduced in this section, their relations and the associated notation.




Figure 3.2: Schematic representation of the macro- and micro-structure of the state space \mathcal{E} and the rate matrix Q that results from appropriately grouping and sorting the states.

3.2.4. Construction of the rate matrix

We now propose a method for the construction of the migrations sub-blocks A_{ip} , coalescence sub-blocks B_{ip} and indexing sub-maps \mathcal{L}_{ip} by combining the algorithms of the previous section. The method is structured in two separate stages.

Discovery and migration stage. The goal of this stage is to generate all the states in a constructive manner by following all possible migration paths, collecting at the same time the information of the transition rates in the migration sub-blocks A_{ip} .

We know from the previous section that the state space is partitioned by the relation C from (3.2.10) into P sub-sub-spaces denoted \mathcal{E}_{ip} , with no possible migration between them. We thus proceed to *discover* the states by initializing the sets \mathcal{E}_{ip} with the class representatives of \mathcal{E}/C . These are the states that have all their live lineages in the first deme, so they can

always be defined regardless of the value of n by placing their associated contribution vector in the first column:

$$\psi \mapsto \alpha = (\psi^\top \ 0 \ 0 \ \cdots \ 0)_{k \times n}. \quad (3.2.15)$$

The contribution vectors ψ are obtained by converting the partitions (a_1, \dots, a_ω) generated by algorithm P to product notation (3.1.7). See lines [3...6] in Algorithm Q.

The next step is to generate (and add to \mathcal{E}_{ip}) all the states that can be reached via migration from the initial state (3.2.15) using Algorithm M, then all the states that can be reached from those new states, and so on until no new states are being visited or added. This operation is accomplished using a *queue* data structure, or any generic data structure that implements the FIFO (first in, first out) interface, where elements are added at one end of the queue, and removed from the other end.

For our use, we may represent a queue of states as a tuple of length $0 \leq \omega < \infty$ with elements in \mathcal{E} , and two associated functions: enq and deq, which add (enqueue) and remove (dequeue) elements from the queue, respectively. More precisely, we say that $U \in \mathcal{E}^\omega$ is a queue of size ω ; and for $\mathcal{U} = \cup_{\omega \geq 0} \mathcal{E}^\omega$ we define the maps

$$\begin{aligned} \text{enq} : \mathcal{U} \times \mathcal{E} &\rightarrow \mathcal{U}, \\ ((\alpha^1, \alpha^2, \dots, \alpha^\omega), \alpha) &\mapsto (\alpha^1, \alpha^2, \dots, \alpha^\omega, \alpha); \end{aligned}$$

$$\begin{aligned} \text{deq} : \mathcal{U} \setminus \mathcal{E}^0 &\rightarrow \mathcal{E} \times \mathcal{U}, \\ (\alpha^1, \alpha^2, \dots, \alpha^\omega) &\mapsto (\alpha^1, (\alpha^2, \dots, \alpha^\omega)). \end{aligned}$$

For a formal treatment of the queue abstract data structure (and other common data structures in computer science), see Lehmann and Smyth (1981).

Each time a state β is found to be reachable via migration from state α , we accumulate the corresponding migration rate m in coefficient $A_{ip}(\mathcal{L}_{ip}(\alpha), \mathcal{L}_{ip}(\beta))$ of the migration sub-block. The index mapping functions \mathcal{L}_{ip} are also progressively constructed. We represent them as subsets of the cartesian product $\mathcal{E}_{ip} \times \{1, \dots, E_{ip}\}$, which contains tuple (α, ℓ) if and only if $\mathcal{L}_{ip}(\alpha) = \ell$. Building these data structures is sufficient to fully define the state space, since the state spaces \mathcal{E}_{ip} are simply the domains of the maps, which we reference using the notation $\mathcal{E}_{ip} = D_{\mathcal{L}_{ip}}$. Algorithm A summarizes the proposed procedure for this stage.

⚙️ Implementation Notes

An efficient implementation of Algorithm A should not repeatedly reallocate and rebuild matrix A_{ip} as indicated in step [11]. Instead, the sparsity of these matrices should be exploited by using a space-efficient representation like the DOK (dictionary of keys) format, for instance. In such cases, step [11] simply adds the key-value pair $((\ell_0, \ell_1), m)$ to the dictionary.

Coalescence stage. The goal for this stage is to compute the coalescence sub-blocks B_{ip} , and consequently update the diagonal entries of sub-blocks A_{ip} .

Algorithm A: For evaluating routine **MigrationSubBlock**(\mathcal{M}, α^0), which given the pairwise migration rates \mathcal{M} of the island model and a protostate α^0 (3.2.15) in \mathcal{E}_{ip} , builds and returns $(A_{ip}, \mathcal{L}_{ip})$.

```

1  $A_{ip} \leftarrow (\mathbb{0})_{1 \times 1}$ ;  $\mathcal{L}_{ip} \leftarrow \{(\alpha^0, 1)\}$ ;  $U \leftarrow (\alpha^0)_{1 \times 1}$ 
2 while  $\text{size}(U) > 0$ :
3    $(\alpha, U) \leftarrow \text{deq}(U)$ 
4    $\ell_0 \leftarrow \mathcal{L}_{ip}(\alpha)$ 
5   for each  $(\beta, m)$  in MigrationDestinations( $\mathcal{M}, \alpha$ ):
6     if  $\beta \in D_{\mathcal{L}_{ip}}$ :
7        $\ell_1 \leftarrow \mathcal{L}_{ip}(\beta)$ 
8        $A_{ip}(\ell_0, \ell_1) \leftarrow A_{ip}(\ell_0, \ell_1) + m$ 
9     else:
10       $\ell_1 \leftarrow |\mathcal{L}_{ip}| + 1$ ;  $\mathcal{L}_{ip} \leftarrow \mathcal{L}_{ip} \cup \{(\beta, \ell_1)\}$ 
11       $A_{ip} \leftarrow \begin{pmatrix} A_{ip} & \mathbb{0} \\ \mathbb{0} & \mathbb{0} \end{pmatrix}_{\ell_1 \times \ell_1}$ ;  $A_{ip}(\ell_0, \ell_1) \leftarrow m$ 
12       $U \leftarrow \text{enq}(U, \beta)$ 
13       $A_{ip}(\ell_0, \ell_0) \leftarrow A_{ip}(\ell_0, \ell_0) - m$ 
14 return  $(A_{ip}, \mathcal{L}_{ip})$ 

```

Since at this point all the states \mathcal{E}_{ip} are known, as well as the index maps \mathcal{L}_i and \mathcal{L}_{ip} , the procedure simply consists in visiting, for every state α , all possible states β that can be reached by coalescence. This is accomplished by using Algorithm C and accumulating the coalescence rates in coefficient $B_{iq}(\sigma_{ip} - \sigma_i + \mathcal{L}_{ip}(\alpha), \mathcal{L}_{i+1, q}(\beta))$ for the corresponding value of q . The detailed procedures are presented in lines [9 \cdots 18] of Algorithm Q. In general, Algorithm Q presents the complete construction of the Markov process, showing the details of how to combine all algorithms specified so far. Here, we omit some of the more menial details of the process, which includes how to compute the sizes E_i and E_{ip} given the existing information, and how to assemble the final rate matrix Q and index map \mathcal{L} by arranging their respective sub-parts A_{ip} , B_{ip} and \mathcal{L}_{ip} .

Algorithm Q: For evaluating routine $\text{RateMatrix}(k, n, \mathcal{M}, \mathcal{S})$, which given the model parameters, builds the rate matrix of the Markov process and the associated data structures.

```

1  $\mathcal{L}_1 \leftarrow \{\}; \mathcal{L}_2 \leftarrow \{\}; \dots; \mathcal{L}_k \leftarrow \{\}$ 
2 for each  $a = (a_1, a_2, \dots, a_\omega)$  in  $\text{Partitions}(k)$ :
3    $\alpha \leftarrow (\mathbb{0})_{k \times n}$ 
4   for each  $j$  in  $\{1, \dots, \omega\}$ :
5      $\alpha_{a_j, 1} \leftarrow \alpha_{a_j, 1} + 1$ 
6    $i \leftarrow k - \omega + 1$ 
7    $p \leftarrow |\mathcal{L}_i| + 1; \mathcal{L}_i \leftarrow \mathcal{L}_i \cup \{(\psi_\alpha, p)\}$ 
8    $(A_{ip}, \mathcal{L}_{ip}) \leftarrow \text{MigrationSubBlock}(\mathcal{M}, \alpha)$ 
9 for each  $i$  in  $\{1, \dots, k - 1\}$ :
10  for each  $q$  in  $\{1, \dots, P_{i+1}\}$ :
11     $B_{iq} \leftarrow (\mathbb{0})_{E_i \times E_{iq}}$ 
12  for each  $p$  in  $\{1, \dots, P_i\}$ :
13     $\sigma \leftarrow E_{i1} + E_{i2} + \dots + E_{i, p-1}$ 
14    for each  $(\alpha, \ell_0)$  in  $\mathcal{L}_{ip}$ :
15      for each  $(\beta, m)$  in  $\text{CoalescenceDestinations}(\mathcal{S}, \alpha)$ :
16         $q \leftarrow \mathcal{L}_{i+1}(\psi_\beta); \ell_1 \leftarrow \mathcal{L}_{i+1, q}(\beta)$ 
17         $B_{iq}(\sigma + \ell_0, \ell_1) \leftarrow B_{iq}(\sigma + \ell_0, \ell_1) + m$ 
18         $A_{ip}(\ell_0, \ell_0) \leftarrow A_{ip}(\ell_0, \ell_0) - m$ 
19 return  $(Q, \mathcal{L})$ 

```

3.3. Computing the expected SFS

In this section we show, given the Markov process defined previously, how to compute the expected SFS for an arbitrary model structure and sample configuration. This computation is reduced to a linear algebra problem for which we propose an efficient solution in Algorithm X.

In the infinite-sites model, all mutations occur on unique sites, and as such the number of mutations is equal to the number of segregating sites. When a mutation occurs in a lineage with i descendants in the sample (a lineage of weight i), the corresponding segregating site has i copies of the mutated allele. Additionally, since we model mutations as a Poisson process of intensity θ along each branch of the genealogy, the total number of mutations that occur in lineages of weight i has an expected value of θ times the expected total length of such lineages, which we denote as ζ_i :

$$\mathbb{E}(\xi_i) = \theta \mathbb{E}(\zeta_i).$$

In the example of Figure 3.1, ζ_1 , ζ_2 and ζ_3 would be the total length of red, blue and green branches, respectively.

Since we are interested in the normalized AFS $\bar{\xi} = \xi / \|\xi\|_1$, the parameter θ is cancelled away, so we focus exclusively on the quantities $\mathbb{E}(\zeta_i)$. These can be computed by adding, for every transient state $\alpha \in \mathcal{E} \setminus \mathcal{E}_k = \mathcal{E}^*$, the number of lineages of weight i (given by $\psi_\alpha(i)$) multiplied by the expected total time spent in the state before absorption, which we denote by T_α . We thus have:

$$\mathbb{E}(\zeta_i) = \sum_{\alpha \in \mathcal{E}^*} \psi_\alpha(i) \mathbb{E}(T_\alpha). \quad (3.3.1)$$

In order to compute the times $\mathbb{E}(T_\alpha)$, we return to the sub-intensity matrix Q^* defined in (3.2.9). This matrix is non-singular, and $G = (-Q^*)^{-1}$, known as the process' Green matrix, has the property that coefficient $G(\ell_0, \ell_1)$ is the expected time that the process spends in state α^{ℓ_1} before absorption, given an initialization at state α^{ℓ_0} (see Bladt and Nielsen (2017), Theorem 3.1.14). With an initial distribution for the transient states $\pi^0 = (\pi^0(1), \pi^0(2), \dots, \pi^0(E^*))^\top$, we can then write $\mathbb{E}(T_{\alpha^\ell})$ as

$$\mathbb{E}(T_{\alpha^\ell}) = \sum_{\iota=1}^{E^*} \pi^0(\iota) G(\iota, \ell), \quad (3.3.2)$$

where $E^* = |\mathcal{E}^*| = E - E_k$. We can combine and rearrange equations (3.3.1) and (3.3.2) to get

$$\mathbb{E}(\zeta_i) = \sum_{\iota=1}^{E^*} \pi^0(\iota) \sum_{\ell=1}^{E^*} G(\iota, \ell) \psi_{\alpha^\ell}(i),$$

which can be further simplified when considering all values of $1 \leq i < k$ as

$$\mathbb{E}(\zeta) = (\pi^0)^\top \cdot G \cdot \Psi, \quad (3.3.3)$$

where Ψ , called the contribution matrix, is an $E^* \times (k-1)$ matrix where coefficient $\Psi(\ell, i) = \psi_{\alpha^\ell}(i)$ is the number of lineages of weight i in state α^ℓ .

3.3.1. Numerical solution

We now turn to the question of how to numerically solve problem (3.3.3). The computation of matrix G involves inverting Q^* , which grows exponentially in size with the number of the samples k (see Figure 3.3), so it is inevitable that the computational cost of the solution will eventually exceed the available resources regardless of the underlying method, given a large enough value of $k = k_{\text{critical}}$. Due to this inherent and unavoidable complexity, our strategy simply consists in exploiting the sparsity and structure of Q (§3.2.3) to delay k_{critical} as much as possible.

We avoid the direct computation of the inverse G . Instead, we compute the product $(\pi^0)^\top \cdot G$ by solving the linear system

$$A \cdot t = \pi^0, \quad (3.3.4)$$

where $A = (-Q^*)^\top$. The unknowns of this system are the coefficients of a vector $t \in \mathbb{R}^{E^* \times 1}$, which contains in component ℓ the time $\mathbb{E}(T_{\alpha^\ell})$ given initial distribution π^0 . Following the notation conventions of §3.2.3, we denote by t_i and π_i^0 the segments of vectors t and π^0 corresponding to states in \mathcal{E}_i , for $i = 1, \dots, k-1$; and by t_{ip} and π_{ip}^0 the sub-segments corresponding to the states in \mathcal{E}_{ip} for $p = 1, \dots, P_i$. We can thus write system (3.3.4) in the following block form:

$$\begin{pmatrix} -A_1^\top & & & & \\ -B_1^\top & -A_2^\top & & & \\ & \ddots & \ddots & & \\ & & & -B_{k-2}^\top & -A_{k-1}^\top \end{pmatrix} \cdot \begin{pmatrix} t_1 \\ t_2 \\ \vdots \\ t_{k-1} \end{pmatrix} = \begin{pmatrix} \pi_1^0 \\ \pi_2^0 \\ \vdots \\ \pi_{k-1}^0 \end{pmatrix}. \quad (3.3.5)$$

This structure immediately suggests the sequence of sub-problems:

$$\begin{aligned} -A_1^\top \cdot t_1 &= \pi_1^0, \\ -A_i^\top \cdot t_i &= b_i := \pi_i^0 + B_{i-1}^\top \cdot t_{i-1}, \quad i = 2, \dots, k-1. \end{aligned} \quad (3.3.6)$$

We note that in (3.3.6), the problems must be solved sequentially, since the solution segment t_i depends on the previous one t_{i-1} . Assuming thus that $i \geq 2$ and t_{i-1} is known, we now focus on computing t_i . Exploiting the sub-structure of the system matrix A_i (3.2.13), the problem $-A_i^\top \cdot t_i = b_i$ can be decomposed into the following P_i sub-problems (Figure 3.2):

$$-A_{ip}^\top \cdot t_{ip} = b_{ip} := \pi_{ip}^0 + B_{i-1,p}^\top \cdot t_{i-1}, \quad p = 1, \dots, P_i. \quad (3.3.7)$$

We note that the computation of the solution sub-segment t_{ip} does not depend on any other sub-segments of t_i ; only on the already-available segment t_{i-1} . This implies that all P_i problems in (3.3.7) can be solved in parallel. For each one of them we use the Successive Over-Relaxation iterative method: given an initial guess t_{ip}^0 , the scheme (known as SOR(ω)) generates a sequence of solutions $t_{ip}^1, t_{ip}^2, \dots, t_{ip}^k$, where coefficient $t_{ip}^k(\ell)$ is computed using information from $t_{ip}^k(\iota)$ for $\iota < \ell$ and $t_{ip}^{k-1}(\iota)$ for $\iota > \ell$, so there is no need to keep two versions of the solution sub-segment t_{ip} in memory. We use the norm of the residual vector $\text{res}_{ip}^k = -A_{ip}^\top \cdot t_{ip}^k - b_{ip}$ as the convergence metric, so for a given desired precision δ , the stopping criterion is $\|\text{res}_{ip}^k\|_\infty < \delta$. We denote by \hat{t}_{ip} the final approximation t_{ip}^k of the unknown vector t_{ip} . The detailed procedure is given in lines [7 ··· 10] of Algorithm X.

The parameter ω in the SOR method affects the rate of convergence. A necessary condition for convergence is $0 < \omega < 2$, and for system matrices that are M -matrices (as is the case for all the A_{ip}), a sufficient condition for convergence is $0 < \omega \leq 1$ (see §7.2.3 in Hackbusch (1994)). In practice, values of ω larger than 1 can be used to accelerate convergence despite not satisfying the sufficient convergence condition.

🏠 Math Minute

We referenced the concept of M -matrices in the previous argumentation. An M -matrix is

a matrix $A \in \mathbb{R}^{n \times n}$ satisfying the conditions:

- $A(i, i) > 0$ for all $i = 1, \dots, n$;
- $A(i, j) \leq 0$ for all $i \neq j$;
- A is non-singular, with $A^{-1}(i, j) \geq 0$ for all i and j .

It is quite straightforward to verify that all our system matrices A_{ip} are M -matrices. Indeed, all diagonal elements of $-Q^*$ are positive, all the off-diagonal elements are non-negative, and all coefficients of $(-Q^*)^{-1} = G$ are also non-negative since they can be interpreted as expected permanence times.

Algorithm X summarizes our proposed method for computing the expected SFS $\hat{\xi}$ given the rate matrix Q of the Markov process and an initial state distribution π^0 .

Algorithm X: For evaluating routine **ExpectedScaledAFS**($Q, \mathcal{L}, \pi^0, \hat{t}, \omega, \delta$), which given the data structures associated with the Markov process and an anterior distribution of states π^0 , updates an approximate solution \hat{t} to problem (3.3.4) using a $SOR(\omega)$ iterative scheme until $\|\text{res}\|_\infty < \delta$, and returns the corresponding value of $\hat{\zeta}$.

```

1  $\zeta \leftarrow (\mathbb{0})_{1 \times (k-1)}$ 
2 for each  $i$  in  $\{1, 2, \dots, k-1\}$ :
3   for each  $p$  in  $\{1, 2, \dots, P_i\}$  in parallel:
4      $b_{ip} \leftarrow \pi_{ip}^0$ 
5     if  $i > 1$ :
6        $b_{ip} \leftarrow b_{ip} + B_{i-1,p}^\top \cdot \hat{t}_{i-1}$ 
7     repeat
8       for each  $\ell$  in  $\{1, 2, \dots, E_{ip}\}$ :
9          $\hat{t}_{ip}(\ell) \leftarrow (1 - \omega)\hat{t}_{ip}(\ell) - \omega(b_{ip}(\ell) + \sum_{\nu \neq \ell} A_{ip}(\nu, \ell)\hat{t}_{ip}(\nu)) / A_{ip}(\ell, \ell)$ 
10      until  $\|A_{ip}^\top \cdot \hat{t}_{ip} + b_{ip}\|_\infty < \delta$ 
11       $\psi \leftarrow \mathcal{L}_i^{-1}(p)$ 
12      for each  $h$  in  $\{1, 2, \dots, k-1\}$ :
13         $\zeta(h) \leftarrow \zeta(h) + \psi(h) \|\hat{t}_{ip}\|_1$ 
14 return  $\zeta$ 

```

Implementation Notes

The read and write operation in line [13] must be implemented as an atomic operation due to the fact that ζ is shared across parallel threads. Otherwise, the values of $\zeta(h)$ could become corrupted as one thread reads a value, and another thread immediately writes another incompatible value to it. If atomic operations are not available, then lines [11 \dots 13] must be moved to a separate, non-parallel, p loop.

3.3.2. Error analysis

Algorithm X guarantees that after each final iteration of the SOR scheme, the residues $\text{res}_{i,p} = -A_{i,p}^\top \cdot \hat{t}_{i,p} - b_{i,p}$ satisfy $\|\text{res}_{i,p}\|_\infty < \delta$. It is then easy to verify that the global residue of system (3.3.4) given by $\text{res} = A \cdot \hat{t} - \pi^0$ satisfies $\|\text{res}\|_\infty < \delta$. This is a consequence of the fact that the vectors $\text{res}_{i,p}$ are simply sub-segments of res , and that for $X \subset \mathbb{R}$ we have $\max_{x \in X} |x| < \delta \Leftrightarrow \forall x \in X, |x| < \delta$. We can thus say that given $\delta > 0$, Algorithm X computes $\hat{\zeta} = \hat{t}^\top \cdot \Psi$, an approximation of $\mathbb{E}(\zeta)$ (3.3.3) where

$$\|A \cdot \hat{t} - \pi^0\|_\infty < \delta.$$

However, the “end user” of Algorithm X is not necessarily interested in \hat{t} or the residue, but rather in

$$\mathbb{E}(\bar{\xi}) = \frac{\mathbb{E}(\zeta)}{\|\mathbb{E}(\zeta)\|_1},$$

so it would be more meaningful to control the norm of the error $\text{err} = \mathbb{E}(\bar{\xi}) - \hat{\xi}$, where $\hat{\xi}$ is the approximation of $\mathbb{E}(\bar{\xi})$ given by Algorithm X:

$$\hat{\xi} = \frac{\hat{\zeta}}{\|\hat{\zeta}\|_1}.$$

Our goal for this section is therefore to find a way of controlling $\|\text{err}\|_\infty$ by controlling $\|\text{res}\|_\infty$.

📌 Math Minute

Some preliminary results about p -norms and their induced matrix norms:

- If $0 < p < q \leq \infty$, then for any $x \in \mathbb{R}^n$:

$$\|x\|_p \leq n^{\frac{1}{p} - \frac{1}{q}} \|x\|_q \quad \text{and} \quad \|x\|_q \leq \|x\|_p, \quad (3.3.8)$$

where we convene that $1/q = 0$ for the case $q = \infty$. This result follows from Hölder's inequality for \mathbb{R}^n vectors. No tighter bounds are possible in the general case, since equality is achieved with $x = (1, 1, \dots, 1)$ and $x = (1, 0, \dots, 0)$ respectively.

- For any non-null $x, y \in \mathbb{R}^n$ we have

$$\left\| \frac{x}{\|x\|} - \frac{y}{\|y\|} \right\| \leq \frac{2\|x - y\|}{\|x\|}, \quad (3.3.9)$$

regardless of the norm $\|\cdot\|$. Proof:

$$\begin{aligned} \left\| \frac{x}{\|x\|} - \frac{y}{\|y\|} \right\| &= \left\| \frac{x}{\|x\|} - \frac{y}{\|x\|} + \frac{y}{\|x\|} - \frac{y}{\|y\|} \right\| \\ &= \left\| \frac{1}{\|x\|}(x - y) + \left(\frac{1}{\|x\|} - \frac{1}{\|y\|} \right) y \right\| \\ &\leq \frac{\|x - y\|}{\|x\|} + \frac{|\|x\| - \|y\||}{\|x\|\|y\|} \|y\| \leq \frac{2\|x - y\|}{\|x\|}. \end{aligned}$$

- Any norm in \mathbb{R}^n can be used to define a norm in $\mathbb{R}^{m \times n}$ as follows:

$$\begin{aligned} \|A\| &= \sup \{ \|A \cdot x\| : x \in \mathbb{R}^n, \|x\| = 1 \} \\ &= \sup \{ \|A \cdot x\| / \|x\| : x \in \mathbb{R}^n, \|x\| \neq 0 \}. \end{aligned}$$

Matrix norms defined this way are called *induced* norms. Their most useful property is a direct consequence of the definition:

$$\|A \cdot x\| \leq \|A\| \|x\|,$$

In addition to the standard norm properties, induced norms also satisfy the sub-multiplicative property:

$$\|A \cdot B\| \leq \|A\| \|B\|.$$

Two notable induced norms:

$$\|A\|_1 = \max_{1 \leq j \leq n} \sum_{i=1}^m |a_{ij}|, \quad \|A\|_\infty = \max_{1 \leq i \leq m} \sum_{j=1}^n |a_{ij}|.$$

Notice that $\|A\|_\infty = \|A^T\|_1$.

A lower bound for $\|\text{err}\|_\infty$ can be derived as follows:

$$\begin{aligned}
\|\text{err}\|_\infty &\leq \|\text{err}\|_1 \\
&= \left\| \frac{\mathbb{E}(\zeta)}{\|\mathbb{E}(\zeta)\|_1} - \frac{\hat{\zeta}}{\|\hat{\zeta}\|_1} \right\|_1 \leq \frac{2\|\mathbb{E}(\zeta) - \hat{\zeta}\|_1}{\|\mathbb{E}(\zeta)\|_1} \\
&= \frac{2\|\Psi^\top \cdot (t - \hat{t})\|_1}{\|\mathbb{E}(\zeta)\|_1} = \frac{2\|\Psi^\top \cdot (A^{-1}\pi^0 - \hat{t})\|_1}{\|\mathbb{E}(\zeta)\|_1} \\
&= \frac{2\|\Psi^\top \cdot A^{-1} \cdot (\pi^0 - A \cdot \hat{t})\|_1}{\|\mathbb{E}(\zeta)\|_1} = \frac{2\|\Psi^\top \cdot A^{-1} \cdot \text{res}\|_1}{\|\mathbb{E}(\zeta)\|_1} \\
&\leq \frac{2\|\Psi^\top\|_1 \|A^{-1}\|_1 \|\text{res}\|_1}{\|\mathbb{E}(\zeta)\|_1} \leq 2E^* \frac{\|\Psi\|_\infty \|G\|_\infty \|\text{res}\|_\infty}{\|\mathbb{E}(\zeta)\|_1}.
\end{aligned}$$

We note that since $\|\Psi\|_\infty$ is the largest possible row sum of Ψ , and the rows of Ψ are the contribution vectors of the states of \mathcal{E}^* , then $\|\Psi\|_\infty = \max_{\alpha \in \mathcal{E}^*} |\alpha| = k$. We also note that $\|G\|_\infty = \|(Q^*)^{-1}\|_\infty$. With this we have

$$\|\text{err}\|_\infty \leq \frac{2kE^* \|(Q^*)^{-1}\|_\infty}{\|\mathbb{E}(\zeta)\|_1} \|\text{res}\|_\infty. \quad (3.3.10)$$

In order to continue we need an upper bound for $\|\mathbb{E}(\zeta)\|_1$. In what follows we note that since $\Psi \geq 0$ and $G^\top = A^{-1} \geq 0$, then all of the coefficients of $G \cdot \Psi$ are also non-negative.

$$\begin{aligned}
\|\mathbb{E}(\zeta)\|_1 &= \|(\pi^0)^\top \cdot G \cdot \Psi\|_1 = \sum_{i=1}^{k-1} \left| \sum_{\ell=1}^{E^*} \pi^0(\ell) (G \cdot \Psi)(\ell, i) \right| \\
&= \sum_{\ell=1}^{E^*} \pi^0(\ell) \sum_{i=1}^{k-1} (G \cdot \Psi)(\ell, i) \\
&= \sum_{\ell=1}^{E^*} \pi^0(\ell) \sum_{i=1}^{k-1} \sum_{\iota=1}^{E^*} G(\ell, \iota) \Psi(\iota, i) \\
&= \sum_{\ell=1}^{E^*} \pi^0(\ell) \sum_{\iota=1}^{E^*} G(\ell, \iota) \sum_{i=1}^{k-1} \Psi(\iota, i) \\
&= \sum_{\ell=1}^{E^*} \pi^0(\ell) \sum_{\iota=1}^{E^*} G(\ell, \iota) |\alpha^\iota|.
\end{aligned}$$

Since $\iota \leq E^*$, α^ι is a transient state and therefore has at least two live lineages, thus:

$$\begin{aligned}
\|\mathbb{E}(\zeta)\|_1 &\geq \sum_{\ell=1}^{E^*} 2\pi^0(\ell) \sum_{\iota=1}^{E^*} G(\ell, \iota) \\
&\geq \sum_{\ell=1}^{E^*} 2\pi^0(\ell) G(\ell, \ell).
\end{aligned}$$

The coefficient $G(\ell, \ell)$ is $\mathbb{E}(T_{\alpha^\ell} | X_0 = \alpha^\ell)$, the expected total time the process spends in state α^ℓ prior to absorption, given that the process started in that same state. This expected

time can be computed as the product of the total expected number of visits to state α^ℓ $\mathbb{E}(N_{\alpha^\ell}|X_0 = \alpha^\ell)$, multiplied by the expected duration of each visit, which is an exponentially distributed time with parameter $-Q(\ell, \ell)$. An upper bound for $G(\ell, \ell)$ thus can be obtained by observing that $\mathbb{E}(N_{\alpha^\ell}|X_0 = \alpha^\ell) \geq 1$, since at least one visit is expected at state α^ℓ in lieu of it being the initial state. Additionally, we note that:

$$-Q(\ell, \ell) \leq \sum_{\iota=1}^{E^*} |Q(\ell, \iota)| \leq \max_{1 \leq \ell \leq E^*} \sum_{\iota=1}^{E^*} |Q(\ell, \iota)| = \|Q^*\|_\infty.$$

This implies that the expected duration of one visit to state α^ℓ is greater than $1/\|Q^*\|_\infty$. With this we have:

$$\|\mathbb{E}(\zeta)\|_1 \geq \frac{2}{\|Q^*\|_\infty} \sum_{\ell=1}^{E^*} \pi^0(\ell) = 2 \frac{\|\pi^0\|_1}{\|Q^*\|_\infty}.$$

Returning then to (3.3.10):

$$\begin{aligned} \|\text{err}\|_\infty &\leq \frac{2kE^* \|(Q^*)^{-1}\|_\infty}{\|\mathbb{E}(\zeta)\|_1} \|\text{res}\|_\infty \\ &\leq \frac{kE^* \|Q^*\|_\infty \|(Q^*)^{-1}\|_\infty}{\|\pi^0\|_1} \|\text{res}\|_\infty \\ &= \frac{kE^* \kappa(Q^*)}{\|\pi^0\|_1} \|\text{res}\|_\infty, \end{aligned} \tag{3.3.11}$$

where $\kappa(Q^*)$ is the condition number of the matrix Q^* . These results can be summarized in the following

Lemma 2. For any $\varepsilon > 0$, it is sufficient to choose δ as

$$\delta = \frac{\|\pi^0\|_1}{kE^* \kappa(Q^*)} \varepsilon \tag{3.3.12}$$

to have

$$\|\text{res}\|_\infty < \delta \Rightarrow \|\text{err}\|_\infty < \varepsilon.$$

Having a clear criterium for how to control the global error of the normalized expected SFS $\mathbb{E}(\bar{\xi})$, we can proceed to present the main and final algorithm of this chapter:

Algorithm R: For evaluating routine **ExpectedAFS**($k, n, \mathcal{M}, \mathcal{S}, \alpha^0, \omega, \varepsilon$), which given the model parameters and an initial sampling state $\alpha^0 \in \mathcal{E}_1$, returns an approximation of $\mathbb{E}(\bar{\xi})$ with global error $\|\text{err}\|_\infty < \varepsilon$ using a **SOR**(ω) iterative scheme.

- 1 $(Q, \mathcal{L}) \leftarrow \mathbf{RateMatrix}(k, n, \mathcal{M}, \mathcal{S})$
 - 2 $\pi^0 \leftarrow (\mathbb{0})_{1 \times E^*}; \pi^0(\mathcal{L}^{-1}(\alpha^0)) \leftarrow 1$
 - 3 $\delta \leftarrow \varepsilon \|\pi^0\|_1 / (kE^* \kappa(Q^*))$
 - 4 $t^0 \leftarrow (\mathbb{0})_{E^* \times 1}$
 - 5 $\hat{\zeta} \leftarrow \mathbf{ExpectedScaledAFS}(Q, \mathcal{L}, \pi^0, t^0, \omega, \delta)$
 - 6 **return** $\hat{\zeta} / \|\hat{\zeta}\|_1$
-

3.3.3. Early termination

The following discussion is concerned with the eventual use of Algorithm R in the context of demographic inference, where early termination is desirable.

Given a value of ε , Algorithm R computes a numerical approximation of $\mathbb{E}(\bar{\xi})$ with an error $\|\text{err}\|_\infty < \varepsilon$. However, the internal solution method we use for each of the P sub-problems (3.3.7) is iterative, meaning that progressively-accurate approximations for which t_{ip} are being computed, which in turn could be used to compute progressively-accurate approximations of $\hat{\zeta}$. Obtaining early approximations of $\hat{\zeta}$ (i.e., approximations where $\|\text{err}\|_\infty$ is not smaller than ε yet) could be useful for implementing more flexible stopping criteria, where we allow the method to converge to the desired high accuracy ε only if the SFS is approaching a previously designated *target* SFS; and otherwise terminate the process prematurely, thus saving computational resources.

The key insight for our proposal comes from the observation that Algorithm X operates by simply *updating* or *replacing* a starting solution t^0 with a more accurate one \hat{t} , for which $\|\text{res}\|_\infty < \delta$. This means that it is possible to continue refining the solution by repeatedly invoking the routine, each time passing as t^0 the final result \hat{t} from the previous invocation, and a smaller value of δ .

One possible realization of this idea is as follows: suppose we are interested in observing d intermediate solutions $\hat{\zeta}$ prior to the convergence with $\|\text{err}\| < \varepsilon$, then we may define d intermediate tolerances:

$$\varepsilon_1 = \varepsilon^{1/d} \quad \dots \quad \varepsilon_\kappa = \varepsilon^{\kappa/d} \quad \dots \quad \varepsilon_d = \varepsilon, \quad (3.3.13)$$

along with the corresponding $\delta_\kappa = \delta^{\kappa/d}$, and invoke Algorithm X d times: once for each of the δ_κ . After each invocation we evaluate the corresponding early approximation of $\mathbb{E}(\hat{\xi})$, and decide whether to continue or terminate the process. The sequence (3.3.13), while certainly not unique, satisfies the desirable properties that the ε_κ are in descending order (assuming $\varepsilon < 1$), and are uniformly log-spaced. This implies that the subsequent invocations of Algorithm X, which improve the accuracy of the solution from δ_κ to $\delta_{\kappa+1}$, should take similar times to execute.

There is a tradeoff to be balanced between execution time and level of control when choosing the value of d , since the products $B_{i-1,p} \cdot \hat{t}_{i-1}$ and $\Psi^\top \cdot \hat{\zeta}$ (lines [6] and [12...13] of Algorithm X, respectively) have to be computed d times each. See Figure 3.10.

The following revision of Algorithm R takes an additional parameter (d), and visits d

increasingly-accurate approximations of $\mathbb{E}(\bar{\xi})$.

Algorithm R': For evaluating an overload of routine **ExpectedAFS**($k, n, \mathcal{M}, \mathcal{S}, v, \omega, d, \varepsilon$) allowing early termination.

```

: :
4  $\hat{t} \leftarrow (\mathbb{0})_{E^* \times 1}$ 
5 for each  $\kappa$  in  $\{1, \dots, d\}$ :
6    $\hat{\zeta} \leftarrow \mathbf{ExpectedScaledAFS}(Q, \mathcal{L}, \pi^0, \hat{t}, \omega, \delta^{\kappa/d})$ 
7   visit  $\hat{\zeta} / \|\hat{\zeta}\|_1$ 

```

We can see in Figure 3.10 the tradeoff between computation time and error introspection granularity in an n -island model. In general, there are diminishing returns with values of d larger than 10, although the exact optimal value depends heavily on the application. In all cases, the fastest results are achieved with the fewer introspection steps.

3.4. Model specialization: Symmetrical n -island

Despite the efficient algorithms presented in the previous sections, the state space grows too quickly and becomes computationally intractable even for small values of k . For this reason, we believe that the most sensible application for these methods is through *model specializations*. A model specialization is a way of compressing the state space by taking advantage of the symmetries present in a specific demographic model. In this section we use the n -island model as a case study for introducing model specializations.

The n -island model represents the worst case scenario regarding the size of the state space. Indeed, since all islands are reachable from all other islands, the migration matrix \mathcal{M} is completely dense, which in turn implies that the transition rate matrix Q is the largest and least sparse possible. On the other hand, the n -island model is the most symmetrical demographic model possible: all islands have the same sizes and connectivity, thus they are only distinguishable through the number and configuration of lineages present. Consider for instance the following states in a symmetrical 3-island model with $k = 4$ sampled lineages:

$$\begin{array}{c} \alpha^1 \\ \begin{array}{|c|c|c|} \hline 2 & & \\ \hline & 1 & \\ \hline & & \\ \hline & & \\ \hline \end{array} \end{array}, \begin{array}{c} \alpha^2 \\ \begin{array}{|c|c|c|} \hline 2 & & \\ \hline & & 1 \\ \hline & & \\ \hline & & \\ \hline \end{array} \end{array}, \begin{array}{c} \alpha^3 \\ \begin{array}{|c|c|c|} \hline & 2 & \\ \hline 1 & & \\ \hline & & \\ \hline & & \\ \hline \end{array} \end{array}, \begin{array}{c} \alpha^4 \\ \begin{array}{|c|c|c|} \hline & 2 & \\ \hline & & 1 \\ \hline & & \\ \hline & & \\ \hline \end{array} \end{array}, \begin{array}{c} \alpha^5 \\ \begin{array}{|c|c|c|} \hline & & 2 \\ \hline 1 & & \\ \hline & & \\ \hline & & \\ \hline \end{array} \end{array}, \begin{array}{c} \alpha^6 \\ \begin{array}{|c|c|c|} \hline & & 2 \\ \hline & 1 & \\ \hline & & \\ \hline & & \\ \hline \end{array} \end{array}. \quad (3.4.1)$$

The states in (3.4.1) are different only because we have numbered the islands, but since the islands are indistinguishable, they can all be grouped in the same equivalence class. This equivalence class is characterized by one of the islands having two lineages of weight 1 and another island having one lineage of weight 2. This class can be represented by a state matrix

where the columns no longer indicate the island number, but are simply an arbitrary indexing of the non-empty islands. We can define a unique class representative by sorting the columns of the equivalent state matrices. We use the reverse-lexicographic ordering on the columns (read from top to bottom). With this ordering, the class representative of the equivalence set (3.4.1) would be α^1 .

We will assume the existence of a function **ClassRepresentative** that maps \mathcal{E} onto itself, and assigns to every state matrix its class representative. In this section we present modified versions for some of the previously derived algorithms in order for them to work for the n -island model specialization.

The results shown in this section were obtained with the software SISiFS (short for Symmetrical Islands Site Frequency Spectrum). The implementation was done using the D programming language, and can be found in the public repository github.com/arredondos/sisifs.

3.4.1. The rate matrix, revisited

We saw in §3.2.2 that the model state space is discovered by generating all possible migration and coalescence events (algorithms M and C, respectively). In order to adapt these for a symmetrical n -island specialization, all we have to do is to only visit the states class-representatives, which as discussed previously, are obtained by sorting the columns of the state matrices. Therefore, the adaptation will consist of, given a destination state β , generating its class representative β^* and visiting it instead.

Algorithm sM: For evaluating routine **SymmetricalMigrationDestinations**(m, α), which given the migration rate m and a state matrix α of size $k \times n$, visits all the tuples (β, m) where β is a class-representative state satisfying (3.2.3) for symmetrical island models, and m is the corresponding transition rate.

```

1 for each  $i$  in  $\{1, 2, \dots, k - 1\}$ :
2   for each  $j$  in  $\{1, 2, \dots, n\}$ :
3     if  $\alpha_{ij} \geq 1$ :
4       for each  $g$  in  $\{1, 2, \dots, j - 1, j + 1, \dots, n\}$ :
5          $\beta \leftarrow \alpha - e_{ij} + e_{ig}$ ;  $\beta^* \leftarrow \mathbf{ClassRepresentative}(\beta)$ 
6         visit  $(\beta^*, \alpha_{ij} m)$ 

```

This simple modification allows for a large compression of the state space. See Figure 3.3 for a comparison of the number of states of the Markov systems associated to the symmetrical n -island model, and the most general demographic model possible, where all the states are

present.

Algorithm sC: For evaluating routine **SymmetricalCoalescenceDestinations**(s, α), which given the deme size s of a symmetrical island model and a state matrix α , visits all the tuples (β, c) where β is a class-representative state satisfying (3.2.4) for this model, and c is the corresponding coalescence rate.

```

1 for each  $j$  in  $\{1, 2, \dots, n\}$ :
2   for each  $i$  in  $\{1, 2, \dots, k - 1\}$ :
3     if  $\alpha_{ij} \geq 1$ :
4       if  $\alpha_{ij} \geq 2$ :
5          $\beta \leftarrow \alpha - 2e_{ij} + e_{2ij}$ ;  $\beta^* \leftarrow \mathbf{ClassRepresentative}(\beta)$ 
6          $c \leftarrow \alpha_{ij}(\alpha_{ij} - 1)/(2s)$ 
7         visit  $(\beta^*, c)$ 
8       for each  $h$  in  $\{i + 1, \dots, k - 1\}$ :
9         if  $\alpha_{hj} \geq 1$ :
10           $\beta \leftarrow \alpha - e_{ij} - e_{hj} + e_{i+h,j}$ ;  $\beta^* \leftarrow \mathbf{ClassRepresentative}(\beta)$ 
11           $c \leftarrow \alpha_{ij} \alpha_{hj}/s$ 
12          visit  $(\beta^*, c)$ 

```

The simplification of the Markov system provided by the n -island model specialization goes beyond simply reducing the number of states. Indeed, since now all migration events happen with the same rate m , and all coalescence events with the same rate $1/s$, the coefficients of the rate matrix also become greatly simplified. We can observe in Algorithm sM (line 6) that the coefficients of the migration blocks A_i are of the form $\sum_{\ell, \iota} \alpha_{\ell, \iota} m = m C_{\ell, \iota}$. Similarly for the coalescence blocks B_i (see algorithm sC, lines 7 and 12), the transition rates are of the form $\frac{1}{s} C_{\ell, \iota}$. Consequently, we can decompose the rate matrix Q of a symmetrical n -island model in the following form:

$$Q = m Q_m + \frac{1}{s} Q_c,$$

The matrices Q_m and Q_c hold the coefficients $C_{\ell, \iota}$ for the migration and coalescence blocks, respectively, and they are invariant given k and n fixed. We can use this representation in order to greatly reduce the computational cost of building the rate matrix (algorithm Q) for the same values of k and n but many different values of m and s . See figure 3.4 for a comparison of the computational resources required to execute algorithm Q and those required to build the rate matrix given Q_m and Q_c .

All time and memory benchmarks presented in this and the next sections were performed on a Ryzen 3700X desktop system with 32GB of RAM. As can be seen in Figure 3.4, the RAM capacity was the eventual limiting factor in how big a sample could be processed ($k = 26$ in this case). We note that larger samples could be processed by having systems with larger main memory compute the symbolic rate matrix representation remotely and then distribute this file to the local system where the analyses are being conducted. The size of this file, and




Figure 3.3: The size of the state space in various Markov processes associated with n -island models. Panel (a) shows the number of states as a function of the sample size k . As a point of reference, we include in this panel the number of states in the isolation-with-migration model of Kern and Hey (2017), which computes the exact expected joint SFS of two populations (black plot, data extracted from Figure 2). Panel (b) shows the same information as a function of the number of islands n . We see that with the state space compression (dashed plots), the number of states does not depend on the number of demes once $n \geq k$, as expected.

the RAM cost to load it is also displayed in Figure 3.4.




Figure 3.4: Computational costs associated with the Markov state system for various n -island models as a function of the sample size k . Panel (a) shows the maximum space required in main memory to generate and use the transition rate matrix as per Algorithm Q (blue plot). In the case that a symbolic representation of this matrix is available, then the space required in main memory to use it for computing the expected SFS as per Algorithm X is shown in the green plot. The grey plot shows the space required for storing this symbolic representation of the rate matrix. Panel (b) shows the time in seconds that it takes to generate the rate matrix as per Algorithm Q (blue plot) and the time it takes to load this matrix from storage (orange plot). We note that the space and time costs depicted in the blue plots need only be incurred once for each unique pair of (k, n) values.

3.4.2. The expected SFS in the n -island model

In this section we focus on the problem of computing the expected SFS in the n -island model. Specifically, we are interested in validating the correctness of algorithms X and R', with the n -island-specific modifications introduced in §3.4.1, as well as benchmarking their performance.

We begin by taking a look at the expected SFS itself. Figure 3.5 show the plot of $\mathbb{E}(\zeta)/\|\mathbb{E}(\zeta)\|$ for various parameter values and sampling patterns. Overall, we see that the sampling vector is a key determining factor for the shape of the SFS. Depending on the model parameters, it is relatively common for all or most of the lineages that are sampled in the same deme to coalesce among them before any of them migrates to a different deme. This creates an excess of mutations shared between all the lineages of that deme. This effect can be clearly seen in the figure. For instance, in panel (b) there is a spike at frequency 13 since there are two islands that start with 13 lineages each, and in panel (c) where the sampling

vector is $[14, 6, 6]$, there are spikes in the expected SFS at frequencies 14, 6, $14+6$ and $6+6$. Another observation is that this effect is amplified when the migration value is small, since it increases the probability that coalescence happens before migration in a deme. Likewise, a smaller number of islands also amplifies this effect of structure, although to a lesser degree than M .

We can quantify the effect of structure on the SFS by measuring the deviation of the expected SFS in these n -island models from the expected SFS under panmixia (1.4.3). We present these results in Figure 3.6. We notice the same spikes near the sampling frequencies, but also that for many of the structure parameters, there is a smaller-than-expected (as compared to a panmictic model) number of unique mutations (singletons and doubletons) and a higher-than-expected number of shared mutations.

These results are not entirely new (see for instance Figure 4.11 in Hein et al. (2004)). However, we can now compute these figures exactly (up to machine precision), and very quickly. Figure 3.7 shows the required time to compute run Algorithm X for some parameter values of the n -island. The sizes of the associated Markov chains can be referenced in Figure 3.3. For instance, we observe that for a sample of size $k = 26$, the rate matrix has over 8 million states, and the implementation of Algorithm X in SISiFS iteratively solves the associated linear system to a precision of $\varepsilon = 10^{-6}$ in under 2 seconds, if the migration rate is not too large. The migration rate M affects the performance of the method because the condition number of the system matrices $\kappa(A_{ip})$ increases rapidly with large values of M .

The effect of the high condition number for large M can also be observed as numerical instability in Figure 3.8. This figure showcases the convergence of the expected SFS under various n -island models to the expected SFS under panmixia, for which the closed-form expression is known (1.4.3) when $M \rightarrow \infty$. This convergence is expected since island models with a very high migration rate are known to behave like panmictic models with a population size equal to the sum of the island sizes. We verify this fact by making M increasingly larger and comparing the two expected SFSs by computing the $\|\cdot\|_\infty$ -based distance between them. We observe that the algorithm converges exponentially to the correct solution in all but the largest cases ($k \geq 24$). This result serves as a validation of Algorithm X, but also indicates that in practice, the degradation of speed and accuracy due to very large values of M does not impose a great limitation on the method, given that the expected SFS can be well-approximated by (1.4.3) in such cases.

As an additional form of validation, we conducted simulations using the ms software and computed the empirical SFS for samples of various sizes under the n -island model. We simulated independent chromosomes continuously with a given value of θ until we achieved a certain number of observed segregating sites. After that, all the information at these sites was aggregated and an empirical SFS was computed. The goal is to compare the obtained results with the ones given by our methods. This approach validates both algorithms Q and X. Figure 3.9 shows the results of these comparisons. We find that estimating the expected SFS with this method is a poor substitute for exactly computing it using our approach for the tested sampling sizes ($k \leq 26$), both from the numerical accuracy perspective as well as from

the required computation time (we note that even though Figure 3.9 does not show the times required to complete these ms simulations, our method was so many orders of magnitudes faster for the same accuracies, that we felt that a direct comparison was unnecessary).

Figure 3.10 showcases the early-termination feature of Algorithm R' (see §3.3.3). Here, we can see the trade-off between the speed of convergence and the flexibility afforded by the error introspection. Increasing the value of the parameter d allows for additional intermediate results of the expected SFS to be available. This could be useful for terminating the computation of the expected SFS if it is not converging towards the desired curve. For instance, in the context of a hypothetical demographic inference application, the computation of the SFS could be aborted before reaching the accuracy specified by ε if the expected SFS is not converging towards the target (observed) SFS, and the search algorithm could then focus on testing other parameters.

As expected, the fastest results are achieved with $d = 1$, but since the convergence is often achieved in one or two iterations, it also allows for minimum flexibility. This is the recommended value for one-time computations. In general, the optimum value of the d parameter will be heavily dependant on the application that Algorithm R' is being used in.




Figure 3.5: Expected SFS of the n -island model for various n -island models. All the displayed SFSs correspond to a sample of size $k = 26$ haploids, and the different panels show the effect of different sampling patterns. The curves corresponding to different number of islands are displayed in different colors, and different migration values are displayed in different line styles. We omit the $n = 5$ plots in panel (d) because they are not compatible with the corresponding sampling pattern.

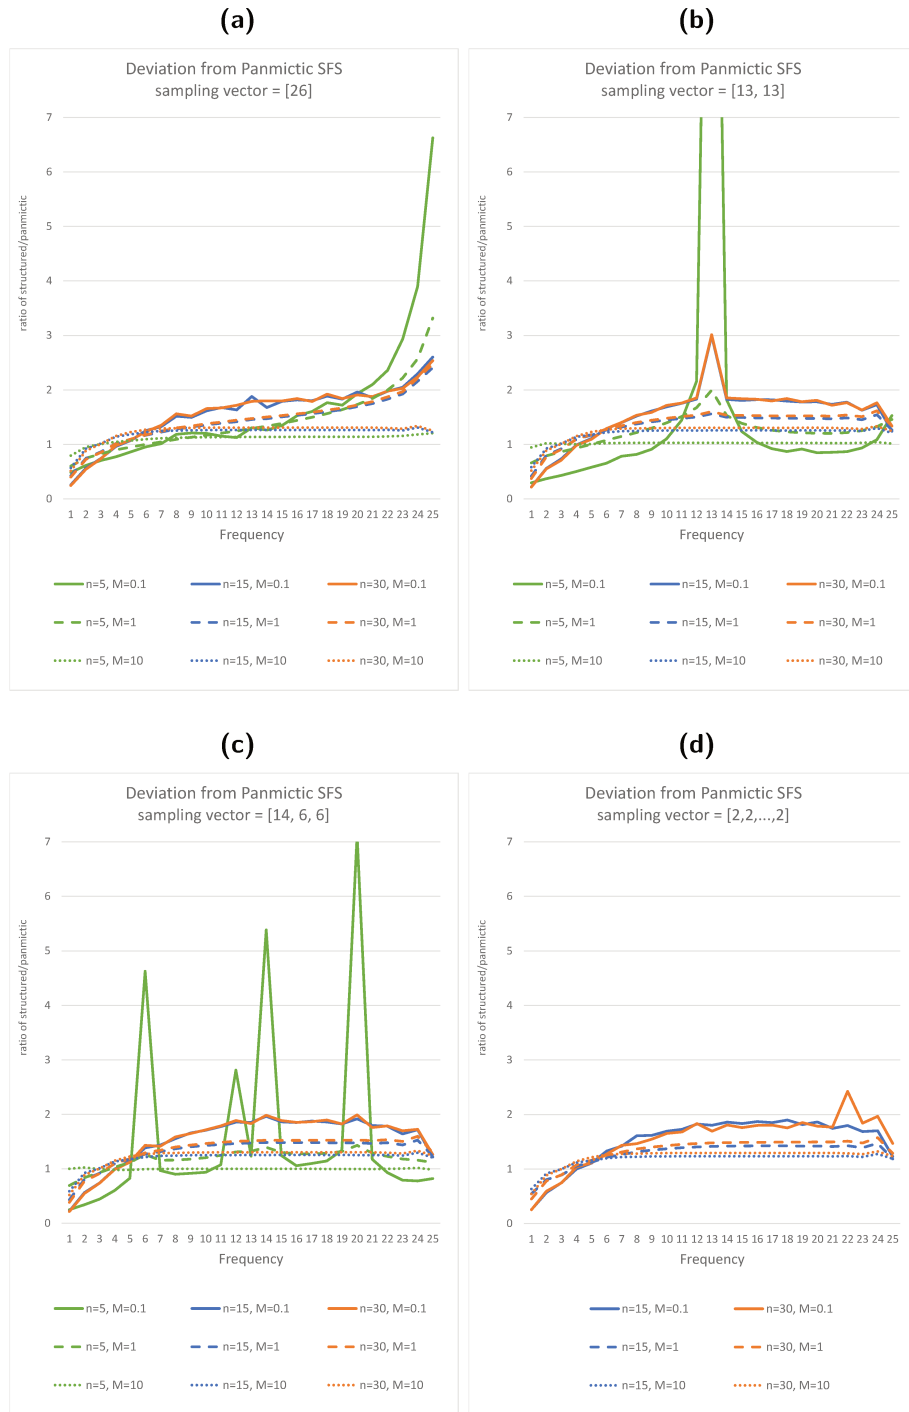


Figure 3.6: Ratio of the expected SFS in n -island models to the expected SFS in panmictic models, for various model parameters. All the displayed SFSs correspond to a sample of size $k = 26$ haploids, and the different panels show the effect of different sampling patterns. The curves corresponding to different number of islands are displayed in different colors, and different migration values are displayed in different line styles. We omit the $n = 5$ plots in panel (d) because they are not compatible with the corresponding sampling pattern.




Figure 3.7: Time to compute the expected SFS in various n -island models as a function of the sample size k . Each plot shows the time required in seconds to perform a single run of Algorithm X for different values of the migration rate M . We note that significantly increasing the value of M results in slower convergence. This can be attributed to the increase in the condition number of the rate matrix, which is correlated with the numerical difficulty of the problem. The number of demes is fixed at $n = 30$ so that the size of the state space only depends on the value of k , therefore the time complexity can equivalently be read as being a function of the size of the state space. The error tolerance was fixed at $\epsilon = 10^{-6}$.

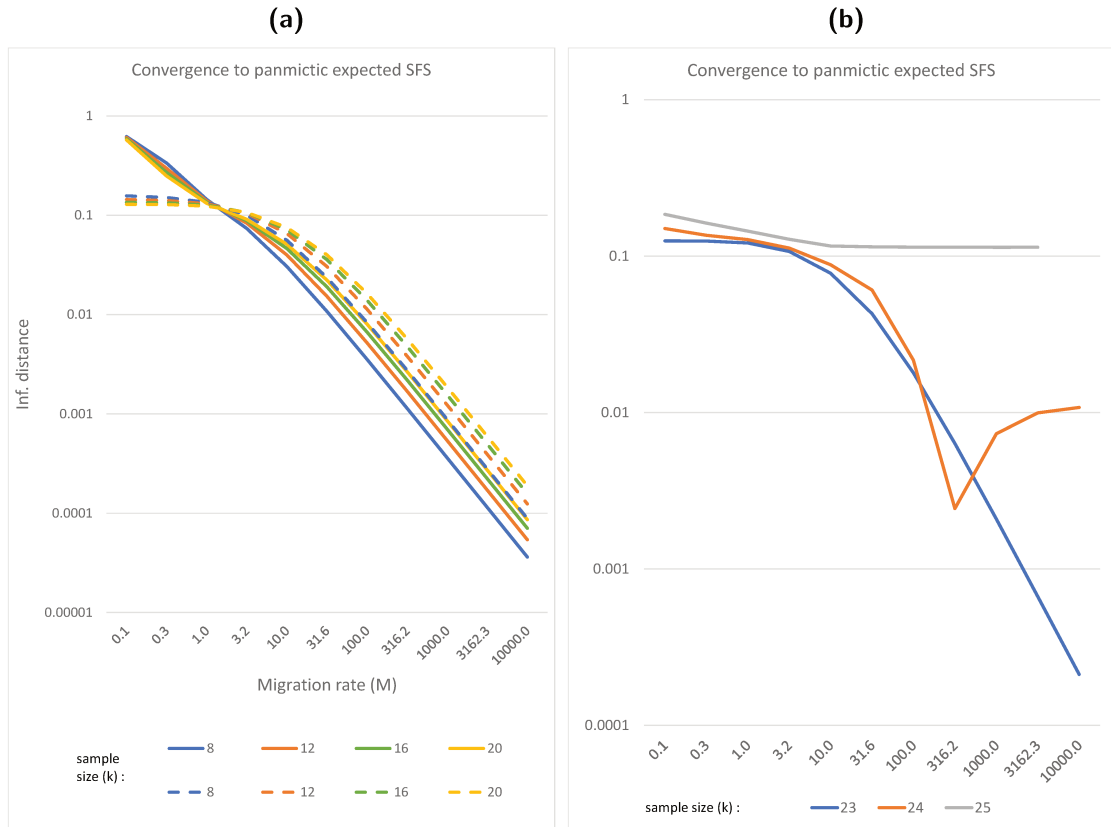


Figure 3.8: Distance between the expected SFS under panmixia and the expected SFS under various n -island models as a function of the migration rate M . Panel (a) shows the $\|\cdot\|_\infty$ distance between the expected SFS as computed by Algorithm X and the one expected under panmixia (1.4.3) for various sample sizes k . The exponential convergence is expected under the presupposition that an n -island model with very high migration rate functions equivalently to a panmictic model where the size of the population is the sum of the sizes of the demes. The different plot clustering is due to different sampling schemes: $(k, 0, \dots, 0)$ for the solid plots and $(k/2, k/2, 0, \dots, 0)$ for the dashed plots. Panel (b) shows the same information for larger sample sizes. We observe that the numerical accuracy of the method degrades for very large values of M in these cases. This may be attributed to the increase in the condition number of the rate matrix, which is associated with the numerical difficulty of the problem.





Figure 3.9: Distance between the expected SFS under one n -island model and the empirical SFS obtained from *ms* simulation as a function of the number of segregating sites. For the representative n -island model we chose $n = 30$ islands with a migration rate of $M = 1$ and $k = 6$ lineages sampled in the same deme. The plots show the $\|\cdot\|_\infty$ distance between the expected SFS for this model and the one obtained by building the observed SFS after many rounds of *ms* simulations under the same structured model. The different plots correspond to different θ values used in the simulations. We notice that lower values of θ result in more accurate empirical SFSs as expected, however the simulation approach only achieves acceptable accuracy after a very large number of simulations (note that for lower values of θ , a greater number of simulations is required to achieve the same amount of segregating sites).

(a)



(b)

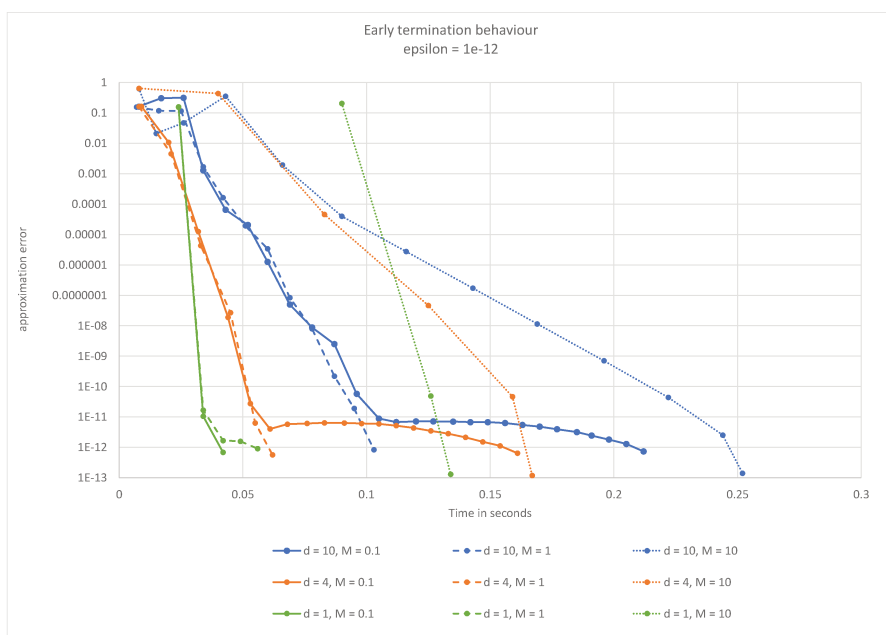



Figure 3.10: Method error as a function of the computation time.



Chapter 4.

Conclusions and future work

In this final section of the thesis we summarize the main results from the previous chapters in a forward-looking context, where the natural steps for the continuation of the research are outlined.

4.1. The IICR and the inference of structure

In Chapter 2 we presented an inference framework that is able to accurately infer structure parameters (number of islands and their sizes) within a symmetrical island model given an IICR estimate like the PSMC. It is also able to date up to five events of changes in migration rate (i.e., six components) with good precision and consistency, as long as the underlying model is compatible with a symmetrical island model with constant population size.

A comprehensive validation effort using simulated data support these results. Indeed, the normalized RMSD (nRMSD) (Figure 2.4) of the simulated vs. inferred scenario parameters is zero for stationary scenarios ($c = 1$), and increases linearly with the number of components. The M parameters reach a value of about 0.5 at six components, and we see that the first (most recent) and last (most ancient) components are better estimated than the middle ones. It is likely that these middle components are being affected by what we called the component misidentification phenomenon, where the migration rate for a given component is being inferred for a neighboring one (previous or following), thus contributing to a negative correlation between the simulated and inferred migration values for a given component. This can happen when, for instance, a given demographic event (with an associated change of migration rate) does not have a large-enough effect on the shape of the IICR curve. Regardless of this behavior, the number of islands n and the reference effective size N are consistently well estimated, reaching an nRMSD of about 0.1 in the worst cases. The event time parameters (t_i) exhibit the worst nRMSD values, varying between 1 and 2 in the worst cases. Although in these cases, the fact that time is log-spaced and spans several orders of magnitude causes outliers to have a disproportionate contribution to this statistic.

One novel aspect of our approach is that the number of demes gets inferred as one of the model parameters, and it is in fact the best estimated parameter, which is in agreement with Mazet et al. (2015) that used information from the distribution of T_2 values and a likelihood approach. These authors however, only analysed stationary models. Here we found that other parameters were also well-estimated when the number of components was low, but we

also observed that the estimated value of n scaled well with increasing model complexity. A similar consistency can be observed with the deme size parameter N (see Figure 2.4). We give up some flexibility in the model by keeping the number of demes constant throughout the history of the population, so the timed demographic events cannot represent splits or joining of populations even though such events are likely to have taken place in the history of species. Additionally, in the n -island model we do not account for possible asymmetrical gene flow or different deme sizes, even when the theoretical framework does allow for such representations. However, it is a more challenging problem to validate due to the fact that during any given component, changing both the migration rate and the deme size have confounding effects on the IICR curve which can be hard to separate. This requires a dedicated study with a different methodology which will be explored in a future work.

4.2. A note on human evolution

An application of our IICR-based inference method to five publicly available human PSMCs suggests that the backwards long term history of the sampled individuals, when accounting for possible recent expansions and the noise introduced by the PSMC method, can be accurately modelled in the framework of a symmetrical island model of approximately 10 to 12 demes with varying levels of connectivity through time. Only one of the five samples (Yoruba) displayed less consistent evidence of this finding, which may indicate that more complex models (possibly including asymmetric gene flow, spatial modelling of the environment, or changes in deme sizes) could be needed to explain the full complexity of the data.

These findings regarding changes in connectivity and number of islands are in agreement with the results of Rodríguez et al. (2018), in which a hand-fitting approach of the IICRs was used to arrive at an estimate of 10 islands with a similar value of N and a comparable period featuring a significant increase of gene flow between 600 Kya and 2 Mya. Note that the timing in years and the deme sizes in Rodríguez et al. (2018) differ due to the change in mutation rate.

We also compared our results with the tree model for human evolution published by Noskova et al. (2019) (the C3PO model), which is a revision of the model from Gutenkunst et al. (2009) and represents a simplified model of human evolution (Jouganous et al., 2017; Kamm et al., 2019). The C3PO model proposes an ancestral human population that experiences two splits: an old one that resulted in the current African “population” and another more recent one that resulted in the current European and Asian “populations”. The parameters of this model include the times of these events, the population size history of these populations and their ancestral branches, and the migration rates between them. The summary statistic targeted by these methods is the AFS, and we see that a fitting AFS does not guarantee a fitting IICR and vice versa (Chikhi et al., 2018; Beichman et al., 2017). Indeed, the IICRs of the populations from the C3PO model do not resemble those of the real humans. Likewise, when we use the C3PO model to generate IICR curves, and infer the corresponding demographic history using SNIF, these inferred n -island models do not resemble those inferred directly from the human IICRs, and the parameter values are less consistent across different runs.

These findings suggest that tree models fitted with the SFS, like those considered here, do not offer a definitive answer to the question of human evolution and other families of models should be explored (Goldstein and Chikhi, 2002; Scerri et al., 2018, 2019). It remains to be seen however how well models inferred with our method fit the real SFS of their respective human populations. A more general treatment of this question is a good candidate for future research. However, in §2.4.3 we compared the SFS of a sample of 216 humans from the Yoruba population (Lapierre et al., 2017) to the one inferred by the GADMA method from Noskova et al. (2019) and the one corresponding to three variations of the inferred demographic model by our method (see Figure 2.25). These simulations suggest that existing SFSs could be easily fitted with a structured model similar to those inferred by SNIF, but in which we would allow for a recent population size change.

4.3. Computing the exact SFS for structured populations

The SFS has long been an important statistic due to its capacity to succinctly summarize a large amount of information about the genetic diversity of a sample and the ease with which it can be computed using real data.

Comparable approaches to utilize the SFS with structured populations (Gutenkunst et al., 2009; Excoffier et al., 2013; Kern and Hey, 2017) focused on what is known as the joint SFS, where information regarding the origin of the lineages is preserved along with the frequency of the mutations. This approach limits the number of populations that can be modeled, since a high-dimensional joint SFS is both computationally intractable and uninformative in practice, since one would expect that most of the coefficients would be null due to lack of information.

In this work we show that by focusing on the aggregate SFS of structured populations, it is possible to have a summary statistic that allows for complex models of population structure and is sensitive to the model parameters (see figures 3.6 and 3.5). We also show that by using phase-type theory, it is possible to exactly compute this SFS for moderate sample sizes very quickly. We did this by introducing a generalization of Herbots's structured coalescent that tracks information regarding the ancestry of the lineages as well as their location in the demes. The challenge of the approach is tied to its computational requirements, since the state space of the Markov process grows almost exponentially with the sample size. To combat this effect, we recommend the usage of model specializations. These are Markov processes that are derived from the general process described in §3.2 that include state-reduction optimizations specific to the demographic model in use. These state reductions are achieved by exploiting the symmetries often exhibited in structured models. We explore the n -island model specialization in §3.4. One obvious way to continue this line of research is to explore other model specializations. A great candidate would be the stepping-stone family of models, which would enable us to draw more spatial-aware conclusions from the potential demographic inference applications.

The performance results shown in figures 3.3 to 3.7 are focused on the n -island model specialization. We do not know yet how these techniques will perform under different models.

On one hand, the n -island is the most symmetrical of the structured models, therefore it allows for the greatest compression on the state space, so from this perspective, other models would be expected to require larger state spaces and thus more computational resources. On the other hand, the n -island model has the densest migration matrix of all structured models (note that in the stepping stone model for instance, the migration matrix itself is sparse). This means that other models will generate even sparser rate matrices. Given that the computational cost in sparse-algebra methods are associated with the number of non-null coefficients, this could counter the effect of the increased state space. It is hard to estimate how these two opposite effects will ultimately affect the performance of the methods. This is why more specializations have to be developed and tested.

4.4. Structured demographic inference using multi-sample statistics

One immediate use case for the methods presented here is to use the SFS as a summary statistic for parameter inference in structured models.

We have shown that, at least for the case of the n -island model, we can achieve good performance and accuracy when computing the SFS. This is in stark contrast with simulation-based estimates of the SFS, which are significantly slower for all tested sample sizes and less accurate (see Figure 3.9). This opens the opportunity for an approach similar to that of Chapter 2, where the summary statistic in use would be the SFS. There are two important questions that would need to be researched here: how to manage the sampling vector and how to choose the distance function.

The sampling vector is an important aspect of any structured demographic model that incorporates multiple samples. We have seen that it can have a profound impact on the SFS of otherwise identical n -island models (Figure 3.5), and such is the case as well for the $IICR_k$ (see the differences between $IICR_s$ and $IICR_d$ in Chikhi et al. (2018)).

One initial approach could be to assume that the sampling vector is known, although this defeats the purpose of an inferential approach. Alternatively, we could attempt to completely infer the sampling vector. This would add a great burden of dimensionality to the parameter space which could render the method impractical. Perhaps a compromise can be achieved between these two extreme approaches, where any existing information about the samples is incorporated in the model (for instance, lineages 1 and 2 are always from the same deme, as are lineages 3 and 4, etc.). In models that are spatially aware, such as the many variants of the stepping stones model, an alternative approach to managing the sampling vector is possible. We can assume that the actual (physical) sampling locations are known, given a map of the sampling region. We can then associate each sample with the nearest model deme corresponding to that region of the map. This approach allows the sampling vector to be constructed organically from the rest of the model parameters. See Figure 4.1 for an example.

Regarding the question of what distance to use for inference, this will depend on the exact nature of the method. For purely SFS based methods, log-likelihood metrics have been

successfully used in the past (Gutenkunst et al., 2009; Excoffier et al., 2013). However, a multi-sample demographic-inference method need not be SFS-only based. In cases where full-genome information is available for the sampled individuals, it is conceivable to combine the predictive power of both the IICR and the SFS summary statistics. This could be modeled either by using several IICR curves, belonging to each of the sampled individuals. The optimization goal in these scenarios is less clearly defined, as is often the case with multi-objective optimization problems. An initial approach could be based on a linear combination of the distances used for each of the summary statistics.




Figure 4.1: A method for parameterizing spatial structured models. A stepping-stone model is defined by overlaying a grid onto the map of the region of interest. The physical locations of the samples (the red dots in the figure) are assumed to be known, and each one is associated with the nearest deme (the association is represented by a red line connecting the sample with a deme). In this example model, it is possible to specify the number of islands, their connections, and the sampling vector by setting a single parameter: the grid spacing.

Another potential direction is to use multiple IICR curves simultaneously during the inference process. These multiple IICRs may come in the form of more than one IICR sampled from an asymmetrical demographic model (for which the initial sampling deme *does* result in different curves (Chikhi et al., 2018), as opposed to the n -island model where demes are by definition indistinguishable). They may also be in the form of multiple $IICR_k$ curves where k is the number of sampled haploid genomes. Indeed, the IICR of Mazet et al. (2016) was defined for $k = 2$, and this is the IICR that we have been studying in our previous works. However, the concept can be extended to more haploid genomes in the same way that the MSMC method (Schiffels and Durbin, 2013) is an extension of the PSMC to multiple genomes, which takes into consideration the distribution of the coalescent time T_k . The precise concept of the $IICR_k$ is currently being developed in a separate study.

These approaches may prove beneficial in choosing between structured and non-structured models. Indeed, Grusea et al. (2018) shows that using more than one IICR curve can help discriminate between structured and non structured scenarios in the n -island model. Finally,

the incorporation of larger samples not only enables exploring more complex scenarios, but it also allows using other summary statistics to complement the IICR, most notably among them the AFS, which is widely used for the purposes of demographic inference.

4.5. Closing thoughts

We have presented in this work an inference method for automatically estimating demographic parameters under a piecewise stationary symmetrical island model that uses the IICR as its summary statistic. The underlying methodology consists in quantifying the discrepancy between a target IICR and many simulated IICR curves for a large number of candidate scenarios, and using this metric to drive a global optimization process. With a large number of validations we have shown that the method works accurately and consistently for a diverse range of parameter values, and we additionally showed an application to human data that agrees with and improves upon previously published results using similar approaches. We believe that despite its current scope, our IICR-based method can be of great value during the initial exploration of the parameter space for simple models, and thus can also provide a starting point for manually fitting the IICR with models that could express spatial structure and varying N (Rodríguez et al., 2018).

With the goal of extending this method to multi-sample statistics, we have taken the first steps towards incorporating the SFS by providing a method for exactly computing it in the context of structured demographic models. The proposed method allows for an arbitrary number of demes and an arbitrary connectivity matrix between the demes. We also provide an implementation for the special case of the n -island model which was shown to be superior to simulation-based estimates of the SFS in both runtime (for the tested sample sizes) and accuracy. These developments should allow us to incorporate the SFS as a summary statistic in inference applications for demographic models of structure in which the number of islands is unknown and potentially large.

Bibliography

- Al-Mohy, A. H. and Higham, N. J. (2011). Computing the action of the matrix exponential, with an application to exponential integrators. *SIAM journal on scientific computing*, 33(2):488–511.
- Arredondo, A., Mourato, B., Nguyen, K., Boitard, S., Rodríguez, W., Noûs, C., Mazet, O., and Chikhi, L. (2021). Inferring number of populations and changes in connectivity under the n-island model. *Heredity*, 126(6):896–912.
- Beaumont, M. (2004). Recent developments in genetic data analysis: what can they tell us about human demographic history? *Heredity*, 92(5):365–379.
- Beaumont, M. A. (1999). Detecting population expansion and decline using microsatellites. *Genetics*, 153(4):2013–2029.
- Beichman, A. C., Phung, T. N., and Lohmueller, K. E. (2017). Comparison of single genome and allele frequency data reveals discordant demographic histories. *G3: Genes, Genomes, Genetics*, 7(11):3605–3620.
- Bladt, M. and Nielsen, B. F. (2017). *Matrix-exponential distributions in applied probability*, volume 81. Springer.
- Boitard, S., Rodríguez, W., Jay, F., Mona, S., and Austerlitz, F. (2016). Inferring population size history from large samples of genome-wide molecular data—an approximate bayesian computation approach. *PLoS Genetics*, 12(3):e1005877.
- Chikhi, L., Rodríguez, W., Grusea, S., Santos, P., Boitard, S., and Mazet, O. (2018). The IICR (inverse instantaneous coalescence rate) as a summary of genomic diversity: insights into demographic inference and model choice. *Heredity*, 120:13–24.
- de Barros Damgaard, P., Martiniano, R., Kamm, J., Moreno-Mayar, J. V., Kroonen, G., Peyrot, M., Barjamovic, G., Rasmussen, S., Zacho, C., Baimukhanov, N., et al. (2018). The first horse herders and the impact of early bronze age steppe expansions into asia. *Science*, 360(6396).
- Excoffier, L., Dupanloup, I., Huerta-Sánchez, E., Sousa, V. C., and Foll, M. (2013). Robust demographic inference from genomic and snp data. *PLoS Genetics*, 9(10):e1003905.
- Fay, J. C. and Wu, C.-I. (2000). Hitchhiking under positive darwinian selection. *Genetics*, 155(3):1405–1413.

- Fu, Y.-X. (1995). Statistical properties of segregating sites. *Theoretical population biology*, 48(2):172–197.
- Gillespie, J. H. (2004). *Population genetics: a concise guide*. JHU Press.
- Goldstein, D. B. and Chikhi, L. (2002). Human migrations and population structure: what we know and why it matters. *Annual Review of Genomics and Human Genetics*, 3(1):129–152.
- Griffiths, R. C. and Tavaré, S. (1998). The age of a mutation in a general coalescent tree. *Stochastic Models*, 14(1-2):273–295.
- Grusea, S., Rodríguez, W., Pinchon, D., Chikhi, L., Boitard, S., and Mazet, O. (2018). Coalescence times for three genes provide sufficient information to distinguish population structure from population size changes. *Journal of Mathematical Biology*, 78(1-2):189–224.
- Gutenkunst, R. N., Hernandez, R. D., Williamson, S. H., and Bustamante, C. D. (2009). Inferring the joint demographic history of multiple populations from multidimensional SNP frequency data. *PLoS Genetics*, 5(10):e1000695.
- Hackbusch, W. (1994). *Iterative solution of large sparse systems of equations*, volume 95. Springer.
- Hardy, G. H. and Ramanujan, S. (1918). Asymptotic formulæ in combinatory analysis. *Proceedings of the London Mathematical Society*, 2(1):75–115.
- Hein, J., Schierup, M., and Wiuf, C. (2004). *Gene genealogies, variation and evolution: a primer in coalescent theory*. Oxford University Press, USA.
- Herbots, H. M. J. D. (1994). *Stochastic models in population genetics: genealogy and genetic differentiation in structured populations*. PhD thesis.
- Hobolth, A., Siri-Jegousse, A., and Bladt, M. (2019). Phase-type distributions in population genetics. *Theoretical Population Biology*, 127:16–32.
- Hudson, R. R. (2004). ms a program for generating samples under neutral models.
- Hudson, R. R. (2015). A new proof of the expected frequency spectrum under the standard neutral model. *Plos one*, 10(7):e0118087.
- Johri, P., Riall, K., and Jensen, J. D. (2020). The impact of purifying and background selection on the inference of population history: problems and prospects. *BioRxiv*.
- Jouganous, J., Long, W., Ragsdale, A. P., and Gravel, S. (2017). Inferring the joint demographic history of multiple populations: beyond the diffusion approximation. *Genetics*, 206(3):1549–1567.

- Kamm, J., Terhorst, J., Durbin, R., and Song, Y. S. (2019). Efficiently inferring the demographic history of many populations with allele count data. *Journal of the American Statistical Association*, pages 1–16.
- Kaplan, E. L. and Meier, P. (1958). Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association*, 53(282):457–481.
- Kern, A. D. and Hey, J. (2017). Exact calculation of the joint allele frequency spectrum for isolation with migration models. *Genetics*, 207(1):241–253.
- Knuth, D. E. (2005). *The Art of Computer Programming, Volume 4, Fascicle 3: Generating All Combinations and Partitions*. Addison-Wesley Professional.
- L Graham, R. (1994). *Concrete mathematics: a foundation for computer science*. Addison-Wesley.
- Lapierre, M., Lambert, A., and Achaz, G. (2017). Accuracy of demographic inferences from the site frequency spectrum: the case of the yoruba population. *Genetics*, 206(1):439–449.
- Lehmann, D. J. and Smyth, M. B. (1981). Algebraic specification of data types: A synthetic approach. *Mathematical systems theory*, 14(1):97–139.
- Li, H. and Durbin, R. (2011). Inference of human population history from individual whole-genome sequences. *Nature*, 475(7357):493–496.
- Liu, X. and Fu, Y.-X. (2015). Exploring population size changes using SNP frequency spectra. *Nature Genetics*.
- Mazet, O., Rodríguez, W., and Chikhi, L. (2015). Demographic inference using genetic data from a single individual: Separating population size variation from population structure. *Theoretical Population Biology*, 104:46–58.
- Mazet, O., Rodríguez, W., Grusea, S., Boitard, S., and Chikhi, L. (2016). On the importance of being structured: instantaneous coalescence rates and human evolution—lessons for ancestral population size inference? *Heredity*, 116(4):362.
- Moler, C. and Van Loan, C. (2003). Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later. *SIAM review*, 45(1):3–49.
- Nielsen, R. (2000). Estimation of population parameters and recombination rates from single nucleotide polymorphisms. *Genetics*, 154(2):931–942.
- Noskova, E., Ulyantsev, V., Koepfli, K.-P., O'Brien, S. J., and Dobrynin, P. (2019). GADMA: Genetic algorithm for inferring demographic history of multiple populations from allele frequency spectrum data. *bioRxiv*.

- Prado-Martinez, J., Sudmant, P. H., Kidd, J. M., Li, H., Kelley, J. L., Lorente-Galdos, B., Veeramah, K. R., Woerner, A. E., O'Connor, T. D., Santpere, G., et al. (2013). Great ape genetic diversity and population history. *Nature*, 499(7459):471–475.
- Rodríguez, W., Mazet, O., Grusea, S., Arredondo, A., Corujo, J. M., Boitard, S., and Chikhi, L. (2018). The IICR and the non-stationary structured coalescent: towards demographic inference with arbitrary changes in population structure. *Heredity*, 121(6):663–678.
- Rodríguez, W., Mazet, O., Grusea, S., Arredondo, A., Corujo, J. M., Boitard, S., and Chikhi, L. (2021). The IICR and the non-stationary structured coalescent: towards demographic inference with arbitrary changes in population structure. (Correction to *Heredity* 2018, vol. 121(6), p. 663–678). *Heredity*, 126(3):xxx–xxx.
- Scally, A. and Durbin, R. (2012). Revising the human mutation rate: implications for understanding human evolution. *Nature Reviews Genetics*, 13(10):745–753.
- Scerri, E. M., Chikhi, L., and Thomas, M. G. (2019). Beyond multiregional and simple out-of-Africa models of human evolution. *Nature Ecology & Evolution*, 3(10):1370–1372.
- Scerri, E. M. L., Thomas, M. G., Manica, A., Gunz, P., Stock, J. T., Stringer, C., Grove, M., Groucutt, H. S., Timmermann, A., Rightmire, G. P., d'Errico, F., Tryon, C. A., Drake, N. A., Brooks, A. S., Dennell, R. W., Durbin, R., Henn, B. M., Lee-Thorp, J., deMenocal, P., Petraglia, M. D., Thompson, J. C., Scally, A., and Chikhi, L. (2018). Did our species evolve in subdivided populations across Africa, and why does it matter? *Trends in Ecology & Evolution*.
- Schiffels, S. and Durbin, R. (2013). Inferring human population size and separation history from multiple genome sequences. *Nature Genetics*, 8(46):919–925.
- Sherlock, C. (2021). Direct statistical inference for finite markov jump processes via the matrix exponential. *Computational Statistics*, pages 1–25.
- Siva, N. (2008). 1000 genomes project. *Nature biotechnology*, 26(3):256–257.
- Stanley, R. P. (2011). Enumerative combinatorics volume 1 second edition. *Cambridge studies in advanced mathematics*.
- Storn, R. and Price, K. (1997). Differential evolution—a simple and efficient heuristic for global optimization over continuous spaces. *Journal of Global Optimization*, 11(4):341–359.
- Tajima, F. (1983). Evolutionary relationship of dna sequences in finite populations. *Genetics*, 105(2):437–460.
- Tajima, F. (1989). Statistical method for testing the neutral mutation hypothesis by dna polymorphism. *Genetics*, 123(3):585–595.

- Tavaré, S. (2004). Part I: Ancestral inference in population genetics. In *Lectures on probability theory and statistics*, pages 1–188. Springer.
- Wakeley, J. (1999). Nonequilibrium migration in human history. *Genetics*, 153(4):1863–1871.
- Wakeley, J. and Hey, J. (1997). Estimating ancestral population parameters. *Genetics*, 145(3):847–855.
- Watterson, G. (1975). On the number of segregating sites in genetical models without recombination. *Theoretical population biology*, 7(2):256–276.
- Weissman, D. B. and Hallatschek, O. (2017). Minimal-assumption inference from population-genomic data. *Elife*, 6:e24836.
- Wright, S. (1931). Evolution in Mendelian populations. *Genetics*, 16(2):97.
- Wright, S. (1943). Isolation by distance. *Genetics*, 28(2):114.

Appendix A.

Additional results of validation using exact simulated IICRs

A.1. Unscaled IICR

In this section we show in figures A.1 to A.6 the validation results of simulating and then inferring from 400 randomly generated demographic scenarios with unscaled IICRs and varying number of components c . Each figure consists of as many sub-panels as there are free parameters for that model, and in each one the simulated values are in the horizontal axis and the inferred ones in the vertical axis.

In the lower-right corner of each sub-panel we display the normalized root-mean-square deviation (nRMSD) of the simulated versus inferred vectors. Additionally, we indicate in grey the region of 10% relative error (50% for the t_i parameters). The percentage of tests that fall within this and other margins of error is summarised in Figure A.7.




Figure A.1: Scatter plots of all the simulated (horizontal axis) versus inferred (vertical axis) parameter values for scenarios of $c = 1$ component.




Figure A.2: Scatter plots of all the simulated (horizontal axis) versus inferred (vertical axis) parameter values for scenarios of $c = 2$ components.




Figure A.3: Scatter plots of all the simulated (horizontal axis) versus inferred (vertical axis) parameter values for scenarios of $c = 3$ components.




Figure A.4: Scatter plots of all the simulated (horizontal axis) versus inferred (vertical axis) parameter values for scenarios of $c = 4$ components.




Figure A.5: Scatter plots of all the simulated (horizontal axis) versus inferred (vertical axis) parameter values for scenarios of $c = 5$ components.




Figure A.6: Scatter plots of all the simulated (horizontal axis) versus inferred (vertical axis) parameter values for scenarios of $c = 6$ components.




Figure A.7: Percent of tests within a given relative error in unscaled scenarios. For each parameter of the demographic scenarios, we counted the instances where the inferred value was within a certain percent of the simulated one. We say that an inferred parameter \bar{p} is within $x\%$ of the simulated value p if $-x\hat{p} < 100(\hat{p} - p) < xp$. Note that not all parameters are present in all scenarios (for instance, the parameter M_4 only appears in scenarios with $c \geq 5$ components).

A.2. Scaled IICR

In this section we show in figures A.8 to A.13 the validation results of simulating and then inferring from 400 randomly generated demographic scenarios with scaled IICRs and varying number of components c . Each figure consists of as many sub-panels as there are free parameters for that model, and in each one the simulated values are in the horizontal axis and the inferred ones in the vertical axis. The deme size panel (N) is different because the simulated values for N was always $N = 1000$, therefore the horizontal axis indicates the test number (1 to 400) and the vertical axis the inferred value.

In the lower-right corner of each sub-panel we display the normalized root-mean-square deviation (nRMSD) of the simulated versus inferred vectors. Additionally, we indicate in grey the region of 10% relative error (50% for the t_i parameters). The percentage of tests that fall within this and other margins of error is summarised in Figure A.14.




Figure A.8: Scatter plots of all the simulated (horizontal axis) versus inferred (vertical axis) parameter values for scenarios of $c = 1$ component.




Figure A.9: Scatter plots of all the simulated (horizontal axis) versus inferred (vertical axis) parameter values for scenarios of $c = 2$ components.

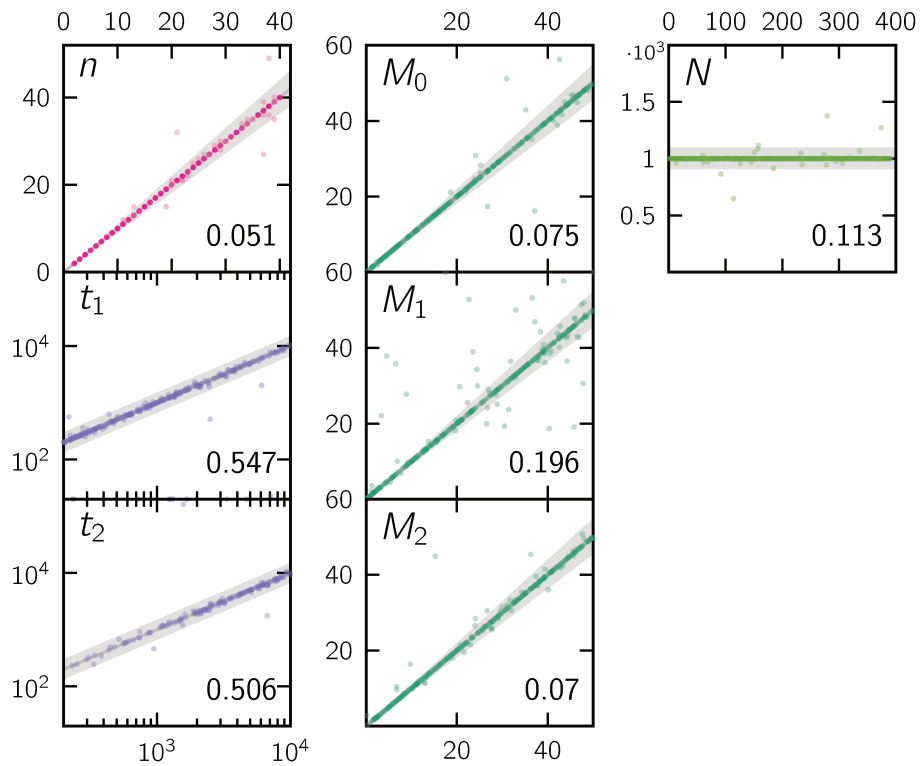


Figure A.10: Scatter plots of all the simulated (horizontal axis) versus inferred (vertical axis) parameter values for scenarios of $c = 3$ components.




Figure A.11: Scatter plots of all the simulated (horizontal axis) versus inferred (vertical axis) parameter values for scenarios of $c = 4$ components.




Figure A.12: Scatter plots of all the simulated (horizontal axis) versus inferred (vertical axis) parameter values for scenarios of $c = 5$ components.

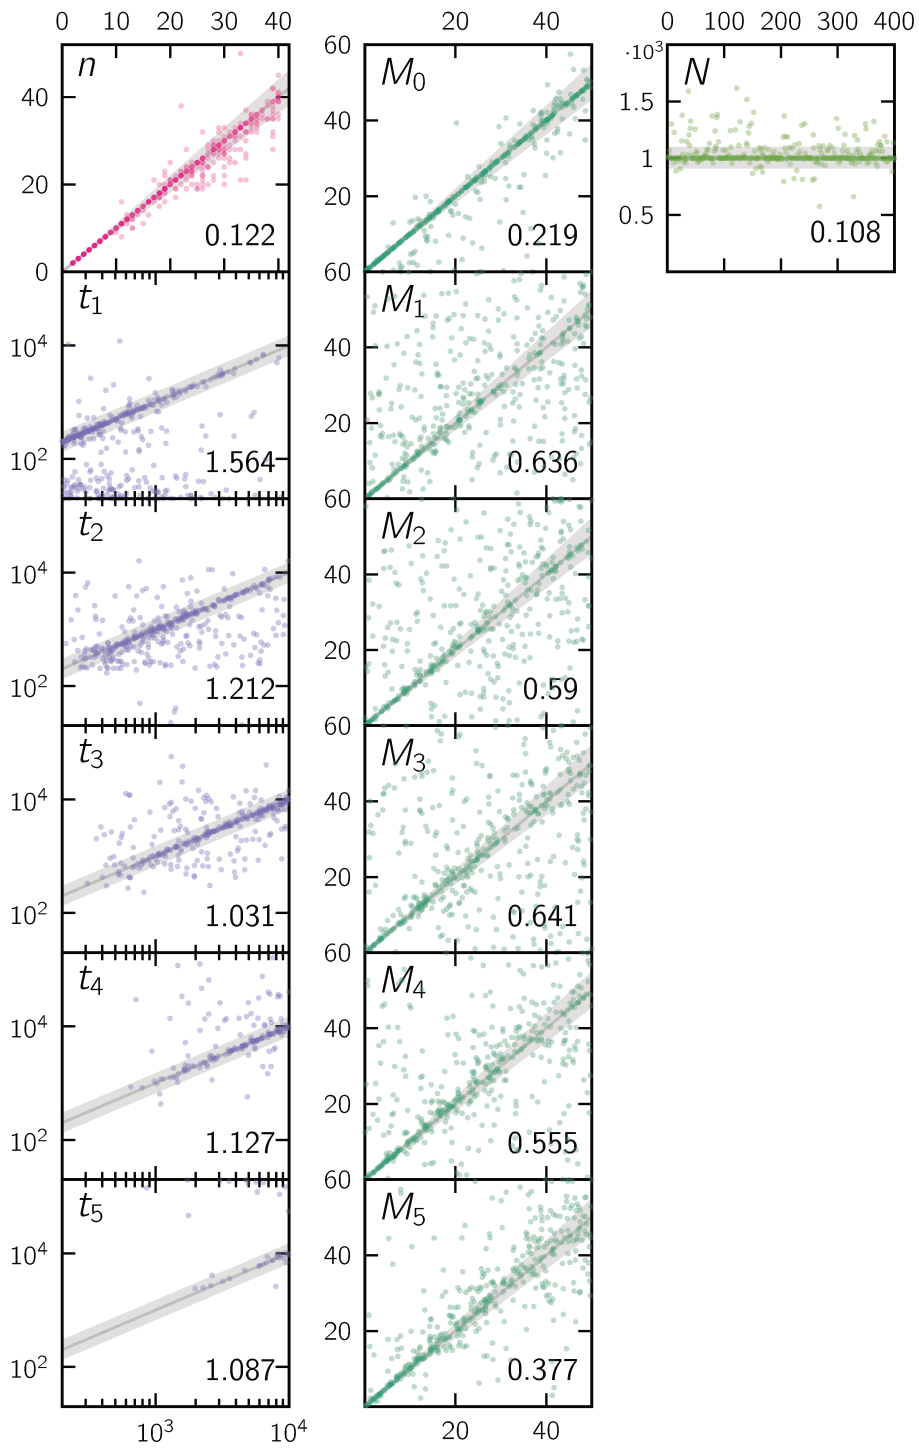


Figure A.13: Scatter plots of all the simulated (horizontal axis) versus inferred (vertical axis) parameter values for scenarios of $c = 6$ components.




Figure A.14: Percent of tests within a given relative error in scaled scenarios. For each parameter of the demographic scenarios, we counted the instances where the inferred value was within a certain percent of the simulated one. We say that an inferred parameter \bar{p} is within $x\%$ of the simulated value p if $-x\bar{p} < 100(\hat{p} - p) < x\bar{p}$. Note that not all parameters are present in all scenarios (for instance, the parameter M_4 only appears in scenarios with $c \geq 5$ components).

Appendix B.

Additional results of validation using T-sim IICRs

In this section we show the validation results of simulating and then inferring from 100 randomly generated demographic scenarios with scaled IICRs and varying number of components c . In all cases, the scenario parameters were drawn from the following finite sets:

$$\begin{aligned} n &\in \{2, 5, 10, 15, 20\}, \\ t_i &\in \{0.1, 0.5, 1, 2, 5, 10, 20, 50\} \quad \forall i, \\ M_i &\in \{0.1, 0.2, 0.5, 1, 2, 5, 10, 20, 50\} \quad \forall i, \\ s_i &= 1 \quad \forall i, \\ N &= 1000. \end{aligned} \tag{B.0.1}$$

Unlike in the previous sections, the simulated IICRs are T-sim IICRs, meaning that the values are not exact due to the stochastic nature of the underlying ms simulation. For each value of c from $c = 1$ to $c = 5$ we show the aggregate connectivity graph for all the simulations as well as the IICR and parameters of two individual scenarios from the set.

B.1. Scenarios with 1 component




Figure B.1: Connectivity graph of 100 inferred demographic histories using simulated T-sim IICRs of $c = 1$ components with simulated parameters randomly drawn from (B.0.1) and represented here by the dashed gray lines in the connectivity graph and the bold black circles in the islands and reference size plots.




Figure B.2: IICRs, connectivity graph, number of islands and reference size for one of the 100 simulated scenarios with $c = 1$ component and 10 independent inferences. The inferred scenario corresponds to $n = 5$, $M = 20$ and $N = 1000$.




Figure B.3: IICRs, connectivity graph, number of islands and reference size for one of the 100 simulated scenarios with $c = 1$ component and 10 independent inferences. The inferred scenario corresponds to $n = 10$, $M = 1$ and $N = 1000$.

B.2. Scenarios with 2 components




Figure B.4: Connectivity graph of 100 inferred demographic histories using simulated T-sim IICRs of $c = 2$ components with simulated parameters randomly drawn from (B.0.1) and represented here by the dashed gray lines in the connectivity graph and the bold black circles in the islands and reference size plots.




Figure B.5: IICRs, connectivity graph, number of islands and reference size for one of the 100 simulated scenarios with $c = 2$ components and 10 independent inferences. The inferred scenario corresponds to $n = 15$, $t = 1$, $M = (50, 5)$ and $N = 1000$.




Figure B.6: IICRs, connectivity graph, number of islands and reference size for one of the 100 simulated scenarios with $c = 2$ components and 10 independent inferences. The inferred scenario corresponds to $n = 15$, $t = 0.1$, $M = (50, 0.1)$ and $N = 1000$.

B.3. Scenarios with 3 components




Figure B.7: Connectivity graph of 100 inferred demographic histories using simulated T-sim IICRs of $c = 3$ components with simulated parameters randomly drawn from (B.0.1) and represented here by the dashed gray lines in the connectivity graph and the bold black circles in the islands and reference size plots.




Figure B.8: IICRs, connectivity graph, number of islands and reference size for one of the 100 simulated scenarios with $c = 3$ components and 10 independent inferences. The inferred scenario corresponds to $n = 20$, $t = (0.1, 5)$, $M = (50, 10, 2)$ and $N = 1000$.

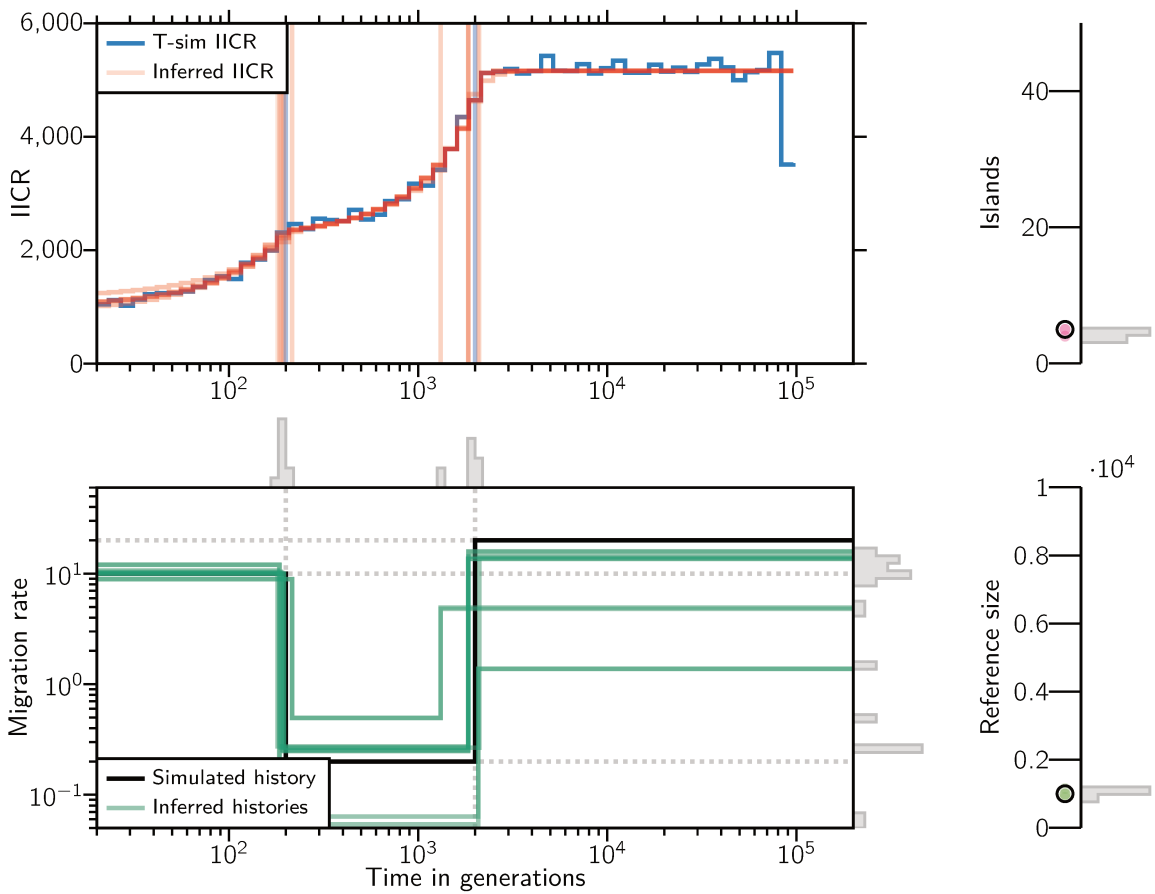


Figure B.9: IICRs, connectivity graph, number of islands and reference size for one of the 100 simulated scenarios with $c = 3$ components and 10 independent inferences. The inferred scenario corresponds to $n = 5$, $t = (0.1, 1)$, $M = (10, 0.2, 20)$ and $N = 1000$.

B.4. Scenarios with 4 components




Figure B.10: Connectivity graph of 100 inferred demographic histories using simulated T-sim IICRs of $c = 4$ components with simulated parameters randomly drawn from (B.0.1) and represented here by the dashed gray lines in the connectivity graph and the bold black circles in the islands and reference size plots.




Figure B.11: IICRs, connectivity graph, number of islands and reference size for one of the 100 simulated scenarios with $c = 4$ components and 10 independent inferences. The inferred scenario corresponds to $n = 5$, $t = (1, 10, 50)$, $M = (5, 50, 10, 0.2)$ and $N = 1000$.

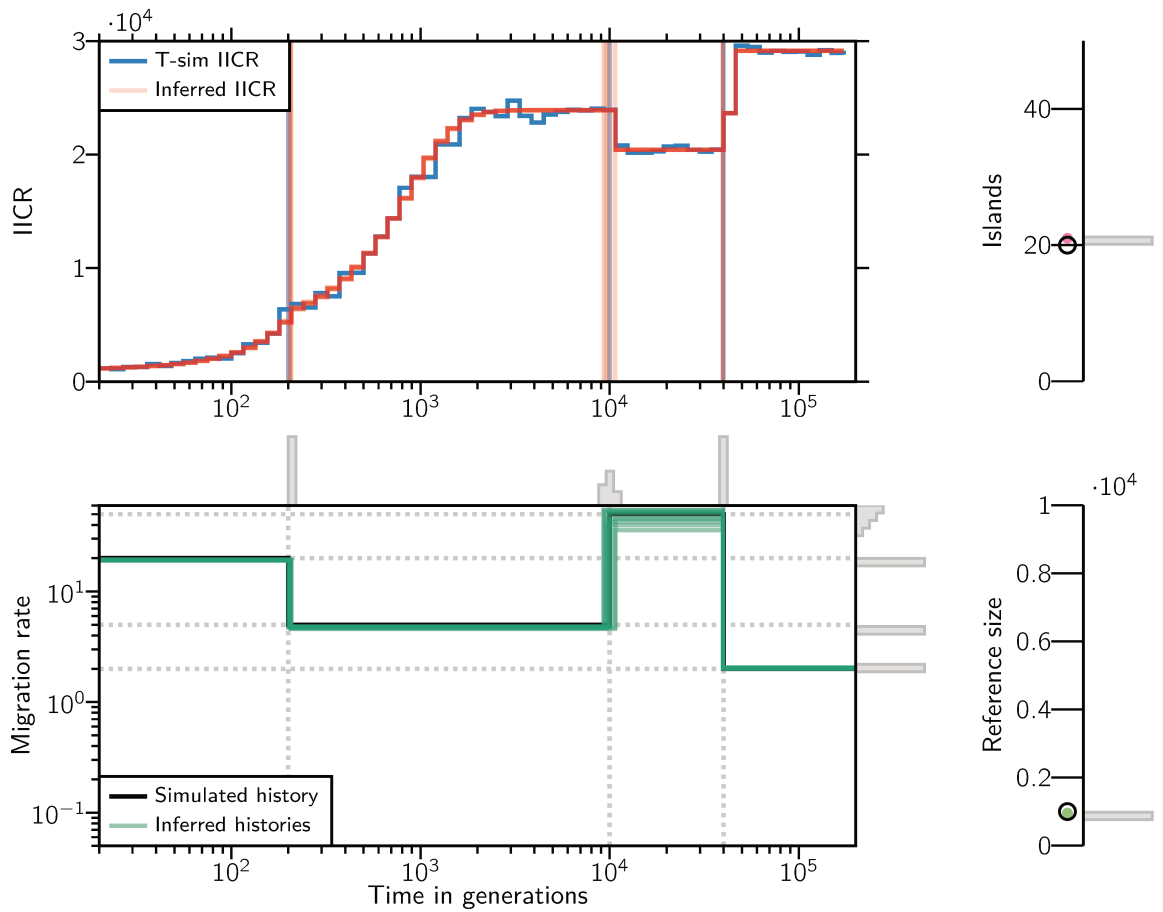


Figure B.12: IICRs, connectivity graph, number of islands and reference size for one of the 100 simulated scenarios with $c = 4$ components and 10 independent inferences. The inferred scenario corresponds to $n = 20$, $t = (0.1, 5, 20)$, $M = (20, 5, 50, 2)$ and $N = 1000$.

B.5. Scenarios with 5 components




Figure B.13: Connectivity graph of 100 inferred demographic histories using simulated T-sim IICRs of $c = 5$ components with simulated parameters randomly drawn from (B.0.1) and represented here by the dashed gray lines in the connectivity graph and the bold black circles in the islands and reference size plots.

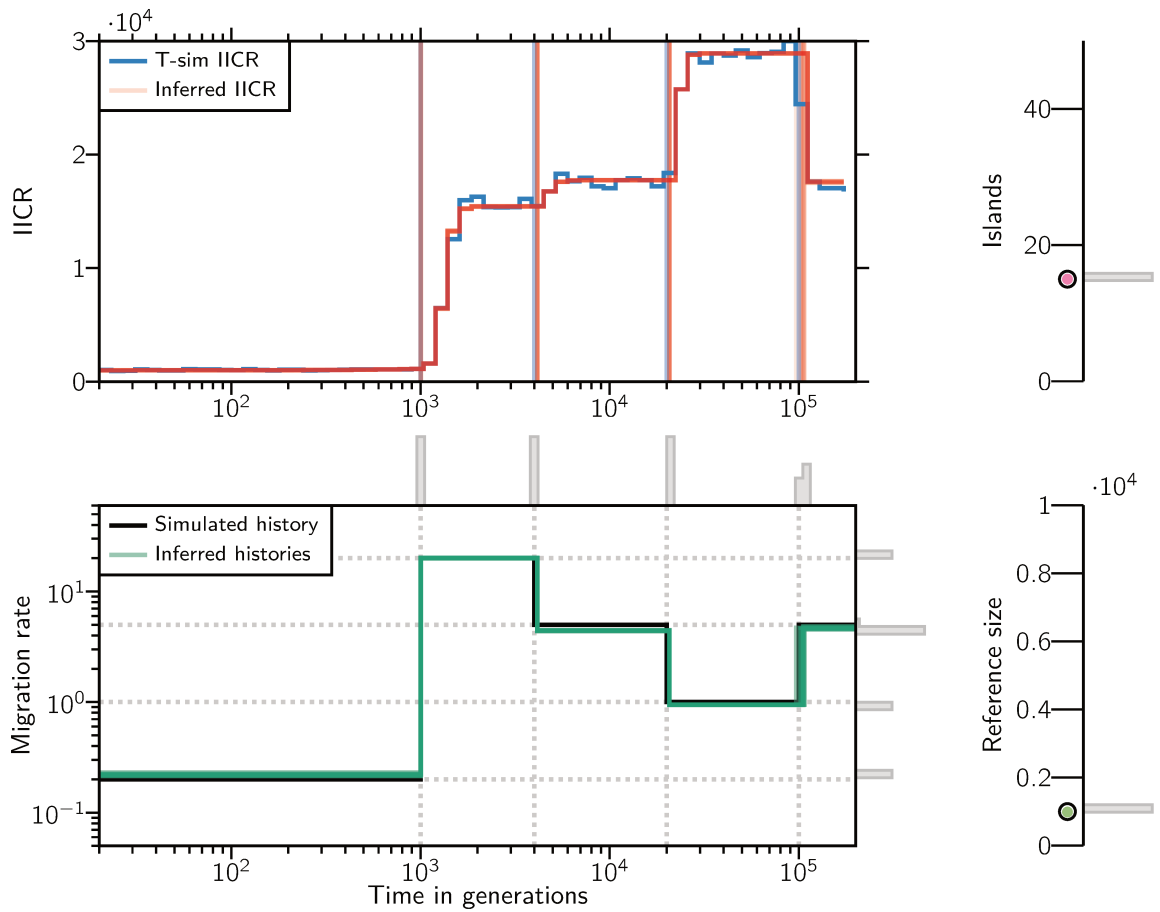


Figure B.14: IICRs, connectivity graph, number of islands and reference size for one of the 100 simulated scenarios with $c = 5$ components and 10 independent inferences. The inferred scenario corresponds to $n = 15$, $t = (0.5, 2, 10, 50)$, $M = (0.2, 20, 5, 1, 5)$ and $N = 1000$.




Figure B.15: IICRs, connectivity graph, number of islands and reference size for one of the 100 simulated scenarios with $c = 5$ components and 10 independent inferences. The inferred scenario corresponds to $n = 5$, $t = (0.1, 1, 5, 20)$, $M = (5, 0.5, 2, 50, 0.1)$ and $N = 1000$.