

# Estimating Large Correlation Matrices for International Migration

Jonathan J. Azose and Adrian E. Raftery  
University of Washington

Working Paper no. 154  
Center for Statistics and the Social Sciences  
University of Washington

May 26, 2016

## Abstract

The United Nations is the major organization producing and regularly updating probabilistic population projections for all countries. International migration is a critical component of such projections, and between-country correlations are important for forecasts of regional aggregates. However, there are 200 countries and only 12 data points, each one corresponding to a five-year time period. Thus a  $200 \times 200$  correlation matrix must be estimated on the basis of 12 data points. Using Pearson correlations produces many spurious correlations. We propose a maximum *a posteriori* estimator for the correlation matrix with an interpretable informative prior distribution. The prior serves to regularize the correlation matrix, shrinking *a priori* untrustworthy elements towards zero. Our estimated correlation structure improves projections of net migration for regional aggregates, producing narrower projections of migration for Africa as a whole and wider projections for Europe. A simulation study confirms that our estimator outperforms both the Pearson correlation matrix and a simple shrinkage estimator when estimating a sparse correlation matrix.

**Keywords:** Correlation, High-dimensional matrices, International Migration, World Population Prospects.

# Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
1.1	Illustrative example . . . . .	3
1.2	Background . . . . .	5
<b>2</b>	<b>Methods</b>	<b>8</b>
2.1	MAP correlation estimate . . . . .	9
2.1.1	Data Likelihood . . . . .	9
2.1.2	Prior . . . . .	10
2.1.3	Posterior . . . . .	11
2.2	Estimating $\tilde{R}$ . . . . .	11
2.3	Solving the minimization problem . . . . .	12
2.4	Selecting the regularization parameter $\lambda$ . . . . .	13
<b>3</b>	<b>Results</b>	<b>14</b>
3.1	Application to migration . . . . .	14
3.1.1	Data . . . . .	14
3.1.2	Selection of $P$ . . . . .	14
3.1.3	Selection of the regularization parameter, $\lambda$ . . . . .	16
3.1.4	Projection and evaluation . . . . .	16
3.2	Simulation study . . . . .	22
3.2.1	Simulation details . . . . .	22
3.2.2	Initial run to select $\lambda$ . . . . .	24
3.2.3	Evaluation of repeated estimation of $R$ . . . . .	24
<b>4</b>	<b>Discussion</b>	<b>25</b>
<b>A</b>	<b>Determining step size</b>	<b>32</b>
<b>B</b>	<b>Inflation of correlation estimates</b>	<b>33</b>

## List of Tables

1	Results of Kolmogorov-Smirnov test for correlations . . . . .	15
2	Continuous ranked probability score for continental migration projections . .	21
3	Evaluation of correlation matrix estimates from simulation study . . . . .	26

## List of Figures

1	Net migration rates for six countries . . . . .	4
2	Estimated correlations among forecast errors for migration . . . . .	5
3	Regularization criterion as a function of $\lambda$ . . . . .	16
4	Elements of the correlation matrix before and after regularization . . . . .	17
5	Probabilistic projections of net migration for continents . . . . .	18
6	Probabilistic projections of net migration for regional aggregates . . . . .	20
7	Simulation study: Elements of the correlation matrix before and after regularization . . . . .	23
8	Simulation study: Regularization criterion . . . . .	24

# 1 Introduction

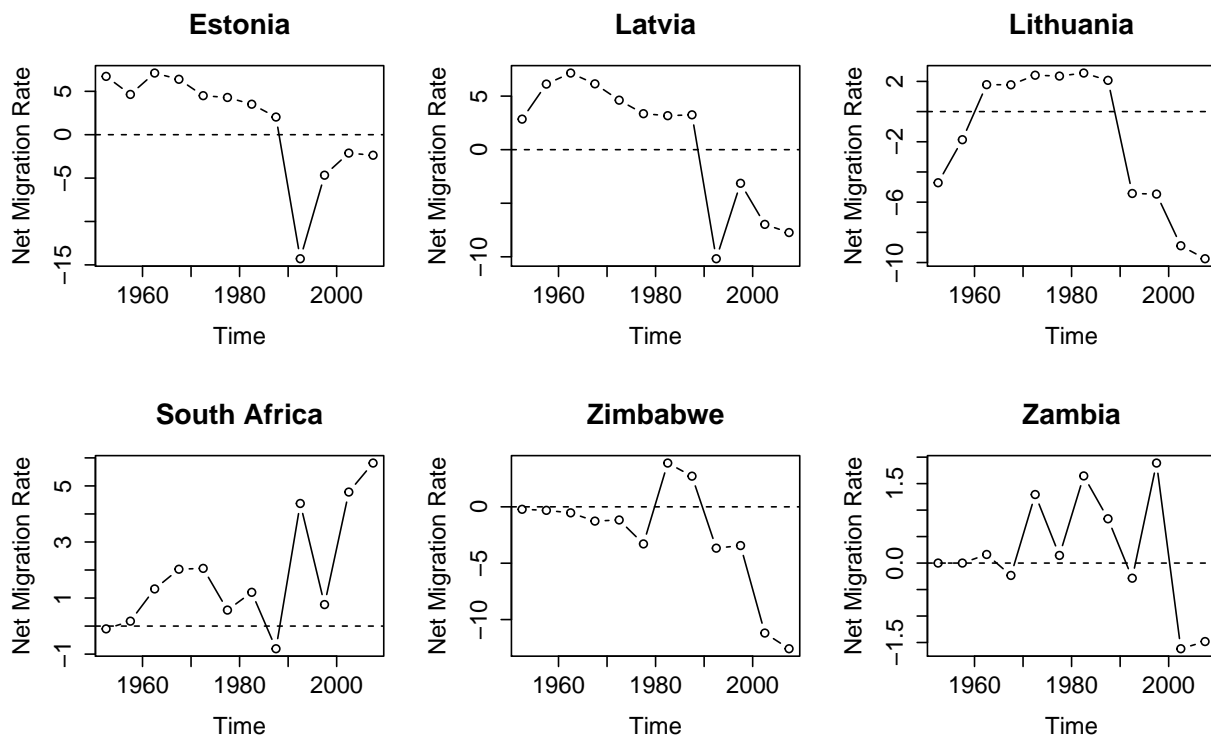
International migration is a major contributor to population change, but is hard to project, making proper quantification of uncertainty especially important. Existing global models for migration are well-calibrated marginally, i.e. for individual countries (Azose & Raftery, 2015), but typically rely on an unrealistic modeling assumption that forecast errors are uncorrelated across countries. If correlations exist, but are not modeled, the resulting projections may still be well calibrated for countries individually, but can under- or overestimate variance in projections of migration for regions that span multiple countries. We present a method for estimating a correlation matrix from a small number of data points that uses informative priors, shrinking elements of the correlation matrix which we expect *a priori* to be small. In applying this method to migration, we choose priors based on empirical evidence of non-zero correlations among classes of countries which are “close” to one another according to a variety of distance covariates. Our method improves projections of migration for regional aggregates while mitigating the issue of spurious correlations that arises from trying to estimate a large correlation matrix based on many short time series.

## 1.1 Illustrative example


In this section we focus on six selected countries—Estonia, Latvia, Lithuania, South Africa, Zimbabwe, and Zambia—to highlight the need for regularization of the correlation matrix.

Migration rates in Estonia, Latvia, and Lithuania over the period from 1950 to 2010 look quite similar (top row of Figure 1.) All three countries share a spike in out-migration during the 1990–1995 time period, which appears as a large negative forecast error in a first-order autoregressive (AR(1)) model. This sudden jump in out-migration among the Baltic states shares a common cause, namely the fall of the Soviet Union, which both induced westward migration and prompted many ethnic Russians to return to Russia (Fassmann & Munz, 1994; Okólski, 1998)

Meanwhile, several countries in Southern Africa also experienced big shifts in migration rates during the 1990–1995 time period (bottom row of Figure 1.) From 1990 to 1995, South Africa received substantially more in-migration than it had in previous decades, while Zimbabwe and Zambia both switched from being net receivers of migrants to net senders. For these three countries, at least some of the change in migration was due to political shifts related to the end of South Africa’s apartheid policy. For example, the number of legal entrants to South Africa who overstayed their visas grew dramatically during the 1990s, with many such entrants coming from other countries of the Southern African Development Community (Crush, 1999).



**Figure 1:** Net migration rates (net annual migrants per thousand individuals) for six countries.



**Figure 2:** Estimated correlations among forecast errors for migration. Left panel shows Pearson correlation estimates. Right panel shows our regularized estimates.

Because all six countries experienced pronounced changes in migration rates during the same time period, the usual Pearson estimates of the correlation in forecast errors are relatively large for these six countries (left panel of Figure 2.) Knowledge of world affairs, however, suggests that some of these correlations may be spurious. There are plausible explanations for the correlations within the three Baltic nations and within the Southern African nations, but the cross-regional correlations are suspect. In fact, the cross-regional correlations seem to have arisen largely from a coincidental synchrony in the timing of disparate geopolitical events, and do not represent correlations that we would expect to continue to exist in future migration data. Our method is designed to shrink these seemingly spurious cross-regional correlations, producing the estimated correlation matrix shown in the right panel of Figure 2. Cross-regional correlations decrease substantially in magnitude, while correlations within regions remain largely unchanged.

## 1.2 Background

Country-specific projections of international migration are an important input in policy-making decisions (Bijak, Kupiszewska, Kupiszewski, Saczuk, & Kicingier, 2007; Brown & Bean, 2012). Projected migration figures are commonly used in long-term planning of social welfare programs (U.S. Social Security Administration, 2013; Wright, 2010). However, projection of migration is difficult—Bijak and Wiśniowski (2010) describe migration as “barely

predictable”—and global modeling of migration remains somewhat rudimentary. The United Nations Population Division produces global projections of fertility, mortality, and migration for all countries (United Nations, 2012). For most countries, the 2012 revision of the World Population Prospects (WPP) deterministically projects net migration to persist at current levels until 2050 and decline linearly thereafter.

To produce fully probabilistic population projections, one must incorporate probabilistic projections of fertility and mortality with a global probabilistic model of migration. It follows from the demographic balancing equation that the contribution of migration to population change is given by *net* migration (that is, in-migration minus out-migration.) Probabilistic models exist for both net migration (Azose & Raftery, 2015; Azose, Ševčíková, & Raftery, 2016) and in- and out-migration separately (Wiśniowski, Smith, Bijak, Raymer, & Forster, 2015). Both of these models are Bayesian hierarchical autoregressive models which treat forecast errors in migration as independent across countries, conditional on model parameters. This leads to projections that are well calibrated for individual countries, but may not be for multi-country aggregates. Our method aims to relax this independence assumption.

It is worth noting that a strong correlation in migration rates themselves need not translate to a strong correlation in forecast errors. For example, from 1960 through 2000, Mexico was consistently either the largest or second-largest source of migration flows to the US, with nearly 5 million Mexicans migrating to the US during the 1990’s (Abel, 2013). While we estimate that net migration rates for the USA and Mexico have a correlation of -0.56 based on quinquennial WPP data from 1950-2010, we estimate a correlation in forecast errors of only -0.07. That is, most of the relationship between the USA and Mexico is already captured by the autoregressive model parameters, and the “random” components of migration rates for the two countries are nearly independent conditional on the AR(1) model.

In this high-dimensional setting with short time series, the empirical correlation matrix is a poor estimator, in that it can include many spuriously large estimated correlations. Our goal is to use regularization to improve an empirical correlation matrix for forecast errors in migration. There is a large body of literature on regularized estimation of covariance matrices, with applications in genomics, image processing, and finance, among other fields (Fan, Han, & Liu, 2014). The novelty of our method is that it allows the incorporation of available prior information in an easily interpretable way.

Existing covariance estimators based on penalized likelihood maximization are typically maximum *a posteriori* (MAP) estimates under some prior belief about covariance, but these formulations are not well suited to specifying beliefs directly about elements of the correlation matrix. Perhaps the most similar method to ours is that of Bien and Tibshirani (2011), which allows informative priors on elements of the covariance matrix rather than the correlation



matrix. Their method is not directly applicable to our setting, as our goal is to augment existing marginal variances with a suitable correlation structure. Other proposed MAP estimators include the graphical lasso (Friedman, Hastie, & Tibshirani, 2008), which can be used to place an informative prior on the inverse covariance, and the method of Chi and Lange (2014), which penalizes covariance estimates that have very large or very small eigenvalues. An extreme example is given by Chaudhuri, Drton, and Richardson (2007), who provide a method for covariance estimation in the presence of known zeroes. Zhang and Zou (2012) propose a variant on penalized likelihood maximization that replaces the negative log-likelihood with a simpler loss function.

A related class of covariance estimators relies on shrinkage of an empirical covariance matrix towards a simpler estimator, typically trading some bias for lower mean squared error (Ledoit & Wolf, 2003, 2004, 2012). A strength of these methods is that so long as the empirical covariance matrix is positive semi-definite and the shrinkage target is positive definite, a linear combination of the two will naturally be positive definite. Applying a shrinkage method to the migration setting would be difficult, as the elements we would like to penalize do not define a positive definite shrinkage target.

A form of regularization that is straightforward to implement is applying thresholding directly to elements of a covariance or correlation matrix (Bickel & Levina, 2008a; El Karoui, 2008); these authors show that a hard-thresholded covariance matrix is consistent in operator norm. Generalized thresholding (Antoniadis & Fan, 2001), developed in the context of wavelet applications, provides a class of related regularized estimators. A key difficulty with such estimators is that care must be taken to ensure that the resulting estimator is positive definite. In some problems, this can be handled by selecting a thresholding constant from an appropriate range (Fan, Liao, & Mincheva, 2013). Unfortunately, such an approach is not easily adapted to our problem. The structure of the elements we wish to penalize is such that we can tolerate only a small amount of shrinkage of all penalized elements before our estimated correlation matrix loses positive definiteness.

One fully Bayesian treatment is proposed by Liechty, Liechty, and Müller (2004), who include substantive prior information by specifying clusters of correlations which they expect to be similar. This is unfortunately unsuitable to our setting, since geographical and cultural proximity can give rise to either positive or negative correlations. A. Huang, Wand, et al. (2013) describe a computationally attractive *non-informative* prior on covariances which does not easily extend to the informative priors we would like to include. Other fully Bayesian treatments are given by Barnard, McCulloch, and Meng (2000), who propose a prior on the correlation matrix which is either marginally or jointly uniform, and Leonard and Hsu (1992) and Deng and Tsui (2013), who propose Bayesian estimation of the logarithm of the

covariance matrix, which is unfortunately hard to interpret.

In scenarios where there is a natural ordering to the variables, it is often reasonable to make the assumption that large values of  $|i - j|$  imply near independence or conditional independence. When this is the case, one can regularize by banding or tapering of the covariance or inverse covariance matrix (Bickel & Levina, 2008b; Chen, Xu, Wu, et al., 2013; Fan, Huang, & Li, 2007; Furrer & Bengtsson, 2007; Levina, Rothman, Zhu, et al., 2008). These approaches are not suitable to our problem, as there is no natural ordering of countries.

Good overviews of other methods in covariance estimation are given by Fan, Liao, and Liu (2015) and Pourahmadi (2011).

## 2 Methods

We start with an established, well-calibrated autoregressive model on net migration rates for all countries (Azose & Raftery, 2015). This model has the form:

$$\mathbf{g}_t - \boldsymbol{\mu} = \text{diag}(\boldsymbol{\phi})(\mathbf{g}_{t-1} - \boldsymbol{\mu}) + \boldsymbol{\varepsilon}_t, \quad (1)$$

$$\boldsymbol{\varepsilon}_t \stackrel{iid}{\sim} \mathcal{N}_C(\mathbf{0}, \text{diag}(\boldsymbol{\sigma}) \cdot I_C \cdot \text{diag}(\boldsymbol{\sigma})), \quad (2)$$

$$\phi_c \stackrel{iid}{\sim} U(0, 1), \quad (3)$$

$$\mu_c \stackrel{iid}{\sim} N(\lambda, \tau^2), \quad (4)$$

$$\sigma_c^2 \stackrel{iid}{\sim} IG(a, b). \quad (5)$$

Notationally,  $\mathbf{g}_t$  is a length- $C$  vector of net migration rates for all countries during the time period from  $t$  to  $t + 1$ , where  $C$  is the number of countries analyzed. The quantities  $\boldsymbol{\mu}$ ,  $\boldsymbol{\phi}$ , and  $\boldsymbol{\sigma}$  are vectors of model parameters, and  $\mathbf{0}$  is a length- $C$  vector of zeroes. (We have omitted here the specifics of hyperpriors on  $a$ ,  $b$ ,  $\lambda$ , and  $\tau$ , which Azose and Raftery selected to reflect the ranges of plausible values.) Notably, forecast errors in their model are treated as independent, conditional on the model’s other parameters. Our method augments this model with an estimated correlation structure. Although this paper focuses on the migration context, the same technique could be applied to probabilistic models of other demographic indicators.

From this point forward, we refer to Azose and Raftery’s model as the Bayesian Hierarchical Model with Independent Forecast Errors (BHM+IFE). In principle, the methodology we describe here provides a means of estimating a correlation matrix to be adjoined to any probabilistic model with conditionally independent forecast errors.

The outline of our procedure for estimating a correlation matrix is as follows:

1. From the BHM+IFE model, draw a posterior sample of  $m$  realizations of model parameters,  $\boldsymbol{\mu}^{(1)}, \boldsymbol{\phi}^{(1)}, \boldsymbol{\sigma}^{(1)}, \dots, \boldsymbol{\mu}^{(m)}, \boldsymbol{\phi}^{(m)}, \boldsymbol{\sigma}^{(m)}$ .
2. Convert the estimated forecast errors from the posterior sample of model parameters to a single empirical correlation matrix,  $\tilde{R}$ .
3. Combine the empirical correlation matrix with informative priors on correlations to obtain a maximum *a posteriori* (MAP) correlation estimate,  $\hat{R}$ .

This procedure can be viewed as performing a single step of the Monte Carlo EM (MCEM) algorithm (Wei & Tanner, 1990).

The posterior sampling in stage 1 can be performed using any reasonable sampling procedure. In practice, we performed our posterior sampling with a combination of Gibbs sampling and Metropolis-Hastings steps.

In the following sections, we first discuss the details of obtaining an MAP estimator (Section 2.1) and then the question of what to use for an empirical correlation matrix (Section 2.2). This is followed by an algorithm for computing the MAP estimator (Section 2.3), and finally discussion of a criterion for selecting a regularization parameter (Section 2.4).

## 2.1 MAP correlation estimate

Our goal is to estimate the correlation structure,  $R$ , of forecast errors,  $\boldsymbol{\varepsilon}_t$ . We assume a model of the form

$$\boldsymbol{\varepsilon}_t \stackrel{iid}{\sim} \mathcal{N}_C(\mathbf{0}, \Sigma), \quad (6)$$

where the variance matrix,  $\Sigma$ , decomposes into standard deviations,  $\boldsymbol{\sigma}$ , and a correlation matrix,  $R$ , as  $\Sigma = \text{diag}(\boldsymbol{\sigma}) \cdot R \cdot \text{diag}(\boldsymbol{\sigma})$ . To determine a MAP estimator for  $R$ , we express the posterior distribution for  $R$  as a product of likelihood and prior.

### 2.1.1 Data Likelihood

Equation (6) implies a likelihood function for  $R$  of the form

$$p(\boldsymbol{\varepsilon}_1, \dots, \boldsymbol{\varepsilon}_{T-1} | R, \boldsymbol{\sigma}) \propto_R \det(R)^{-(T-1)/2} \exp\left(-\frac{1}{2} \sum_{t=1}^{T-1} \boldsymbol{\varepsilon}_t' \text{diag}(\boldsymbol{\sigma})^{-1} R^{-1} \text{diag}(\boldsymbol{\sigma})^{-1} \boldsymbol{\varepsilon}_t\right), \quad (7)$$

restricted to the space  $\Omega$  of valid correlation matrices (i.e. positive semi-definite matrices with ones on the diagonal.) Matrix trace identities simplify this likelihood to

$$p(\boldsymbol{\varepsilon}_1, \dots, \boldsymbol{\varepsilon}_{T-1} | R, \boldsymbol{\sigma}) \propto_R \det(R)^{-(T-1)/2} \exp\left(-\frac{1}{2} \text{tr}(R^{-1} \tilde{R})\right), \quad (8)$$

where

$$\tilde{R} := \frac{1}{T-1} \sum_{t=1}^{T-1} \text{diag}(\boldsymbol{\sigma})^{-1} \boldsymbol{\varepsilon}_t \boldsymbol{\varepsilon}_t' \text{diag}(\boldsymbol{\sigma})^{-1}. \quad (9)$$

The evidence from the data is encapsulated in  $\tilde{R}$ , which is something akin to an empirical correlation matrix. Note that  $\tilde{R}$  would be a sufficient statistic for  $R$  if the  $\boldsymbol{\varepsilon}_t$ 's and  $\boldsymbol{\sigma}$  were known. In fact neither the  $\boldsymbol{\varepsilon}_t$ 's nor  $\boldsymbol{\sigma}$  are known, and  $\tilde{R}$  must be replaced with a sensible estimate in order to proceed. Details of the estimation of  $\tilde{R}$  are given in Section 2.2.

### 2.1.2 Prior

Our choice of prior distribution on  $R$  is motivated by a desire to incorporate informative prior beliefs about which country pairs are likely to be nearly uncorrelated. As such, we choose a prior of the form

$$\pi(R) \propto_R \prod_{0 \leq i < j \leq C} \exp(-\lambda P_{ij} |R_{ij}|), \quad (10)$$

again restricted to  $\Omega$ . The matrix  $P$  with entries  $P_{ij}$  is a penalty matrix that encodes the extent to which we believe that countries  $i$  and  $j$  may be correlated. In our application to migration, we constrain all the entries in  $P$  to be equal to 0 or 1, although in general  $P$  may be allowed to have arbitrary non-negative entries. The parameter  $\lambda$  is an overall regularization parameter that encodes how strongly we want to penalize correlations.

The key benefit of this prior is its ease of interpretability. Setting  $P_{ij} = 1$  expresses a belief that  $R_{ij}$  should be close to zero, with the strength of that belief controlled by  $\lambda$ . Setting  $P_{ij} = 0$  implies that all values of  $R_{ij}$  are equally believable, *a priori*. Other penalized likelihood estimators have been proposed, corresponding to MAP estimators under implied priors on precision (Friedman et al., 2008), covariance (Bien & Tibshirani, 2011), or eigenvalues of the covariance matrix (Chi & Lange, 2014). None of these allow one to specify prior beliefs about correlations directly.

Note that under this specification, the prior distribution of the correlation  $R_{ij}$  is either uniform or truncated Laplace conditional on the rest of the correlation matrix, but marginal distributions will not be uniform or double exponential. Although it is possible to specify a marginally uniform prior on all elements of the correlation matrix (Barnard et al., 2000), we know of no way to specify a distribution that is marginally uniform for some elements and marginally peaked at zero for others.

Because the prior density is a product of Laplace densities on correlations, we will refer to our eventual correlation estimator as the LPoC (Laplace Prior on Correlations) estimator. Augmenting the BHM+IFE with the LPoC correlation estimate produces the BHM+LPoC model.

### 2.1.3 Posterior

Combining the likelihood and prior, we obtain the log posterior distribution for  $R$ , equal to

$$\log p(R|\boldsymbol{\varepsilon}_1, \dots, \boldsymbol{\varepsilon}_{T-1}, \boldsymbol{\sigma}) = -\frac{T-1}{2} \log \det(R) - \frac{T-1}{2} \text{tr}(R^{-1}\tilde{R}) - \frac{\lambda}{2} \|P * R\|_1 + c(\boldsymbol{\varepsilon}_1, \dots, \boldsymbol{\varepsilon}_{T-1}, \boldsymbol{\sigma}), \quad (11)$$

where  $*$  denotes elementwise matrix multiplication, and  $\|\cdot\|_1$  gives the sum of the absolute value of the elements of a matrix.

Thus, finding the MAP estimator for  $R$  is equivalent to solving the minimization problem

$$\text{Minimize}_{R \in \Omega} \left\{ \log \det(R) + \text{tr}(R^{-1}\tilde{R}) + \frac{1}{T-1} \lambda \cdot \|P * R\|_1 \right\} \quad (12)$$

Algorithmic details of a numerical solution are given in Section 2.3.

Note that if the penalty parameter,  $\lambda$ , is zero, then this minimization problem yields the maximum likelihood estimator (MLE) of  $R$  conditional on  $\boldsymbol{\sigma}$ . So long as  $\tilde{R}$  is itself positive definite, this MLE is just  $\tilde{R}$ , the empirical correlation matrix. Similarly, if  $\lambda$  is held fixed as  $T$  grows, the penalty term in (12) goes to zero and the LPoC estimator converges to  $\tilde{R}$ . Since  $\tilde{R}$  is consistent for  $R$ , the LPoC estimator is also consistent.

## 2.2 Estimating $\tilde{R}$

Since the forecast errors and model parameters of the BHM+IFE model are unknown, we do not have access to the true value of  $\tilde{R}$ . Instead we use an estimate of  $\tilde{R}$ . For practical reasons, we would prefer to have  $\tilde{R}$  itself be a valid correlation matrix so that (12) will have a known analytic solution in the limiting scenarios where  $T$  grows or  $\lambda$  goes to zero. Accordingly, we might choose an estimator  $\tilde{R}^{basic}$  with elements defined by

$$\tilde{R}_{ij}^{basic} := \frac{\sum_{t=1}^{T-1} \hat{\varepsilon}_{i,t} \hat{\varepsilon}_{j,t}}{\sqrt{\sum_{t=1}^{T-1} \hat{\varepsilon}_{i,t}^2} \sqrt{\sum_{t=1}^{T-1} \hat{\varepsilon}_{j,t}^2}}, \quad (13)$$

where  $\hat{\varepsilon}_t$  is the posterior mean of  $\boldsymbol{\varepsilon}_t$  from the BHM+IFE model. This estimate,  $\tilde{R}^{basic}$ , is the MLE for estimating the correlation matrix of a multivariate normal random variable with mean known to be zero and unknown marginal variance terms. By construction,  $\tilde{R}^{basic}$  is guaranteed to be positive semi-definite and to have ones on the diagonal.

However, in our application,  $\tilde{R}^{basic}$  is low rank, since  $T$  is small relative to the dimension of the matrix. For computational reasons, we would prefer to have a strictly positive definite matrix, so we estimate  $\tilde{R}$  by

$$\tilde{R}^{PD} = 0.99 \cdot \tilde{R}^{basic} + 0.01 \cdot I_C. \quad (14)$$

This change can be viewed as augmenting our estimates of  $\boldsymbol{\varepsilon}_t$  with a small amount of additional uncorrelated data.

## 2.3 Solving the minimization problem

We apply a majorize-minimize algorithm similar to that used by [Bien and Tibshirani \(2011\)](#) to the minimization problem in (12). The function being minimized over is the sum of a convex and a concave component. The majorize-minimize algorithm repeatedly iterates through the following steps:

1. Replace the concave component with its tangent plane to obtain a fully convex function.
2. Find the global minimum of the convex function from Step 1.
3. Update the estimate of the tangent plane.

Notationally, we label our starting point for this algorithm as  $R_0$  and subsequent iterations of this majorize-minimize algorithm are denoted with subscripts  $R_1, R_2, \dots$

In (12), the concave component is  $\log \det(R)$ , which we replace with the tangent plane  $\log \det R_i + \text{tr}(R_i^{-1}(R - R_i))$ . After simplifying and removing terms which are constant in  $R$ , the convex minimization problem in the  $i$ th iteration of the algorithm is

$$\text{Minimize}_{R \in \Omega} \left\{ \text{tr}(R_i^{-1}R) + \text{tr}(R^{-1}\tilde{R}) + \lambda \|P * R\|_1 \right\}. \quad (15)$$

Now all of the terms the objective function in (15) are convex, and all but  $\lambda \|P * R\|_1$  are differentiable, so we can apply the generalized gradient descent algorithm ([Beck & Teboulle, 2009](#)). Each generalized gradient descent step takes the form

$$R_{new} = \text{argmin}_{\omega \in \Omega} \left\{ (2t)^{-1} \|\omega - (R_{current} - t(R_i^{-1} - R_{current}^{-1}\tilde{R}R_{current}^{-1}))\|_F^2 + \frac{\lambda}{T-1} \|P * \omega\|_1 \right\}. \quad (16)$$

If the restriction to  $\Omega$  were not present, this problem would have a simple analytic solution, given by

$$R_{new} = \mathcal{S} \left( R_{current} - t(R_i^{-1} - R_{current}^{-1}\tilde{R}R_{current}^{-1}), \frac{\lambda}{T-1} tP \right), \quad (17)$$

where  $\mathcal{S}$  is the element-wise soft-thresholding operator defined by

$$\mathcal{S}(X, \alpha)_{ij} = \text{sign}(X_{ij}) \cdot (|X_{ij}| - \alpha_{ij}) \cdot \mathbb{1}(|X_{ij}| > \alpha_{ij}). \quad (18)$$

(This move is actually restricted to the off-diagonal elements only, as the diagonal elements of a correlation matrix are constrained to equal 1.) Thus, if there were no positive definiteness constraint, each update step would consist of a gradient descent step according to the gradient of the differentiable component followed by soft-thresholding the result.

Although we do have to satisfy a positive definiteness constraint, we can start by trying the update step in (17). If this update results in a valid correlation matrix, then that matrix

is our solution to (16), and we replace  $R_{current}$  with  $R_{new}$ . However, sometimes the soft-thresholded gradient step results in a matrix that is not positive definite. In that case, it is possible to appeal to a slower, iterative solution to (16). One such solution is given by Cui, Leng, and Sun (2016). In practice, as long as we are looking for a solution in the interior of  $\Omega$ , it is good enough to simply reduce step size rather than appealing to the relatively costly iterative algorithm whenever the generalized gradient descent suggestion lies outside of  $\Omega$ .

Step size selection has a large impact on performance and convergence of this algorithm. Details of step size selection are discussed in Appendix A.

## 2.4 Selecting the regularization parameter $\lambda$

Although the penalty matrix  $P$  can be selected on the basis of world knowledge, we are less likely to have genuine prior beliefs about the value of the regularization parameter  $\lambda$ . Accordingly, we need some procedure for selecting a value for  $\lambda$ . In regularization problems, it is common to select the regularization parameter via cross-validation (Bien & Tibshirani, 2011; Chi & Lange, 2014; J. Z. Huang, Liu, Pourahmadi, & Liu, 2006). This approach is too computationally intensive to be feasible for our application. Among shrinkage estimators, it is common to choose the amount of shrinkage in order to minimize an expected loss function (James & Stein, 1961; Ledoit & Wolf, 2003). However, no suitable analytic result exists that allows us to approximately minimize expected loss in our scenario.

Consequently, we developed a heuristic criterion that selects  $\lambda$  in a way that aligns with the goal of our regularization process. Our method’s intent is to shrink the magnitude of penalized elements of the correlation matrix while leaving unpenalized elements more or less unchanged. In practice, although we succeed at bringing penalized elements towards zero, this shrinkage usually comes at the cost of inflating other elements. We have observed that this inflation tends to grow more pronounced as  $\lambda$  grows. For very large values of  $\lambda$ , our estimated correlation matrix may shrink nearly all penalized entries to zero at the expense of inflating a few elements (both penalized and unpenalized) to nearly  $\pm 1$ . This is not a desirable outcome.

Although it may seem counterintuitive at first, the observed inflation is not an artifact of a coding error or poor convergence of our algorithm. A simple reproducible example of inflation in a  $3 \times 3$  matrix is provided in Appendix B. In this low-dimensional setting, standard numerical optimization routines agree with the results from our code and both display inflation of unpenalized elements.

Our criterion for selecting  $\lambda$  compares the off-diagonal elements of  $\tilde{R}$  and  $\hat{R}(\lambda)$ . We choose the value of  $\lambda$  which maximizes the difference between average shrinkage and average

inflation. Formally, our criterion is defined by

$$k(\tilde{R}, \lambda) = \underset{i,j \text{ s.t. } |\hat{R}(\lambda)_{ij}| < |\tilde{R}_{ij}|}{\text{mean}} \left( |\tilde{R}_{ij}| - |\hat{R}(\lambda)_{ij}| \right) - \underset{i,j \text{ s.t. } |\hat{R}(\lambda)_{ij}| > |\tilde{R}_{ij}|}{\text{mean}} \left( |\hat{R}(\lambda)_{ij}| - |\tilde{R}_{ij}| \right). \quad (19)$$

Large positive values of  $k$  are desirable, as they correspond to values of  $\lambda$  for which we induce a lot of shrinkage and not much inflation.

## 3 Results

In this section, we first report results from applying our method to global migration data in Section 3.1. Section 3.2 then provides a simulation study which demonstrates that our method outperforms Pearson correlations and the Ledoit-Wolf shrinkage estimator (Ledoit & Wolf, 2003) in the scenario where the penalty matrix  $P$  is appropriate to the true correlation structure.

### 3.1 Application to migration

#### 3.1.1 Data

We use data on net migration from the 2012 revision of the World Population Prospects (WPP) United Nations (2012). The WPP contains estimates of net migration for all countries in five-year time periods from 1950 until 2010, a total of 12 time periods. We compute the net migration rate  $g_{c,t}$  as the net number of migrants in country  $c$  over the five year period starting at time  $t$ , divided by thousands of individuals in country  $c$  at time  $t$ .

Because we want to express prior beliefs as a function of distance covariates, we restrict the set of modeled countries to the 191-country overlap between the WPP 2012 and the set of countries included in CEPII’s GeoDist database, a database of bilateral distance covariates defined on pairs of countries (Mayer & Zignago, 2011).

#### 3.1.2 Selection of $P$

Our estimation technique requires that we choose a penalty matrix,  $P$ , that reflects our prior beliefs about which country pairs are likely to be correlated. Although it would be possible to elicit expert opinion about each of the roughly 18,000 country pairs, we instead choose a  $P$  that can be characterized in terms of just a few covariates. Our matrix  $P$  penalizes a pair of countries if *none* of the following conditions is met:

1. The two countries are contiguous.
2. The two countries’ most important cities are located less than 3000 km apart.



**Table 1:** Results of Kolmogorov-Smirnov test that empirical correlations are significantly different from the distribution of elements of a sample correlation matrix when the true error structure is uncorrelated.  $p$ -values lower than 0.05 are bolded.

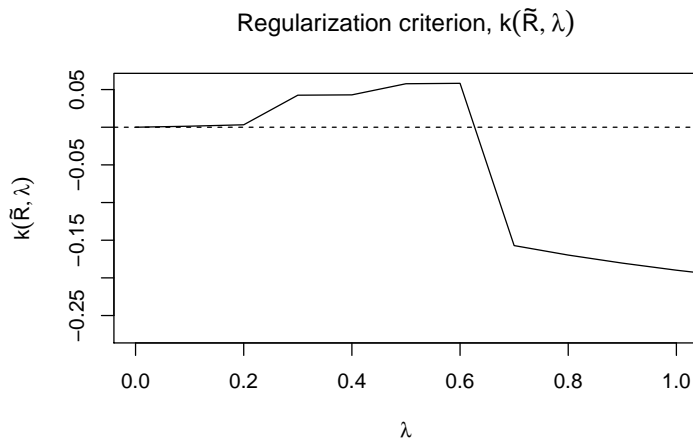
	Covariate	$p$ -value
	Contiguous	<b>0.019</b>
	Common language (official)	0.23
	Common language (spoken by 9% of pop.)	0.58
	Geodesic distance less than 3000 km	<b>0.0003</b>
	Colonial relationship after 1945	0.57
	Common colonizer after 1945	0.11
	Current colonial relationship	<b>0.035</b>
	Ever had a colonial link	0.36
	Same UN Region	<b>0.036</b>

3. The two countries are in the same region according to the United Nations Population Division’s division of the world into 22 regions, based on both geographical contiguity and cultural affinity ([United Nations, 2012](#)).

4. The two countries are currently in a colonial relationship.

This definition of  $P$  is in line with migration theory, which suggests that migrant flows are more likely when monetary and social costs of movement are low ([Harris & Todaro, 1970](#); [Lee, 1966](#); [Sjaastad, 1962](#); [Stark & Bloom, 1985](#)), as will be the case with countries which are geographically proximate or share administrative ties. This definition penalizes 85% of country pairs, leaving 15% unpenalized. The average country is considered to be “close” to 29 other countries, and “distant” from the remaining 161.

In selecting these conditions, we examined nine candidate distance covariates. The first eight such covariates come from CEPII’s GeoDist database ([Mayer & Zignago, 2011](#)), while the ninth is derived from the United Nations division into 22 regions. The left column of [Table 1](#) gives the complete list of covariates considered. As an empirical basis for determining which criteria to include in defining our penalty matrix, we examined the elements of the sample correlation matrix for all pairs of countries meeting each criterion. Using a Kolmogorov-Smirnov test, we tested whether the distribution of these sample correlations was different from the distribution of elements of the sample correlation matrix under a null hypothesis of uncorrelated errors. The right column of [Table 1](#) shows the  $p$ -values from these Kolmogorov-Smirnov tests. Our definition of the penalty matrix  $P$  includes all covariates with a  $p$ -value less than 0.05.



**Figure 3:** Regularization criterion,  $k(\tilde{R}, \lambda)$  as a function of  $\lambda$ . The regularization criterion is the difference between the average shrinkage among shrunk elements of  $\hat{R}(\lambda)$  and average inflation among inflated elements.

### 3.1.3 Selection of the regularization parameter, $\lambda$


We computed values of  $\hat{R}(\lambda)$  for all values of  $\lambda$  from 0 to 3 in increments of 0.1. Figure 3 shows the value of  $k(\tilde{R}, \lambda)$  over a range of  $\lambda$  values. We found that  $k(\tilde{R}, \lambda)$  peaked at  $\lambda = 0.6$ , where we find average shrinkage of 0.13 compared with average inflation of 0.07. Increasing  $\lambda$  from 0.6 to 0.7 induces additional shrinkage, but at the cost of greatly inflating some correlations. Accordingly, we choose  $\hat{R}(0.6)$  as our estimate of  $R$ .

Figure 4 shows the impact of regularization on the correlation matrix. Among penalized elements (top panel), we see significant shrinkage towards zero, although many penalized elements remain large in magnitude, even after regularization. The bottom panel shows the unpenalized elements of the correlation matrix before regularization (solid curve) and after (dashed curve). On average we induce some shrinkage in the unpenalized elements, but the distribution is largely unchanged.

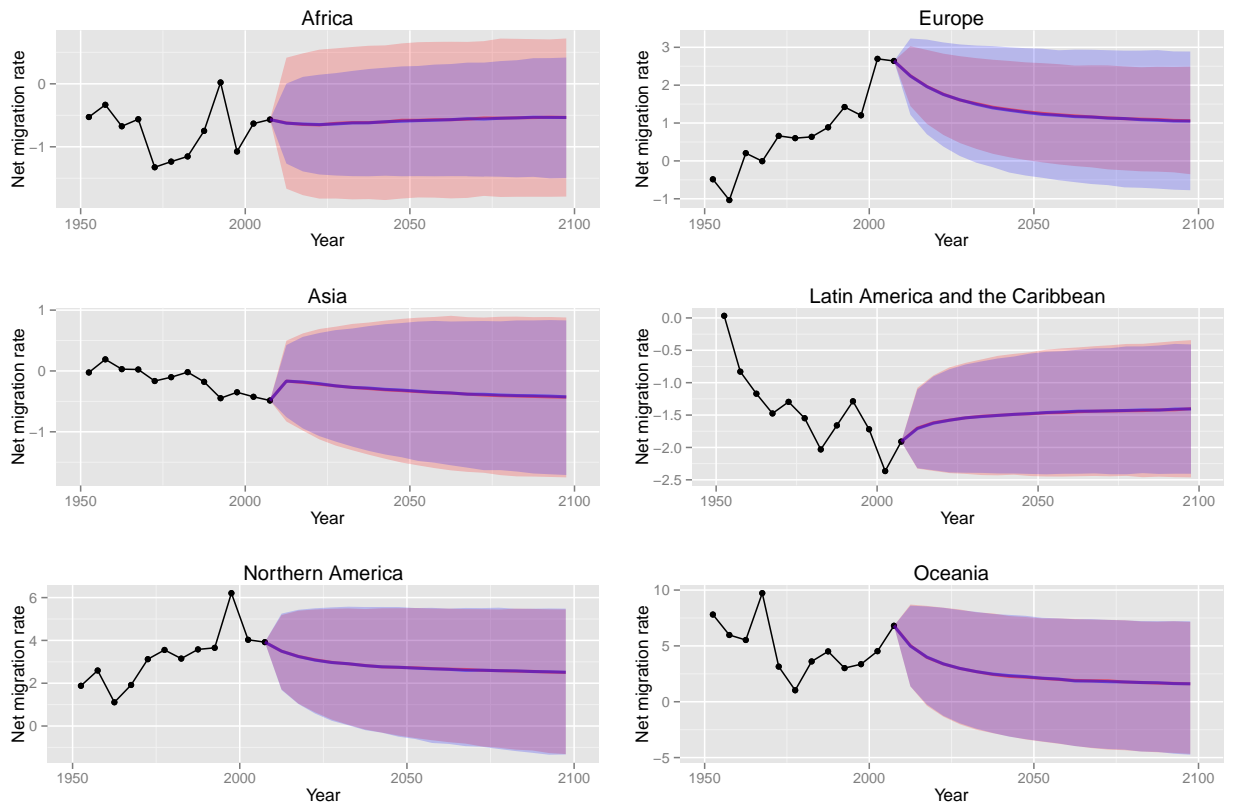
### 3.1.4 Projection and evaluation

We augment the BHM+IFE model with the LPoC estimate  $\hat{R}(0.6)$  to produce probabilistic projections of migration for any collection of countries. Figure 5 contains medians and 80% prediction intervals of projected migration for all continents. In Africa, negative correlations narrow our projections. In Europe, positive correlations cause forecasts to widen. For the other continents, we see little change in projected migration.

For evaluation, we compare true migration rates for regional aggregates in 1995–2010 with projections of the same regional aggregates based only on migration data from 1950–1995.



**Figure 4:** Comparison of elements of the correlation matrix before regularization (solid curves) and after (dashed curves). Top panel shows penalized elements; bottom panel shows unpenalized elements.



**Figure 5:** Medians and 80% prediction intervals for net migration among continents. Projections from the Bayesian hierarchical model with independent forecast errors (BHM+IFE) are given in red. Projections using our estimated correlation matrix (BHM+LPoC) are in blue. Overlap is in purple.

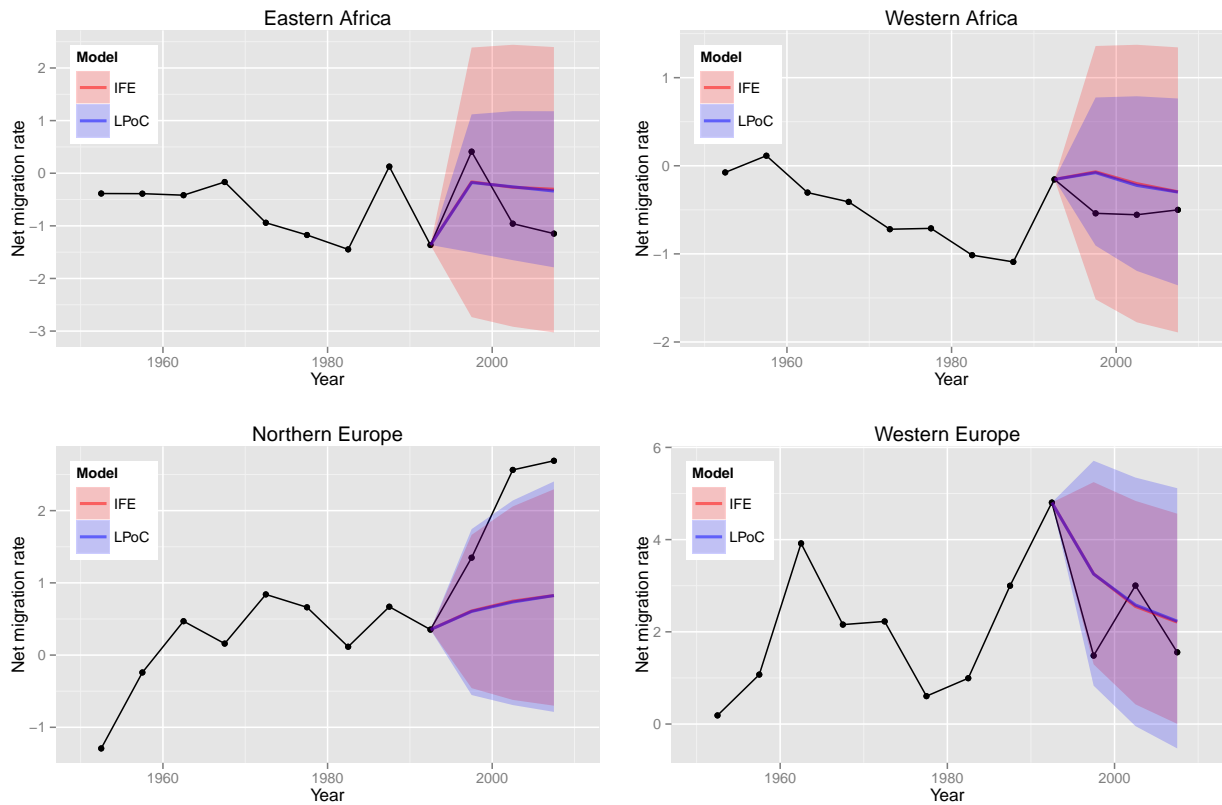
This procedure entails re-estimation of the BHM+IFE model using only the 1950–1995 data, followed by construction of an empirical correlation matrix, selection of  $\lambda$ , and extraction of  $\hat{R}(\lambda)$ . We compare the performance of the BHM+IFE model on regional aggregates to a model using the same sampled values of  $\boldsymbol{\mu}$ ,  $\boldsymbol{\phi}$ , and  $\boldsymbol{\sigma}$ , but augmented with  $\hat{R}(\lambda)$ .

As an evaluation metric, we use the negatively oriented continuous ranked probability score (CRPS) (Gneiting & Raftery, 2007; Hersbach, 2000). The CRPS compares the cumulative distribution function,  $F$ , of a probabilistic forecast to an observation,  $x$ , and is defined by

$$\text{CRPS}(F, x) = \int_{-\infty}^{\infty} (F(y) - \mathbb{1}\{y \geq x\})^2 dy. \quad (20)$$

In our application the two probabilistic forecasts under consideration have the same mean as one other, by design. One approximate way of looking at CRPS in this setting is that when  $g_{c,t}$  is close to the mean of the forecast, we reward  $F$  for having low variance; when  $g_{c,t}$  is far from the mean, we reward  $F$  for having high variance.

Table 2 gives CRPS for projections of aggregate migration for the six continents. Our model improves the quality of projections in Africa and Europe, while projections for the other four continents are more or less unchanged. Figure 6 illustrates the change in projections of net migration in 1995–2010 for four subregions of Africa and Europe. Projections from the BHM+IFE model are in red; projections from BHM+LPoC are in blue. Our method narrows prediction intervals in Eastern and Western Africa, bringing the width of the 80% prediction intervals more into line with the range of observed variability. In both regions, true migration rates for the projected period stayed within our narrower intervals. In contrast, our method widens projections in Northern and Western Europe, where the 80% intervals from the BHM+IFE model either miss or nearly miss capturing some of the observed data points.



**Figure 6:** Medians and 80% prediction intervals for projections of net migration rates for regional aggregates. Projections from the Bayesian hierarchical model with independent forecast errors (BHM+IFE) are given in red. Projections using our estimated correlation matrix (BHM+LPoC) are in blue. Overlap is in purple.

**Table 2:** Continuous ranked probability score for all continents evaluated on projections of 1995-2010, where lower is better. Left column: Projections based on the Bayesian hierarchical model with independent correlation structure (BHM+IFE). Right column: Projections based on the Bayesian hierarchical model with our regularized correlation estimate (BHM+LPoC). Bolded entry in each row indicates the lower value.

	IFE	LPoC
Africa	1.66	<b>1.49</b>
Asia	<b>0.73</b>	0.74
Europe	3.92	<b>3.76</b>
Latin America and the Caribbean	<b>1.62</b>	<b>1.62</b>
Northern America	5.02	<b>4.99</b>
Oceania	8.53	<b>8.49</b>

## 3.2 Simulation study

In this section we show by simulation that our regularization procedure improves correlation estimates in a low-dimensional setting. To match the application of interest, we simulate 12 observed time points from an AR(1) process with correlated errors. For computational tractability, we decrease the number of simulated countries from 191 in the real data to 9 in the simulation. For each of 100 simulations, we perform the following procedure:

1. Generate a set of simulated migration rates  $\mathbf{g}_1, \dots, \mathbf{g}_{12}$  from an AR(1) process with errors correlated as described below.
2. Produce point estimates of  $\boldsymbol{\varepsilon}_1, \dots, \boldsymbol{\varepsilon}_{11}$  via MCMC sampling of  $\boldsymbol{\mu}$ ,  $\boldsymbol{\phi}$ , and  $\boldsymbol{\sigma}$ .
3. Convert  $\boldsymbol{\varepsilon}_t$ 's to a matrix  $\tilde{R}$  using the procedure in Section 2.2.
4. Solve the minimization problem (12) to obtain a regularized estimate for the correlation matrix.

Since the procedure for selecting  $\lambda$  is computationally intensive, we perform this procedure only once and use the same value of  $\lambda$  for all subsequent simulations.

### 3.2.1 Simulation details

We simulate a collection of nine countries with true migration rates governed by the AR(1) process

$$\mathbf{g}_t - \boldsymbol{\mu} = \text{diag}(\boldsymbol{\phi})(\mathbf{g}_{t-1} - \boldsymbol{\mu}) + \boldsymbol{\varepsilon}_t. \quad (21)$$

For simplicity we take  $\boldsymbol{\mu} = \mathbf{0}$ ,  $\boldsymbol{\phi} = \frac{1}{2}\mathbf{1}$ , and

$$\boldsymbol{\varepsilon}_t \stackrel{iid}{\sim} \mathcal{N}_9(\mathbf{0}, \Sigma). \quad (22)$$

We fix  $\Sigma$  to be block diagonal. Compound symmetric correlation structure within each  $3 \times 3$  block is given by

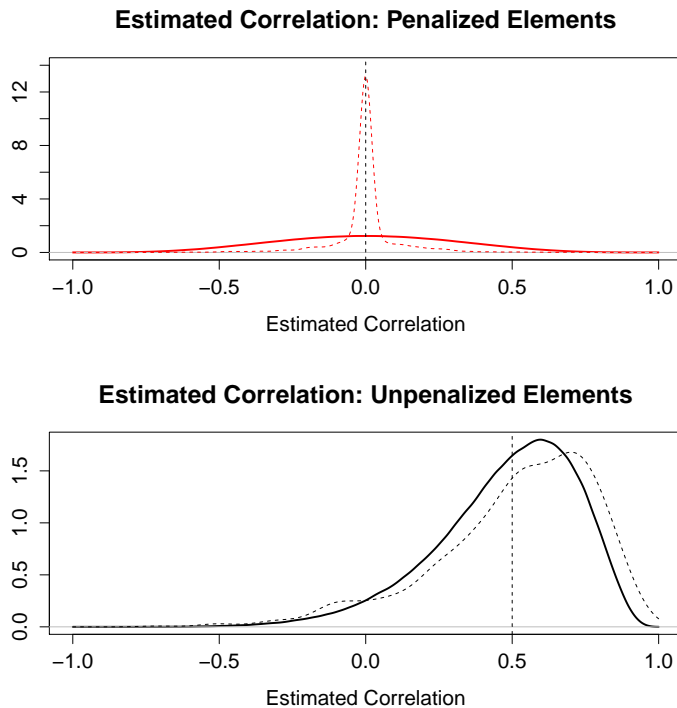
$$\Sigma_{3 \times 3} = \begin{pmatrix} 1 & 0.5 & 0.5 \\ 0.5 & 1 & 0.5 \\ 0.5 & 0.5 & 1 \end{pmatrix}, \quad (23)$$

and the full covariance matrix by

$$\Sigma = \begin{pmatrix} \Sigma_{3 \times 3} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \Sigma_{3 \times 3} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \Sigma_{3 \times 3} \end{pmatrix}. \quad (24)$$

We then simulate observations  $\mathbf{g}_1, \dots, \mathbf{g}_{12}$  and attempt to make inference on the correlation structure of  $\Sigma$ .






**Figure 7:** Simulation study results: Comparison of elements of the correlation matrix before regularization (solid curves) and after (dashed curves). Top panel shows penalized elements; bottom panel shows unpenalized elements. True correlations are indicated with dashed vertical lines.

Because we are basing inference on a small number of time points, Pearson estimates of correlation are highly variable. Solid curves in Figure 7 show the distributions of the off-diagonal elements of the unregularized Pearson correlation matrix in the ideal scenario where the values of  $\varepsilon_t$  can be perfectly estimated. The top panel shows the distribution of the elements for which the true correlation is zero. The bottom panel shows elements for which the true correlation is 0.5. In both cases, high variability makes inference difficult. Our method is designed to decrease variability among estimated correlations for those country pairs where prior knowledge suggests that correlation should be close to zero.

To illustrate a best case scenario, we choose a penalty matrix  $P$  which is well suited to the true correlation structure. The simplest such  $P$  is the one which penalizes the off-diagonal elements of the correlation matrix if and only if the true correlation is zero. That  $P$  is given by

$$P = \begin{pmatrix} \mathbf{0}_{3 \times 3} & \mathbf{1}_{3 \times 3} & \mathbf{1}_{3 \times 3} \\ \mathbf{1}_{3 \times 3} & \mathbf{0}_{3 \times 3} & \mathbf{1}_{3 \times 3} \\ \mathbf{1}_{3 \times 3} & \mathbf{1}_{3 \times 3} & \mathbf{0}_{3 \times 3} \end{pmatrix}. \quad (25)$$



**Figure 8:** Average shrinkage minus average inflation of elements of  $\hat{R}(\lambda)$  as  $\lambda$  varies from 0 to 10. Exact curve in black, Lowess-smoothed curve in red.

### 3.2.2 Initial run to select $\lambda$

Our procedure to select  $\lambda$  is computationally expensive, as it requires us to compute  $\hat{R}(\lambda)$  repeatedly as  $\lambda$  varies. We therefore perform this procedure only once and use the same  $\lambda$  for estimation of  $R$  in all subsequent simulated data sets. Figure 8 plots our  $\lambda$ -selection criterion based on a single simulated data set over the range  $\lambda = 0, 0.1, 0.2, \dots, 10$ . The exact curve, shown in black, exhibits some jumpiness in this low-dimensional setting, a problem which naturally becomes less severe in the high-dimensional setting of interest. Because of this jumpiness, we base our selection of  $\lambda$  on a Lowess-smoothed curve, selecting the maximizing value of  $\lambda = 6.4$ .

### 3.2.3 Evaluation of repeated estimation of $R$

We produced 100 estimates of  $\hat{R}(\lambda = 6.4)$  from 100 different sets of simulated migration rates, all using the same block diagonal correlation structure. Dashed lines in Figure 7 show the distribution of off-diagonal elements of  $\hat{R}$ , split into those elements where the true correlation is 0 and elements where the true correlation is 0.5 (top and bottom panel, respectively).

Our method is successful in shrinking penalized elements towards zero. Among elements where the true correlation is zero, we correctly estimate an exact zero in 62% of cases in this simulation. Among unpenalized elements, our method produces estimates with slightly more variability (the standard deviation is 0.256 for Pearson correlations versus 0.272 for our estimates). Both methods produce estimates for unpenalized elements that are within two standard errors of the true mean value of 0.5. The mean estimated correlation is 0.489 for Pearson correlations (standard error 0.009) versus 0.514 for our estimates (standard error of

0.009). On the whole, the LPoC estimator greatly improves estimates of penalized elements at the expense of slightly increasing variability in unpenalized elements.

Table 3 compares mean absolute error and mean squared error from our method with two competing estimators. We compare our results against both Pearson correlation matrices and correlation matrices that have been regularized using the Ledoit-Wolf method, which shrinks Pearson estimates towards a spherical correlation structure (Ledoit & Wolf, 2003). In the top panel, we estimate  $\epsilon_1, \dots, \epsilon_{11}$  with a Bayesian hierarchical model, as is done in our real application to migration. In the bottom panel, we assume instead a scenario where we have direct access to  $\epsilon_1, \dots, \epsilon_{11}$ , as would be suitable in other applications where the interest is in estimating correlations of directly observed quantities. In both cases, our method provides an overall reduction in mean squared error by at least two thirds when compared against the Pearson sample correlation matrix. A large reduction in error from shrinking penalized elements is offset by a mild increase in error among unpenalized elements. We also outperform the Ledoit-Wolf estimator in terms of overall error.

## 4 Discussion

Our method augments probabilistic projections of migration that are well-calibrated for individual countries, with a correlation structure that reflects prior knowledge of between-country correlations. By combining a high-dimensional empirical correlation matrix with an informative prior that shrinks spurious correlations, we produce an estimated correlation matrix that is in line with migration theory and improves projections of regional aggregates. When compared with a simple model that assumes uncorrelated forecast errors, our method narrows projections of net migration for Africa and widens projections for Europe. Out-of-sample evaluation confirms that these changes produce better probabilistic forecasts as measured by continuous ranked probability score. Mechanically, the novelty of our method is our prior on correlations, which benefits from being interpretable and simple in form, and converts MAP estimation to an  $\ell_1$ -penalized regularization problem which is computationally tractable.

Our analysis focuses on modeling net migration, but an attractive alternative would be to model a full matrix of bilateral migration flows. Such a model would naturally imply correlations in migration—if out-migrants from country  $i$  tend to go to country  $j$ , then net migration in countries  $i$  and  $j$  will be negatively correlated. However, modeling the global bilateral flow matrix is currently not feasible. Flows are hard to estimate, even in countries with good data (De Beer, Raymer, Van der Erf, & Van Wissen, 2010; Raymer, de Beer, & van der Erf, 2011). Abel (2013) produces global estimates of migration flows based on

**Table 3:** Evaluation of correlation matrix estimates from simulation study. “LPoC” refers to our estimator, which uses Laplace priors on correlations. MAE is mean absolute error. MSE is mean squared error. Averages over “all elements” exclude diagonal elements, which are fixed at zero by all methods. The lowest (best) values are shown in bold.

Values of $\epsilon_t$ estimated with MCMC			
	Estimator	MAE	MSE
All elements	Pearson	0.253	0.098
	Ledoit-Wolf	0.193	0.055
	LPoC	<b>0.090</b>	<b>0.028</b>
True correlation = 0	Pearson	0.270	0.109
	Ledoit-Wolf	0.190	0.053
	LPoC	<b>0.049</b>	<b>0.012</b>
True correlation = 0.5	Pearson	0.201	0.066
	Ledoit-Wolf	<b>0.200</b>	<b>0.060</b>
	LPoC	0.214	0.074
True values of $\epsilon_t$ used			
	Estimator	MAE	MSE
All elements	Pearson	0.227	0.079
	Ledoit-Wolf	0.182	0.047
	LPoC	<b>0.078</b>	<b>0.022</b>
True correlation = 0	Pearson	0.244	0.089
	Ledoit-Wolf	0.162	0.039
	LPoC	<b>0.041</b>	<b>0.010</b>
True correlation = 0.5	Pearson	<b>0.176</b>	<b>0.051</b>
	Ledoit-Wolf	0.243	0.073
	LPoC	0.190	0.058

migrant stock data, but for only a small number of time periods at which migrant stock data exist. His method involves minimizing the total number of migrants subject to the available data on migrant stocks. This induces many structural zeroes in his estimates, making modeling difficult. Because of the lack of good data on migration flows, we choose instead to work with net migration rates.

Although our method produces a MAP estimator in the presence of informative priors, we are not able to leverage any of the usual Bayesian machinery to produce a sample from the posterior distribution. While it would in theory be possible to use MCMC methods to produce a posterior sample by updating one element of the correlation matrix at a time, an updating procedure would need to iterate through some 18,000 elements of the correlation matrix, checking for positive definiteness after each proposed step. Such an algorithm is therefore likely to move around the parameter space too slowly to be of any use. In some settings a Laplace approximation centered at the posterior mode can provide a good approximation of marginal posterior distributions (Tierney & Kadane, 1986). However, the double-exponential priors in our setting render this procedure impracticable. Within each orthant of the parameter space, a quadratic approximation to the log likelihood is reasonable, but because of the  $\ell_1$  penalty term, a different quadratic approximation is required for each of the roughly  $2^{18,000}$  orthants, which is not feasible.

Given our interest in combining data with prior beliefs, an inverse Wishart prior on covariance is tempting because it allows easy sampling from the full posterior. However, the inverse Wishart distribution is restrictive in form (Barnard et al., 2000) and does not provide a straightforward way to describe prior beliefs about correlations.

Another tempting alternative is that of Liu, Wang, and Zhao (2014), who give a simple thresholding method for producing a penalized correlation matrix that is guaranteed to be positive definite. Their estimator solves

$$\operatorname{argmin}_{\omega \succ \delta \cdot I} \frac{1}{2} \|\tilde{R} - \omega\|_F^2 + \lambda \|W * \omega\|_{1,\text{off}}, \quad (26)$$

to produce an estimator among the set of valid correlation matrices with minimum eigenvalue no smaller than  $\delta$ . Although the weight matrix,  $W$ , is in principle arbitrary, they use  $W$  to induce greater shrinkage where empirical correlations are weakest, not as a means of conveying prior information. We would be hesitant to replace  $W$  with our penalty matrix  $P$ , as that off-license use of their method would not incorporate prior information in a principled way.

Our method can be generalized to shrink estimated correlations towards non-zero values by replacing the penalty term  $\lambda \|P * R\|_1$  with  $\lambda \|P * (R - S)\|_1$  for some target matrix  $S$ . This may be desirable in cases where heavily structured estimates of correlations are available, as

is the case for modeling of fertility (Fosdick & Raftery, 2014).

Note that we have used the 2012 revision of the WPP here (United Nations, 2012). The more recent 2015 revision (United Nations, 2015) contains one additional data point. It would be of interest to redo the analysis with the newer data, but we expect the results would be similar.

## References

- Abel, G. (2013). Estimating global migration flow tables using place of birth data. *Demographic Research*, 28, 505–546.
- Antoniadis, A., & Fan, J. (2001). Regularization of wavelet approximations. *Journal of the American Statistical Association*, 96(455), 939–955.
- Azose, J. J., & Raftery, A. E. (2015). Bayesian probabilistic projection of international migration. *Demography*, 52(5), 1627–1650.
- Azose, J. J., Ševčíková, H., & Raftery, A. E. (2016). Probabilistic population projections with migration uncertainty. *Proceedings of the National Academy of Sciences*. (Published ahead of print May 23, 2016, doi: 10.1073/pnas.1606119113)
- Barnard, J., McCulloch, R., & Meng, X.-L. (2000). Modeling covariance matrices in terms of standard deviations and correlations, with application to shrinkage. *Statistica Sinica*, 10(4), 1281–1312.
- Beck, A., & Teboulle, M. (2009). A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal of Imaging Sciences*, 2(1), 183–202.
- Bickel, P. J., & Levina, E. (2008a). Covariance regularization by thresholding. *The Annals of Statistics*, 36(6), 2577–2604.
- Bickel, P. J., & Levina, E. (2008b). Regularized estimation of large covariance matrices. *The Annals of Statistics*, 36(1), 199–227.
- Bien, J., & Tibshirani, R. J. (2011). Sparse estimation of a covariance matrix. *Biometrika*, 98(4), 807.
- Bijak, J., Kupiszewska, D., Kupiszewski, M., Saczuk, K., & Kicingier, A. (2007). Population and labour force projections for 27 European countries, 2002–2052: Impact of international migration on population ageing. *European Journal of Population*, 23(1), 1–31.
- Bijak, J., & Wiśniowski, A. (2010). Bayesian forecasting of immigration to selected European countries by using expert knowledge. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 173(4), 775–796.
- Brown, S. K., & Bean, F. D. (2012). Population growth. In J. Gans, E. M. Replogle, & D. J. Tichenor (Eds.), *Debates on U.S. Immigration*. Thousand Oaks, California: SAGE.
- Chaudhuri, S., Drton, M., & Richardson, T. S. (2007). Estimation of a covariance matrix with zeros. *Biometrika*, 94(1), 199–216.
- Chen, X., Xu, M., Wu, W. B., et al. (2013). Covariance and precision matrix estimation for high-dimensional time series. *The Annals of Statistics*, 41(6), 2994–3021.

- Chi, E. C., & Lange, K. (2014). Stable estimation of a covariance matrix guided by nuclear norm penalties. *Computational Statistics & Data Analysis*, *80*, 117–128.
- Crush, J. (1999). Fortress South Africa and the deconstruction of Apartheid’s migration regime. *Geoforum*, *30*(1), 1–11.
- Cui, Y., Leng, C., & Sun, D. (2016). Sparse estimation of high-dimensional correlation matrices. *Computational Statistics & Data Analysis*, *93*, 390–403.
- De Beer, J., Raymer, J., Van der Erf, R., & Van Wissen, L. (2010). Overcoming the problems of inconsistent international migration data: A new method applied to flows in Europe. *European Journal of Population*, *26*(4), 459–481.
- Deng, X., & Tsui, K.-W. (2013). Penalized covariance matrix estimation using a matrix-logarithm transformation. *Journal of Computational and Graphical Statistics*, *22*(2), 494–512.
- El Karoui, N. (2008). Operator norm consistent estimation of large-dimensional sparse covariance matrices. *The Annals of Statistics*, *36*(6), 2717–2756.
- Fan, J., Han, F., & Liu, H. (2014). Challenges of big data analysis. *National Science Review*, *1*(2), 293–314.
- Fan, J., Huang, T., & Li, R. (2007). Analysis of longitudinal data with semiparametric estimation of covariance function. *Journal of the American Statistical Association*, *102*(478), 632–641.
- Fan, J., Liao, Y., & Liu, H. (2015). An overview on the estimation of large covariance and precision matrices. *arXiv preprint arXiv:1504.02995*.
- Fan, J., Liao, Y., & Mincheva, M. (2013). Large covariance estimation by thresholding principal orthogonal complements. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *75*(4), 603–680.
- Fassmann, H., & Munz, R. (1994). European East-West migration, 1945-1992. *International Migration Review*, *28*(3), 520–538.
- Fosdick, B. K., & Raftery, A. E. (2014). Regional probabilistic fertility forecasting by modeling between-country correlations. *Demographic Research*, *30*(35), 1011.
- Friedman, J., Hastie, T., & Tibshirani, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, *9*(3), 432–441.
- Furrer, R., & Bengtsson, T. (2007). Estimation of high-dimensional prior and posterior covariance matrices in Kalman filter variants. *Journal of Multivariate Analysis*, *98*(2), 227–255.
- Gneiting, T., & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, *102*(477), 359–378.
- Harris, J. R., & Todaro, M. P. (1970). Migration, unemployment and development: A two-sector analysis. *American Economic Review*, *60*(1), 126–142.
- Hersbach, H. (2000). Decomposition of the continuous ranked probability score for ensemble prediction systems. *Weather and Forecasting*, *15*(5), 559–570.
- Huang, A., Wand, M. P., et al. (2013). Simple marginally noninformative prior distributions for covariance matrices. *Bayesian Analysis*, *8*(2), 439–452.
- Huang, J. Z., Liu, N., Pourahmadi, M., & Liu, L. (2006). Covariance matrix selection and estimation via penalised normal likelihood. *Biometrika*, *93*(1), 85–98.
- James, W., & Stein, C. (1961). Estimation with quadratic loss. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability* (Vol. 1, pp. 361–379).

- Ledoit, O., & Wolf, M. (2003). Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *Journal of Empirical Finance*, 10(5), 603–621.
- Ledoit, O., & Wolf, M. (2004). A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis*, 88(2), 365–411.
- Ledoit, O., & Wolf, M. (2012). Nonlinear shrinkage estimation of large-dimensional covariance matrices. *The Annals of Statistics*, 40(2), 1024–1060.
- Lee, E. S. (1966). A theory of migration. *Demography*, 3(1), 47–57.
- Leonard, T., & Hsu, J. S. (1992). Bayesian inference for a covariance matrix. *The Annals of Statistics*, 20(4), 1669–1696.
- Levina, E., Rothman, A., Zhu, J., et al. (2008). Sparse estimation of large covariance matrices via a nested lasso penalty. *The Annals of Applied Statistics*, 2(1), 245–263.
- Liechty, J. C., Liechty, M. W., & Müller, P. (2004). Bayesian correlation estimation. *Biometrika*, 91(1), 1–14.
- Liu, H., Wang, L., & Zhao, T. (2014). Sparse covariance matrix estimation with eigenvalue constraints. *Journal of Computational and Graphical Statistics*, 23(2), 439–459.
- Mayer, T., & Zignago, S. (2011). Notes on CEPII’s distances measures: The GeoDist database.
- Nocedal, J., & Wright, S. (2006). *Numerical Optimization*. Springer Science & Business Media.
- Okólski, M. (1998). Regional dimension of international migration in Central and Eastern Europe. *Genus*, 54(1), 11–36.
- Pourahmadi, M. (2011). Covariance estimation: The glm and regularization perspectives. *Statistical Science*, 26(3), 369–387.
- Raymer, J., de Beer, J., & van der Erf, R. (2011). Putting the pieces of the puzzle together: Age and sex-specific estimates of migration amongst countries in the EU/EFTA, 2002–2007. *European Journal of Population*, 27(2), 185–215.
- Sjaastad, L. A. (1962). The costs and returns of human migration. *Journal of Political Economy*, 70(5), 80–93.
- Stark, O., & Bloom, D. E. (1985). The new economics of labor migration. *American Economic Review*, 75(2), 173–178.
- Tierney, L., & Kadane, J. B. (1986). Accurate approximations for posterior moments and marginal densities. *Journal of the American Statistical Association*, 81(393), 82–86.
- United Nations. (2012). *World Population Prospects: The 2012 Revision*. United Nations, New York.
- United Nations. (2015). *World Population Prospects: The 2015 Revision*. United Nations, New York.
- U.S. Social Security Administration. (2013). *The 2013 Annual Report of the Board of Trustees of the Federal Old-age and Survivors Insurance and Federal Disability Insurance Trust Funds*. Board of Trustees, Federal Old-Age and Survivors Insurance and Federal Disability Insurance Trust Funds.
- Wei, G. C., & Tanner, M. A. (1990). A Monte Carlo implementation of the EM algorithm and the poor man’s data augmentation algorithms. *Journal of the American Statistical Association*, 85(411), 699–704.
- Wiśniowski, A., Smith, P. W., Bijak, J., Raymer, J., & Forster, J. J. (2015). Bayesian



- population forecasting: Extending the Lee-Carter method. *Demography*, 52(3), 1035–1059.
- Wright, E. (2010). 2008-based national population projections for the United Kingdom and constituent countries. *Population Trends*, 139(1), 91–114.
- Zhang, T., & Zou, H. (2012). Sparse precision matrix estimation via lasso penalized D-trace loss. *Biometrika*, 99(1), 1–18.

## A Determining step size

Step size selection is necessary in high dimensions for the general gradient descent algorithm to converge quickly enough to be useful. Complex methods for step size selection are available, but we obtained reasonable results with the backtracking line search algorithm, which starts with a large step size and decreases step size whenever a proposed step results in too little improvement in the objective function.

Say we have an objective function  $f(x)$  which we are trying to minimize. The core of the backtracking line search algorithm is as follows (Nocedal & Wright, 2006).

1. Fix a backtracking coefficient  $\beta \in (0, 1)$ , a starting step size  $\alpha_0$ , and a starting location  $x_0$ .
2. Propose a step of length  $\alpha_k$  in direction  $p_k$ . (The backtracking line search algorithm is a generic algorithm that will work regardless of how the direction  $p_k$  is determined.)
3. If the improvement in the objective function is enough to meet the Armijo condition given in (27) below, then take the proposed step. That is, take  $x_{k+1} = x_k + \alpha_k p_k$ . Keep the step size constant (i.e.,  $\alpha_{k+1} = \alpha_k$ ).
4. Otherwise, if there is any improvement in the objective function, take the proposed step, but also decrease the step size for the next iteration (specifically, set  $\alpha_{k+1} = \beta \alpha_k$ ).
5. Otherwise, there must have been no improvement in the objective function. Don't take a step, but do decrease step size. ( $x_{k+1} = x_k$  and  $\alpha_{k+1} = \beta \alpha_k$ .)
6. Repeat steps 2-5 until convergence.

The Armijo condition, which is used to determine whether to decrease step size, is as follows. The Armijo condition is met if the following inequality is satisfied:

$$f(x_k + \alpha_k p_k) \leq f(x_k) + c_1 \alpha_k \nabla f_k^T p_k. \quad (27)$$

( $c_1$  is a constant chosen from  $(0, 1)$  that controls how strictly the change in  $f$  must match the gradient at  $x_k$ .)

In our application, there's a missing component—we can't actually compute the gradient of our objective function. The relevant objective function is given by

$$f(R) = \text{tr}(R_i^{-1} R) + \text{tr}(R^{-1} \tilde{R}) + \lambda \|P * R\|_1. \quad (28)$$

The first two terms in the sum are differentiable, but the third is not.

We rewrite the Armijo condition as

$$f(x_k + \alpha_k p_k) \leq f(x_k) + c_1 \alpha_i p_k^T p_k - c_1 \alpha_i (p_k - \nabla f_k)^T p_k \quad (29)$$

and then approximate  $(p_k - \nabla f_k)$  with

$$-2 \cdot \nabla(\text{tr}(R_i^{-1} R) + \text{tr}(R^{-1} \tilde{R})). \quad (30)$$

## B Inflation of correlation estimates

We provide here an example of our correlation estimation procedure which produces inflation in some unpenalized elements of the correlation matrix. We solved the minimization problem in (12) with three different methods, finding identical answers each time, up to small numerical tolerances. Those methods are:

1. Estimate  $R$  using our code, which appeals to the generalized gradient descent algorithm.
2. Estimate  $R$  using a black-box numerical optimization algorithm, which has access to the function we're minimizing, but not its derivative.
3. Estimate  $R$  by finding an analytic expression for the gradient of the function we're minimizing, and solve for a point where the gradient is zero.

One case in which inflation manifests if we take our evidence from the data to be given by

$$\tilde{R} = \begin{pmatrix} 1 & 0.8 & 0.5 \\ 0.8 & 1 & 0.1 \\ 0.5 & 0.1 & 1 \end{pmatrix} \quad (31)$$

and the penalty matrix by

$$P = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}. \quad (32)$$

We denote the unknown true correlation matrix by

$$R = \begin{pmatrix} 1 & \rho_1 & \rho_2 \\ \rho_1 & 1 & \rho_3 \\ \rho_2 & \rho_3 & 1 \end{pmatrix}. \quad (33)$$

We fix the regularization parameter at  $\lambda = 0.5$ . The problem is then to estimate the three parameters  $\rho_1$ ,  $\rho_2$ , and  $\rho_3$ .

With all three methods we find an estimate of

$$\hat{\boldsymbol{\rho}} = \begin{pmatrix} \hat{\rho}_1 \\ \hat{\rho}_2 \\ \hat{\rho}_3 \end{pmatrix} = \begin{pmatrix} 0.8211 \\ 0.1542 \\ -0.1813 \end{pmatrix}. \quad (34)$$

Note that the second element, which is penalized, experiences shrinkage towards zero, as expected. The first element is inflated, while the third is both inflated and changes sign.