



TEXAS TECH UNIVERSITY
Libraries™

THE HODRICK-PRESCOTT FILTER: A SPECIAL CASE OF PENALIZED SPLINE SMOOTHING

The Texas Tech community has made this publication openly available. [Please share](#) how this access benefits you. Your story matters to us.

Citation	Paige, R. L., & Trindade, A. A. (2010). The Hodrick-Prescott Filter: A special case of penalized spline smoothing. <i>Electronic Journal of Statistics</i> , 4, 856-874, 819. https://doi.org/10.1214/10-EJS570
Citable Link	https://hdl.handle.net/2346/89336
Terms of Use	CC BY 4.0

Title page template design credit to [Harvard DASH](#).

The Hodrick-Prescott Filter: A special case of penalized spline smoothing

Robert L. Paige*

*Missouri University of Science and Technology
Department of Mathematics & Statistics
Rolla, MO 65409, U.S.A.
e-mail: paigero@mst.edu*

and

A. Alexandre Trindade^{†‡}

*Texas Tech University
Department of Mathematics & Statistics
Lubbock, TX 79409, U.S.A.
e-mail: alex.trindade@ttu.edu*

Abstract: We prove that the Hodrick-Prescott Filter (HPF), a commonly used method for smoothing econometric time series, is a special case of a linear penalized spline model with knots placed at all observed time points (except the first and last) and uncorrelated residuals. This equivalence then furnishes a rich variety of existing data-driven parameter estimation methods, particularly restricted maximum likelihood (REML) and generalized cross-validation (GCV). This has profound implications for users of HPF who have hitherto typically relied on subjective choice, rather than estimation, for the smoothing parameter. By viewing estimates as roots of an appropriate quadratic estimating equation, we also present a new approach for constructing confidence intervals for the smoothing parameter. The method is akin to a parametric bootstrap where Monte Carlo simulation is replaced by saddlepoint approximation, and provides a fast and accurate alternative to exact methods, when they exist, e.g. REML. More importantly, it is also the only computationally feasible method when no other methods, exact or otherwise, exist, e.g. GCV. The methodology is demonstrated on the Gross National Product (GNP) series originally analyzed by Hodrick and Prescott (1997). With proper attention paid to residual correlation structure, we show that REML-based estimation delivers an appropriate smooth for both the GNP series and its returns.

AMS 2000 subject classifications: Primary 62F25; secondary 62G09.

Keywords and phrases: Semiparametric model, parametric bootstrap confidence interval, saddlepoint approximation, econometric smoothing, gross national product.

Received April 2010.

*Supported in part by National Security Agency Grant Number H98230-09-1-0071.

†Supported in part by National Security Agency Grant Number H98230-08-1-0071.

‡Corresponding author.

1. Introduction

Penalized spline models provide a powerful tool for nonparametric smoothing. The essence of the idea is to parametrically represent the nonparametric component as a function of the explanatory variables via a family of flexible basis functions (splines). Fitted values are chosen to minimize the residual sum of squares, but with the addition of a term penalizing the smoothness of the fit (the smoothing parameter). The origins of the method date back to, among others, Parker and Rice (1985), who comment on the use of smoothing with (penalized) splines as an alternative to smoothing splines. Eilers and Marx (1996) defined a closely associated class of P-splines based on B-spline bases. Brumback, Ruppert and Wand (1999) introduce penalized splines on truncated polynomial bases, and comment on an equivalent representation as a normal linear mixed model. These *semiparametric* models have steadily gained in popularity following the monographs by Eubank (1999), and Ruppert, Wand and Carroll (2003). See also Ruppert, Wand and Carroll (2009) for an updated view of the subject.

One important reason for this popularity is the useful feature that the penalized spline can be cast as a linear mixed model, resulting in meaningful expressions for parameters as best linear unbiased predictors. In this framework, the smoothing parameter is expressible as a ratio of variance components. Under the usual Gaussian assumptions, standard mixed model software can be routinely used for estimation, a fact which opens the door for widespread usage. In fact, penalized spline smoothing and its variants are now fairly common in the physical and biological sciences.

The field of econometrics has been slow at catching up with these developments. Since the landmark working paper by Hodrick and Prescott (1981) who proposed a method to extract the trend in econometric time series, the field has hardly looked elsewhere for alternative smoothers. Following publication of this work (Hodrick and Prescott, 1997), the so-called *Hodrick-Prescott Filter* (HPF) is now routinely used, and has consequently been incorporated as a package into E-views and standard statistical software such as SAS, Matlab, and R. One problematic issue with implementation of HPF has been the fact that a subjective choice must be made for the value of the smoothing parameter (α), given that the method minimizes residual sum of squares with a penalty on roughness. Arguing heuristically, Hodrick and Prescott (1981) originally proposed the value of $\alpha = 1,600$ for quarterly data, a suggestion that has heretofore been unquestioningly adhered to by many practitioners.

Over the years, several adjustments and generalizations have been made to the classical HPF. King and Rebelo (1993) discussed HPF from the perspective of both time and frequency domain approaches, motivating it as a generalization of exponential smoothing. Harvey and Jaeger (1993) were among the first to warn users of strict adherence to HPF, arguing that this can “. . . lead investigators to report spurious cyclical behaviour . . .”, advocating instead a structural time-series modeling approach. Trimbur (2006) develops a Bayesian generalization of HPF. Harvey and Trimbur (2008) again re-cast HPF in the state-space

model framework, and by analyzing it in the frequency domain consider the effects of inappropriate smoothing and changing the observation interval. Schlicht (2008) adapts HPF for series containing structural breaks or missing data. Kim, Koh, Boyd, and Gorinevsky (2009) propose a version of HPF based on minimizing absolute errors.

Rather fewer papers have focused on developing methods for estimating the smoothing parameter, thus alleviating the burden of manual tuning. As far as we can determine, Schlicht (2005), and Dermoune, Djehiche, and Rahmania (2008) are the only real attempts in this direction. Both cast the problem as a state-space model, and under the assumption of independent Gaussian errors derive different estimators. Schlicht (2005) also proposes a method of moments estimator. Greiner (2009) seems to be the first paper promoting the penalized spline regression methodology in econometrics, thus veering away from automatic use of HPF.

The primary goal of this paper is to demonstrate that HPF is in fact a special case of penalized spline smoothing. To the best of our knowledge this connection has not yet been made, although Schlicht (2005) comes very close. Since a penalized spline can be cast as a linear mixed model, this then endows HPF with a gamut of data-driven estimation methods. We believe these to be hitherto unknown facts, and therefore aim to connect two diverging lines of research: the penalized splines literature, currently numbering over 200 citations, and the (few) estimation methods recently proposed for HPF, principally the papers by Schlicht (2005) and Dermoune, Djehiche, and Rahmania (2008).

While point estimation is straightforward via maximum likelihood (ML), restricted maximum likelihood (REML), Akaike's information criterion (AIC), or generalized cross-validation (GCV), testing or confidence interval construction for the all-important smoothing parameter in a penalized spline model is more problematic. Crainiceanu and Ruppert (2004), and Crainiceanu, Ruppert, Claeskens, and Wand (2005), develop both exact and asymptotic likelihood ratio tests for ML and REML-based inference. The exact tests use spectral decompositions as the basis for fast simulation algorithms. By inverting these tests, a grid search allows for the computation of corresponding confidence intervals. Because of the substantial point mass at zero, Crainiceanu and Ruppert (2004) report that the usual asymptotic χ^2 distribution is a poor approximation in small samples.

However fast these exact methods for ML and REML may be, Paige and Trindade (2010) report that they are still relatively slow, with computation times of the order of hours to obtain a single interval. Furthermore, there are no known exact tests for AIC and GCV-based inference. The secondary goal of this paper is therefore to further investigate the performance of a faster and more general method to construct confidence intervals for the smoothing parameter recently proposed by Paige and Trindade (2010), which is applicable under a variety of criteria such as ML, REML, AIC, and GCV. This saddlepoint-based bootstrap (SPBB) approach pivots a saddlepoint approximation to the distribution function of the estimator, and as such is akin to a bootstrap percentile method, where simulation is replaced by fast and accurate saddlepoint approx-

imation. This decreases confidence interval computation time down to minutes, rather than hours.

In the present paper it will be further demonstrated that the SPBB method not only competes well with exact methods, e.g. ML and REML, but also offers a computationally feasible alternative where no exact methods exist, e.g. GCV, AIC, and REML with correlated errors, while also delivering a performance that is nearly exact. The determination of an appropriate residual error correlation structure for penalized spline models is an important issue if the data are of a temporal nature (as in econometrics); see for example Opsomer, Wang, and Yang (2001), and Krivobokova and Kauermann (2007). In this regard we will follow closely the findings of Krivobokova and Kauermann (2007) concerning the superiority of REML-based inference.

The remainder of the paper is organized as follows. An overview of penalized spline models is provided in Section 2, where we highlight the linear mixed model formulation and consequent representation of the smoothing parameter as a ratio of variances. Section 3 presents a unified view of estimators for the smoothing parameter as roots of equations that are quadratic forms in normal random variables. In this context, saddlepoint-based bootstrap inference is possible, and we outline the method. The Hodrick-Prescott Filter (HPF) and details of its equivalent formulation as a penalized spline model constitutes the subject of Section 4. We conclude in Section 5 by revisiting the U.S. gross national product series originally analyzed by Hodrick and Prescott (1997), and comparing a variety of data-driven penalized spline models to the classical HPF.

2. Penalized spline models

This section provides a concise summary of penalized spline regression. A more detailed introduction to this material can be found in say, Ruppert *et al.* (2003), and is needed here in order to make a seamless connection with HPF. For a vector of responses $\mathbf{y} = [y_1, \dots, y_n]'$, consider the model

$$y_i = \mu(x_i) + \varepsilon_i \quad i = 1, \dots, n, \tag{2.1}$$

where the mean $\mu(x)$ is a function of the explanatory variable x , and $\boldsymbol{\varepsilon} = [\varepsilon_1, \dots, \varepsilon_n]'$ $\sim (\mathbf{0}, \Sigma_\varepsilon)$ is a vector of zero-mean disturbances with covariance matrix $\Sigma_\varepsilon = \sigma_\varepsilon^2 R$ and correlation matrix R .

In semiparametric regression, a standard choice for $\mu(x)$ is to represent it as $B(x)\boldsymbol{\theta}$, for some sufficiently flexible spline basis $B(x)$ depending on knots $\{\kappa_1, \dots, \kappa_K\}$, and a vector of coefficients $\boldsymbol{\theta}$. A common choice of basis functions is the penalized spline of degree p with respect to the truncated polynomial basis $B(x) = [1, x, \dots, x^p, (x - \kappa_1)_+^p, \dots, (x - \kappa_K)_+^p]$. Defining the coefficient vector $\boldsymbol{\theta} = [\boldsymbol{\beta}, \mathbf{u}]' = [\beta_0, \dots, \beta_p, u_1, \dots, u_K]'$, this leads to the representation for the mean function

$$\mu(x_i) = \beta_0 + \beta_1 x_i + \dots + \beta_p x_i^p + \sum_{k=1}^K u_k (x_i - \kappa_k)_+^p, \tag{2.2}$$

where for any number a , $(a)_+^p$ is equal to a^p if a is positive and zero otherwise, and the knots $\{\kappa_k\}$ are appropriately chosen over the domain of x . With estimation of $\boldsymbol{\theta}$ in mind, (2.1) can then be written in matrix form as

$$\mathbf{y} = X\boldsymbol{\beta} + Z\mathbf{u} + \boldsymbol{\varepsilon} \equiv B\boldsymbol{\theta} + \boldsymbol{\varepsilon}, \tag{2.3}$$

where $B = [X, Z]$, and the design matrices X and Z have $[1, x_i, \dots, x_i^p]$ and $[(x_i - \kappa_1)_+^p, \dots, (x_i - \kappa_K)_+^p]$, respectively, as their i th rows.

In this paradigm $\boldsymbol{\theta}$ is estimated by minimizing the penalized sum of squared errors criterion

$$\hat{\boldsymbol{\theta}}_{\text{PS}} = \arg \min_{\boldsymbol{\theta}} \{(\mathbf{y} - B\boldsymbol{\theta})'R^{-1}(\mathbf{y} - B\boldsymbol{\theta}) + \alpha\boldsymbol{\theta}'D\boldsymbol{\theta}\}, \tag{2.4}$$

where $\alpha = \lambda^{2p} > 0$ is a smoothing parameter controlling the balance between fidelity to the data ($\alpha \approx 0$) and smoothness of the fit ($\alpha \rightarrow \infty$), and D is an appropriately chosen non-negative definite penalty matrix. This leads to the fitted values

$$\hat{\boldsymbol{\theta}}_{\text{PS}} = (B'R^{-1}B + \alpha D)^{-1} B'R^{-1}\mathbf{y}, \quad \text{and} \quad \hat{\mathbf{y}}_{\text{PS}} = B\hat{\boldsymbol{\theta}}_{\text{PS}} \equiv S_{\alpha}\mathbf{y}, \tag{2.5}$$

where $S_{\alpha} = B(B'R^{-1}B + \alpha D)^{-1} B'R^{-1}$ is the smoothing matrix.

Brumback *et al.* (1999) discuss how a penalized spline can be represented as a linear mixed model (LMM). This involves treating both \mathbf{u} and $\boldsymbol{\varepsilon}$ in (2.3) as random vectors

$$\mathbb{E} \begin{bmatrix} \mathbf{u} \\ \boldsymbol{\varepsilon} \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \end{bmatrix}, \quad \text{Cov} \begin{bmatrix} \mathbf{u} \\ \boldsymbol{\varepsilon} \end{bmatrix} = \begin{bmatrix} \Sigma_u & \mathbf{0} \\ \mathbf{0} & \Sigma_{\varepsilon} \end{bmatrix}, \quad \Sigma_u = \sigma_u^2 G, \quad \Sigma_{\varepsilon} = \sigma_{\varepsilon}^2 R. \tag{2.6}$$

The best linear unbiased predictor (BLUP), or posterior Bayes estimate, of \mathbf{y} in this context is $\tilde{\mathbf{y}} = B\tilde{\boldsymbol{\theta}}$, where

$$\tilde{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} \{(\mathbf{y} - B\boldsymbol{\theta})'\Sigma_{\varepsilon}^{-1}(\mathbf{y} - B\boldsymbol{\theta}) + \mathbf{u}'\Sigma_u^{-1}\mathbf{u}\}.$$

If we define $\alpha = \sigma_{\varepsilon}^2/\sigma_u^2$, and let $G = I_K$ and $D = J_{p+1,K}$, where $J_{a,b}$ denotes a diagonal matrix of dimension $a + b$ with the first a elements zero and the last b elements one

$$J_{a,b} = \begin{bmatrix} 0 & 0 \\ 0 & I_b \end{bmatrix}, \tag{2.7}$$

we obtain that $\tilde{\boldsymbol{\theta}} = \hat{\boldsymbol{\theta}}_{\text{PS}}$, i.e. the fitted penalized spline for the mean of \mathbf{y} is precisely the BLUP of \mathbf{y} in the LMM defined above. In particular this means that the optimal values (in the BLUP sense) for the smoothing parameters α and λ are representable in terms of variance components as

$$\lambda = (\sigma_{\varepsilon}^2/\sigma_u^2)^{1/(2p)}, \quad \alpha = \sigma_{\varepsilon}^2/\sigma_u^2 = \lambda^{2p}. \tag{2.8}$$

Since the smoothing parameter α varies between 0 (no smoothing) and ∞ (maximum smoothing), it is difficult to assess how much structure is being imposed on the data for intermediate values. One solution to this is to report the *degrees of freedom of the fit* corresponding to α , $df_{\text{fit}} = \text{tr}(S_{\alpha})$. For a penalized spline model of degree p with K knots it is easily shown that

$$df_{\text{fit}} \rightarrow p + 1 + K, \quad \text{as } \alpha \rightarrow 0, \quad \text{and} \quad df_{\text{fit}} \rightarrow p + 1, \quad \text{as } \alpha \rightarrow \infty.$$

3. Inference for the smoothing parameter

The choice of smoothing parameter greatly affects the quality of the fitted values in a penalized spline model. Ruppert *et al.* (2003) report that the number and location of the knots is in general not as critical a choice. In this section we review some of the more common methods of estimating the smoothing parameter, and present a unified view of these methods where the estimator, $\hat{\alpha}$, is viewed as the root (in α) of a *quadratic estimating equation* in normal random variables (QEE),

$$\mathcal{Q}(\alpha) \equiv \mathbf{y}' A_\alpha \mathbf{y},$$

where \mathbf{y} is a multivariate normal vector, and A_α a conformable matrix depending on α . This leads to the saddlepoint-based bootstrap method of inference for smoothing parameter α introduced by Paige and Trindade (2010). We give only a brief outline of the approach here, but with sufficient detail to permit a straightforward extension to the case of correlated residuals.

The LMM formulation of a penalized spline model suggests that a natural choice for α is given by relation (2.8), provided the variance components are known. In practice, the latter can be estimated via any of the many methods devised in the LMM literature over the years, the most popular being maximum likelihood (ML) and restricted maximum likelihood (REML). (See e.g. Christensen, 1996, Chap. 12, for details on variance component estimation in LMMs.) Thus data-driven selection of the smoothing parameter is straightforward using statistical software, and REML estimates are generally preferable (Harville, 1977).

The REML criterion involves assuming $\mathbf{y}|\mathbf{u} \sim N(X\boldsymbol{\beta} + Z\mathbf{u}, \sigma_\varepsilon^2 R)$, and $\mathbf{u} \sim N(\sigma_u^2 G)$. Reparametrizing with $\sigma_u^2 = \sigma_\varepsilon^2/\alpha$ and partial differentiation with respect to σ_ε^2 and $\boldsymbol{\beta}$, leads to a profile -2 REML log-likelihood to be minimized in α and R ; see for example equation (6) in Krivobokova and Kauermann (2007). When both $G = I_K$ and $R = I_n$, Paige and Trindade (2010) show that differentiation of this criterion leads to a QEE for REML inference on α . A simple generalization of these calculations leads to the following QEE for a given arbitrary correlation R

$$\mathcal{Q}_{\text{REML}}(\alpha) = (n - p) \dot{\sigma}_\alpha^2 + \left\{ -\text{tr} \left[V_\alpha \dot{V}_\alpha^{-1} \right] + \text{tr} \left[(X' V_\alpha^{-1} X)^{-1} X' \dot{V}_\alpha^{-1} X \right] \right\} \delta_\alpha^2, \tag{3.1}$$

where $V_\alpha = R + ZZ'/\alpha$, $\hat{\sigma}_\alpha^2 = (\mathbf{y} - X\hat{\boldsymbol{\beta}}_\alpha)' V_\alpha^{-1} (\mathbf{y} - X\hat{\boldsymbol{\beta}}_\alpha) / (n - p)$, $\hat{\boldsymbol{\beta}}_\alpha = (X' V_\alpha^{-1} X)^{-1} X' V_\alpha^{-1} \mathbf{y}$, and V_α^{-1} and $\dot{\sigma}_\alpha^2$ denote first derivatives with respect to α . (Details of this calculation are provided in Appendix A.) This REML QEE can thus be used to make inference about the BLUP-optimal penalized spline smoothing parameter suggested by the ratio of variances in (2.8).

Other methods for estimating α in a penalized spline model also typically involve the minimization of some criterion function. Two of the more common include Generalized Cross-Validation (GCV), and Akaike's Information Criterion (AIC). In Paige and Trindade (2010) it was also shown that differentiation

in α of their respective criterion functions yield QEEs, respectively

$$\mathcal{Q}_{\text{GCV}}(\alpha) = \mathbf{y}' \left[(I_n - S_\alpha) \dot{S}_\alpha \{1 - n^{-1} \text{tr}(S_\alpha)\} - n^{-1} (I_n - S_\alpha)^2 \text{tr}(S_\alpha) \right] \mathbf{y}, \tag{3.2}$$

and

$$\mathcal{Q}_{\text{AIC}}(\alpha) = \mathbf{y}' \left[(I_n - S_\alpha) \dot{S}_\alpha - n^{-1} (I_n - S_\alpha)^2 \text{tr}(S_\alpha) \right] \mathbf{y}, \tag{3.3}$$

where S_α is the smoothing matrix defined in (2.5).

In Paige, Trindade, and Fernando (2009) we proposed an easy to implement parametric bootstrap percentile method of confidence interval (c.i.) construction for a generic model parameter, α , by indirectly saddlepoint approximating the distribution of an estimator that is representable as the root of a QEE. We termed this the *saddlepoint-based bootstrap* (SPBB) method. Adaptation of SPBB for penalized spline inference is discussed in Paige and Trindade (2010), and the method can therefore be used for making (approximate) inference on the smoothing parameter α via any one of the above criteria: ML, REML, AIC, or GCV.

As described in Paige and Trindade (2010), *exact* finite sample inference on α is also possible by inverting either the likelihood ratio test (LRT) or restricted likelihood ratio test (RLRT), corresponding to ML and REML based inference, respectively (Crainiceanu *et al.*, 2005). This method (which is applicable only when $R = I_n$) is computationally intensive and, like SPBB, involves performing a grid search for the endpoints of the c.i. At each grid value, a large sample must be drawn from the exact distribution of the LRT/RLRT statistic in order to (empirically) calculate the p-value. Note however that the SPBB method provides a quite general solution to the problem of inference, and in particular is able to furnish a c.i. when no (computationally feasible) competing method, exact or otherwise, exists, for example GCV and AIC.

Krivobokova and Kauermann (2007) report that REML is less sensitive to misspecification of the residual correlation structure than either AIC or GCV. In addition, REML permits a more accurate assessment of this correlation, by careful inspection of residual autocorrelation (ACF) and partial autocorrelation (PACF) plots. For these reasons, and when residual correlation is an issue, we will confine our attention to REML. In such cases, we assume (as do Krivobokova and Kauermann, 2007) that R has a Toeplitz structure, corresponding to the autocorrelation of a stationary autoregressive moving average (ARMA) process. Finally, note that no version of the above exact LRT/RLRT test has been devised when $R \neq I_n$, so that in this case SPBB-REML is the only computationally feasible method to construct a c.i. for α .

4. Connections with the Hodrick-Prescott Filter

As discussed in the Introduction, Hodrick and Prescott (1997) proposed a method to estimate the trend in time series data that is now known as the Hodrick-Prescott Filter (HPF). This smoothing technique views the series y_t as the sum

of a trend μ_t and residual component ε_t ,

$$y_t = \mu_t + \varepsilon_t, \quad t = 1, \dots, n.$$

For smoothing parameter α , the trend is then fitted by minimizing the penalized sum of squares

$$\begin{aligned} \hat{\mathbf{y}}_{\text{HP}} &= \arg \min_{\boldsymbol{\mu} \in \mathbb{R}^n} \left\{ \sum_{t=1}^n (y_t - \mu_t)^2 + \alpha \sum_{t=1}^n (\mu_t - 2\mu_{t-1} + \mu_{t-2})^2 \right\} \\ &= \arg \min_{\boldsymbol{\mu} \in \mathbb{R}^n} \{ (\mathbf{y} - \boldsymbol{\mu})'(\mathbf{y} - \boldsymbol{\mu}) + \alpha \boldsymbol{\mu}' \Delta_2' \Delta_2 \boldsymbol{\mu} \} \\ &= (I_n + \alpha \Delta_2' \Delta_2)^{-1} \mathbf{y}, \end{aligned} \tag{4.1}$$

where $\mathbf{y} = (y_1, \dots, y_n)'$, $\boldsymbol{\mu} = (\mu_1, \dots, \mu_n)'$, and Δ_2 is the $(n - 2) \times n$ second order differencing matrix

$$\Delta_2(i, j) = \begin{cases} 1, & j = i, i = 1, \dots, n - 2, \\ -2, & j = i + 1, i = 1, \dots, n - 2, \\ 1, & j = i + 2, i = 1, \dots, n - 2, \\ 0, & \text{otherwise.} \end{cases}$$

The form of the solution in (4.1) has been heuristically compared to that of a cubic smoothing spline with equi-spaced knots

$$\hat{\mathbf{y}}_{\text{CS}} = \arg \min_{\boldsymbol{\mu} \in \mathbb{R}^n} \left\{ (\mathbf{y} - \boldsymbol{\mu})'(\mathbf{y} - \boldsymbol{\mu}) + \alpha \int_1^n (\mu_t^{(2)})^2 dt \right\}, \tag{4.2}$$

since $\Delta_2 \boldsymbol{\mu}$ is the natural discretized version of the second derivative of μ_t appearing in the integrand of (4.2). (See for example the help file for the function `hpfiler` in Matlab, version 7.8.0, R2009a, Econometrics Toolbox). However as far as we can deduce, an exact equivalence has not firmly been established. In fact, the HPF solution is exactly equivalent to (is a special case of) the solution (2.4) obtained from a penalized spline of degree $p = 1$, equi-spaced knots at points $2, \dots, n - 1$, $R = I_n$, and $D = J_{2, n-2}$ as defined by (2.7). We state this fact in the following theorem, and prove it in Appendix B.

Theorem 4.1. *For a time series $\mathbf{y} = (y_1, \dots, y_n)'$ observed at equi-spaced time points $t = 1, \dots, n \geq 4$, consider a penalized smoothing spline model of degree $p = 1$, explanatory variable time $x_t = t$, and $K = n - 2$ knots at $\{\kappa_k = k\}$, $k = 2, \dots, n - 1$, given by*

$$y_t = \beta_0 + \beta_1 t + \sum_{k=2}^{n-1} u_k (t - k)_+ + \varepsilon_t, \quad t = 1, \dots, n.$$

If the parameter vector $\boldsymbol{\theta} = [\boldsymbol{\beta}', \mathbf{u}']' = [\beta_0, \beta_1, u_2, \dots, u_{n-1}]'$ is estimated by minimizing

$$\hat{\boldsymbol{\theta}}_{\text{PS}} = \arg \min_{\boldsymbol{\theta}} \{ (\mathbf{y} - B_{\text{PS}} \boldsymbol{\theta})'(\mathbf{y} - B_{\text{PS}} \boldsymbol{\theta}) + \alpha \mathbf{u}' \mathbf{u} \},$$

where the model is expressible in LMM matrix form as

$$\mathbf{y} = X_{\text{PS}}\boldsymbol{\beta} + Z_{\text{PS}}\mathbf{u} + \boldsymbol{\varepsilon} \equiv B_{\text{PS}}\boldsymbol{\theta} + \boldsymbol{\varepsilon},$$

with $\boldsymbol{\varepsilon} \sim (\mathbf{0}, \sigma_\varepsilon^2 I_n)$, $\mathbf{u} \sim (\mathbf{0}, \sigma_\varepsilon^2 I_K/\alpha)$, $B_{\text{PS}} = [X_{\text{PS}}, Z_{\text{PS}}]$, and

$$X_{\text{PS}} = \begin{bmatrix} 1 & 1 \\ \vdots & \vdots \\ 1 & n \end{bmatrix}, \quad Z_{\text{PS}} = \begin{bmatrix} 0 & 0 & \cdots & 0 \\ 0 & 0 & \cdots & 0 \\ 1 & 0 & \cdots & 0 \\ 2 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ n-2 & n-3 & \cdots & 1 \end{bmatrix},$$

then the penalized spline solution thus obtained,

$$\hat{\mathbf{y}}_{\text{PS}} = B_{\text{PS}}\hat{\boldsymbol{\theta}}_{\text{PS}} = B_{\text{PS}}(B'_{\text{PS}}B_{\text{PS}} + \alpha J_{2,n-2})^{-1}B'_{\text{PS}}\mathbf{y},$$

where $J_{2,n-2}$ is as defined in (2.7) with $a = 2$ and $b = n - 2$, coincides with the HPF solution in (4.1), i.e. $\hat{\mathbf{y}}_{\text{PS}} = \hat{\mathbf{y}}_{\text{HP}}$.

In Appendix C we provide some R code for fitting these models using two different packages: `SemiPar` for penalized splines, and `mFilter` designed specifically for HPF. The reader can verify here that the fit from `mFilter` coincides with that from `SemiPar` with the degree and knots as specified by Theorem 4.1. Note that these packages automatically assume $R = I_n$ and $G = I_K$, as is needed for equivalence of the penalized spline and HPF solutions.

Hodrick and Prescott (1997) came close to such a result, noting that under the assumption of normal identically and independently distributed random trend and residual components with variances σ_u^2 and σ_ε^2 respectively, the conditional expectation of the trend given the observations would solve (4.1) with $\alpha = \sigma_\varepsilon^2/\sigma_u^2$, the inverse of the *signal-to-noise ratio*. But as we note in the proof, Schlicht (2005) is the first to make the connection that HPF can be formulated as a LMM. However, Schlicht (2005) calls this a “stochastic model” and does not seem to be aware that it is in fact a LMM. Apart from the usual maximum (Gaussian) likelihood, he proposes also a method of moments estimate for the variance components. A succession of econometrics papers have built on this, proposing new methods of estimation, e.g. Dermoune *et al.* (2008). However, there is a vast existing body of work on the estimation of fixed and random components in LMMs, see for example Christensen (1996).

Finally, we note that HPF is a type of polynomial mixed model spline; a *penalized spline*. Two other important polynomial mixed model splines that have been proposed in the literature are *smoothing splines*, and *P-splines*. These are similar in form; all involve a parametric part that models the mean function as a linear combination of basis functions and a vector of parameters, and a non-parametric part that penalizes some measure of smoothness of the fitted mean function. Welham, Cullis, Kenward, and Thompson (2007) show that there is a connection between these three types of splines, so that the solution obtained from one type with a given penalty, is identical to that for another type with an appropriately transformed penalty.

5. Application and Monte Carlo study

We consider the quarterly United States Gross National Product (GNP) for the period 1947.1 to 1993.4 (billions of dollars). The (seasonally adjusted) data, available from the Bureau of Economic Analysis of the U.S. Department of Commerce, was analyzed in the landmark paper by Hodrick and Prescott (1997), who suggested the default choice of $\alpha = 1,600$ for smoothing quarterly data. As is typical of macroeconomic series, the time series considered, $\{y_t\}$, are actually the natural logarithms of GNP (log GNP), and consists of 189 observations. The first differences, $\{y_t - y_{t-1}\}$, the *log returns*, will then correspond to a growth rate. The overwhelming linear trend in log GNP obscures any underlying fluctuations, and, as did Hodrick and Prescott (1997), we will remove it by fitting the penalized spline model of Theorem 4.1 with $\alpha = \infty$, corresponding to a linear least-squares fit. Henceforth, all allusions to log GNP correspond to the resulting detrended data, which show a clear cyclical pattern, as was already noted by Hodrick and Prescott (1997) who warned “against interpreting...as a cycle of long duration”. The two series are displayed in the top panels of Figure 1, and can be viewed as two different starting points for plausibly stationary time series modeling.

As discussed in Section 4, Hodrick and Prescott (1997) proposed a method to smooth such data that is now known as HPF, and a penalized spline model with degree and knots as specified by Theorem 4.1 is therefore applicable. In this context any one of the estimation methods for the smoothing parameter α can be used, and we will focus on REML and GCV, providing data-driven alternatives to the $\alpha = 1,600$ choice (which we will term the *HPF smooth*). We will also construct corresponding c.i.’s for the point estimates of α , using both exact and SPBB methods in the case of REML with uncorrelated residuals, and SPBB only in the case of REML with correlated residuals or GCV (since no known exact method exists in these cases). Finally, we will shed some light on the quality of the results by simulating from the model fitted to the log returns.

Analysis of smoothed GNP. From the ACF plots displayed in the bottom panels of Figure 1, it is apparent that taking the autocorrelation structure of the residuals (R) into account will be much more important in log GNP than in the returns. For log GNP, the REML and GCV methods with uncorrelated residuals (REML-IID, GCV-IID) result in essentially identical estimates, e.g. $\hat{\alpha}_{\text{REML-IID}} = 0.16$. When implemented on an Intel Xeon 3.00GHz processor with 5.16GB of RAM, the exact method of Crainiceanu *et al.* (2005) took approximately 8 hours to produce a 95% Exact-REML-IID c.i., whereas the SPBB-REML-IID and SPBB-GCV-IID c.i.’s were determined in approximately 10 minutes. Results are summarized in Table 1, where we omit SPBB-REML-IID since it is similar to Exact-REML-IID. The resulting smoothed fits appear in Figure 2, where it is apparent that REML-IID (and the virtually identical GCV-IID) severely underfits the data. As expected, the HPF smooth is remarkably good.

Inspection of residual ACF/PACF plots from the REML-IID fit reveals, as suspected, substantial autocorrelation. This is in line with the findings of




FIG 1. Time series and sample autocorrelation function (ACF) plots of the detrended log GNP data (left panels) and log returns of GNP (right panels).

Krivobokova and Kauermann (2007), that data-driven methods such as GCV, AIC, and REML tend not to perform well unless this correlation is taken into account (although REML is less sensitive to misspecification thereof). As they suggest, we re-fit via REML with a plausible correlation structure gleaned by inspection of the residual ACF/PACF. In fact, the ACF/PACF of the original log GNP suggests an integrated model such as ARIMA(2,1,0). Such a model for the data may also fit the residuals, but may not be parsimonious. A search over all low-order ARMA(p, q) models ($p + q \leq 5$) in order to identify a simple model for R , reveals that an AR(3) delivers the lowest value of AIC (for the overall fitted LMM; see Krivobokova and Kauermann, 2007, for details on how to fit this with the R function `lme`). The corresponding fitted value of α (at 337.3 now much larger than in the IID case) and its accompanying SPBB-REML-AR(3) c.i. is shown in Table 1. The resulting smoothed data appear in Figure 2. The REML-IID (and GCV-IID) methods produce virtually no smoothing at all, while REML-AR(3) is visibly indistinguishable from the HPF smooth. Note that the amount of structure imposed on log GNP by these models corresponds to minimum and maximum possible values of df_{fit} between $p + 1 = 2$ and $p + 1 + K = 189$.

The results are different for the (qualitatively very different) log returns of GNP which exhibits little autocorrelation. To fit the noise, we limited our search to low-order ARMA(p, q) models with $p + q \leq 1$. An AR(1) delivered the lowest value of AIC, but the smoothed results showed no visibly distinguishable difference from IID fits. We therefore proceeded with $R = I_n$ throughout. The

TABLE 1

Comparison of Exact vs. saddlepoint-based bootstrap (SPBB) point and 95% interval estimates for the smoothing parameter α based on a linear penalized spline model fitted to each of the log GNP and log returns of GNP datasets. Optimality criteria used were restricted maximum likelihood (REML) and generalized cross-validation (GCV), assuming uncorrelated residuals (IID). SPBB-REML-AR(3) assumes an AR(3) correlation structure for the residuals. The right-most column gives equivalent values for α on the df_{fit} scale defined in Section 2

Dataset	Method	α 95% Estimate	df_{fit} 95% Estimate
		(lower, point, upper)	(lower, point, upper)
Log GNP	Exact-REML-IID	(0.07, 0.16, 0.32)	(101.3, 120.7, 144.9)
	SPBB-GCV-IID	(0.04, 0.15, 0.59)	(85.9, 123.0, 159.4)
	SPBB-REML-AR(3)	(97.6, 337.3, 871.4)	(10.5, 13.4, 19.8)
Log Returns of GNP	Exact-REML-IID	(8, 595, 65,302, 330,436)	(3.8, 5.2, 7.9)
	SPBB-REML-IID	(5, 572, 65,302, 463,185)	(3.5, 5.2, 8.7)
	SPBB-GCV-IID	(6,022, 79,164, ∞)	(2.0, 5.0, 8.6)

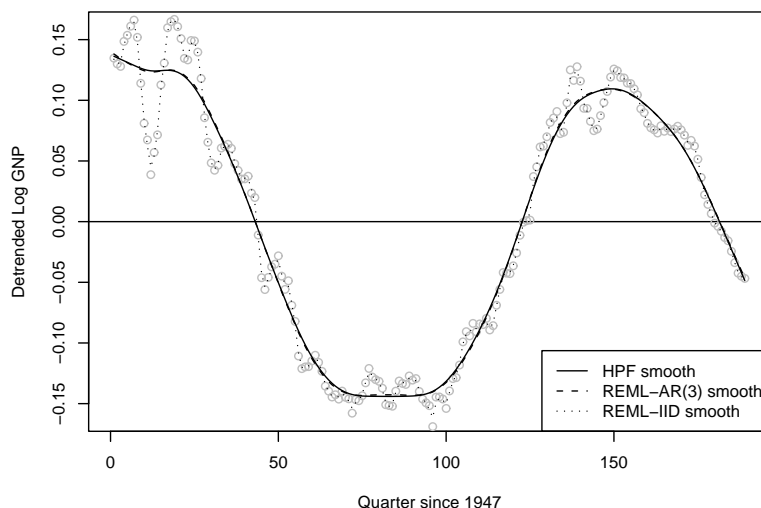


FIG 2. The detrended log GNP data and its smoothed values, with smoothing parameter estimated via HPF (solid), REML-AR(3) (dashed, indistinguishable from HPF), and REML-IID (dotted, indistinguishable from GCV-IID).

α estimates in Table 1 clearly suggest that substantially more smoothing is needed to adequately reveal the trend in the log returns. This is reflected in much smaller df_{fit} values. Figure 3 shows that now it is GCV and REML (indistinguishable) that deliver the most smoothing, and HPF the least. Note that the HPF smooth $\alpha = 1,600$ setting corresponds to df_{fit} values of 11.6 and 11.5, respectively, for log GNP and log returns of GNP.

Simulation results from models fitted to GNP. To assess the quality of the exact vs. SPBB c.i.'s for α in the log returns of GNP, we simulated 10^3 replicates from the REML-fitted penalized spline model over the same design




FIG 3. The log returns of GNP data and its smoothed values, with smoothing parameter estimated via GCV-IID (solid), REML-IID (dashed, indistinguishable from GCV-IID), and HPF (dotted).

TABLE 2

Empirical comparison of Exact vs. saddlepoint-based bootstrap (SPBB) 95% c.i.'s for the smoothing parameter α based on 1,000 simulations from the penalized spline model fitted via REML to the log returns of GNP. Optimality criteria used were restricted maximum likelihood (REML) and generalized cross-validation (GCV). The underage, coverage, and overage probabilities represent the proportion of simulations where the true α fell below, between, and above the c.i. endpoints, respectively

Estimation Method	Median C.I. Length	Underage Prob.	Coverage Prob.	Overage Prob.
Exact-REML	374,800	0.000	0.981	0.019
SPBB-REML	466,300	0.023	0.977	0.000
SPBB-GCV	∞	0.016	0.984	0.000

points. The c.i. upper bounds were declared infinite if they exceeded 10^9 . The results appear in Table 2. The Exact-REML method seems to yield somewhat shorter intervals than SPBB-REML, although both methods had some infinite c.i.'s (about 7%). All GCV c.i.'s had an infinite upper bound, and 19 of them also had zero as the lower bound. The frequency of c.i.'s with an infinite upper bound suggests that a linear fit (unsmoothed) is plausible with 95% confidence according to the GCV criterion (and gives further credence to the upper bound of $\alpha = \infty$ on Table 1). The similarity in empirical coverage probabilities validates the soundness of the SPBB method, providing a faster alternative to Exact-REML, and furnishing the only computationally feasible method to construct a c.i. in the case of GCV.

Discussion. The main message from these two analyses is that blind usage of HPF for econometric smoothing is unnecessary, and may not always deliver the most appropriate fit, even for the quarterly time series data that it was “designed” for. With the insight provided by Theorem 4.1 of this paper that HPF corresponds to a linear penalized spline with knots placed at time points $2, \dots, n - 1$ and uncorrelated residuals, we advocate a more structured approach by paying close attention to the amount of smoothing (α) and residual autocorrelation structure (R). The choice of degree of polynomial (p) and number of knots (K) is not so crucial (Ruppert *et al.*, 2003), since the asymptotic distribution of a penalized spline does not depend on p or K , provided the latter increases sufficiently rapidly (Li and Ruppert, 2008; Wang, Shen, and Ruppert, 2010). With these choices in place, it is then a straightforward matter to use established data-driven methods for estimating all parameters, including the smoothing parameter. Among these, the use of REML is particularly recommended, given the findings by Krivobokova and Kauermann (2007) that it is less sensitive to misspecification of R than either GCV or AIC (although more research needs to be done on this topic). Finally c.i.’s can, if desired, be constructed using SPBB, thus giving a range of plausible smoothed fits.

Appendix A: Derivation of REML QEE (3.1)

From equation (6) in Krivobokova and Kauermann (2007), the $-2 \log$ REML criterion to be minimized in α and R (we continue to assume $G = I_K$ in order to make the standard connection between the penalized spline solution and the BLUP) is

$$L_\alpha(\alpha, R) = (n - p) \log(\hat{\sigma}_\alpha^2) + \log |V_\alpha| + \log |X'V_\alpha^{-1}X|.$$

Standard matrix differentiation with respect to α then leads to equation (3.1), when the following basic expressions are used:

- $\dot{V}_\alpha = -ZZ'/\alpha^2, \quad \dot{V}_\alpha^{-1} = -V_\alpha^{-1}\dot{V}_\alpha V_\alpha^{-1}.$
- $\dot{W}_\alpha = X(X'V_\alpha^{-1}X)^{-1}X'V_\alpha^{-1}, \quad \dot{W}_\alpha = X(X'V_\alpha^{-1}X)^{-1}X'V_\alpha^{-1}.$
- $\dot{\hat{\sigma}}_\alpha^2 = \mathbf{y}' \left[(I_n - W_\alpha)' \dot{V}_\alpha^{-1} (I_n - W_\alpha) - 2\dot{W}_\alpha' V_\alpha^{-1} (I_n - W_\alpha) \right] \mathbf{y} / (n - p).$

Appendix B: Proof of Theorem 4.1

Schlicht (2005) makes the fundamental connection that HPF can be formulated as the LMM

$$\mathbf{y} = X_{\text{HP}}\boldsymbol{\beta} + Z_{\text{HP}}\mathbf{u} + \boldsymbol{\varepsilon} \equiv B_{\text{HP}}\boldsymbol{\theta} + \boldsymbol{\varepsilon},$$

with the $n \times (n - 2)$ matrix $Z_{\text{HP}} = \Delta_2'(\Delta_2\Delta_2')^{-1}$, and X_{HP} is any $n \times 2$ matrix satisfying the conditions: (i) $\Delta_2 X_{\text{HP}} = 0$ and (ii) $X_{\text{HP}}' X_{\text{HP}} = I_2$. By then making the BLUP assumptions (2.6), with additionally $[\mathbf{u}', \boldsymbol{\varepsilon}']$ being multivariate normal, $G = I_K$, and $R = I_n$, he notes that an estimate of $\alpha = \sigma_\varepsilon^2/\sigma_u^2$ can thus be obtained via maximum likelihood. However, Schlicht (2005) calls this

a “stochastic model” and does not seem to be aware that it is in fact a LMM. Nevertheless, the solution can be expressed in two equivalent ways: the classical HPF solution

$$\hat{\mathbf{y}}_{\text{HP}} = (I_n + \alpha \Delta_2' \Delta_2)^{-1} \mathbf{y}, \tag{B.1}$$

and, in light of the LMM expression above, it is also the BLUP of \mathbf{y} , so we must have

$$\hat{\mathbf{y}}_{\text{HP}} = B_{\text{HP}}(B_{\text{HP}}' B_{\text{HP}} + \alpha J_{2,n-2})^{-1} B_{\text{HP}}' \mathbf{y}. \tag{B.2}$$

The equivalence of (B.1) and (B.2) is essentially proved in Theorem 1 of Schlicht (2005). Now, the penalized spline solution is independent of the choice of basis functions used, provided the X matrix corresponds to a basis of the same degree and the Z matrix to the same knot locations (Ruppert *et al.*, 2003). Specifically, defining the $n \times n$ invertible matrix $L = B_{\text{PS}}^{-1} B_{\text{HP}}$, and denoting $L^{-T} = (L')^{-1} = (L^{-1})'$, we have

$$\begin{aligned} \hat{\mathbf{y}}_{\text{HP}} &= B_{\text{HP}} [B_{\text{HP}}' B_{\text{HP}} + \alpha J_{2,n-2}]^{-1} B_{\text{HP}}' \mathbf{y} \\ &= B_{\text{PS}} L [L' (B_{\text{PS}}' B_{\text{PS}} + \alpha L^{-T} J_{2,n-2} L^{-1}) L]^{-1} L' B_{\text{PS}}' \mathbf{y} \\ &= B_{\text{PS}} L L^{-1} [B_{\text{PS}}' B_{\text{PS}} + \alpha L^{-T} J_{2,n-2} L^{-1}]^{-1} L^{-T} L' B_{\text{PS}}' \mathbf{y} \\ &= B_{\text{PS}} [B_{\text{PS}}' B_{\text{PS}} + \alpha K_n]^{-1} B_{\text{PS}}' \mathbf{y} \\ &= \dots \\ &= B_{\text{PS}} (B_{\text{PS}}' B_{\text{PS}} + \alpha J_{2,n-2})^{-1} B_{\text{PS}}' \mathbf{y} \\ &= \hat{\mathbf{y}}_{\text{PS}} \end{aligned}$$

Equivalence will then follow if we can show in the intermediate step above that

$$K_n \equiv L^{-T} J_{2,n-2} L^{-1} = J_{2,n-2}. \tag{B.3}$$

To this end, decompose L^{-1} into the following block structure (with dimensions specified by the subscripts)

$$L^{-1} = \begin{bmatrix} A_{2,2} & B_{2,n-2} \\ C_{n-2,2} & D_{n-2,n-2} \end{bmatrix}.$$

We then have that

$$\begin{aligned} K_n &= L^{-T} J_{2,n-2} L^{-1} \\ &= \begin{bmatrix} A_{2,2}' & C_{n-2,2}' \\ B_{2,n-2}' & D_{n-2,n-2}' \end{bmatrix} \begin{bmatrix} 0_{2,2} & 0_{2,n-2} \\ 0_{n-2,2} & I_{n-2} \end{bmatrix} \begin{bmatrix} A_{2,2} & B_{2,n-2} \\ C_{n-2,2} & D_{n-2,n-2} \end{bmatrix} \\ &= \begin{bmatrix} C'C & C'D \\ D'C & D'D \end{bmatrix}, \end{aligned}$$

and thus it suffices to show that (i) $C = 0_{n-2,2}$, and (ii) $D = I_{n-2}$, in order to establish (B.3). To this end, note that B_{HP} and B_{PS}^{-1} have the block structures, $B_{\text{HP}} = [X_{\text{HP}}, Z_{\text{HP}}]$, and

$$B_{\text{PS}}^{-1} = \begin{bmatrix} E_{2,n-2} \\ P_{n-2,2} \end{bmatrix} \equiv \begin{bmatrix} 2 & -1 & 0 & \dots & 0 \\ -1 & 1 & 0 & \dots & 0 \\ & & \Delta_2 & & \end{bmatrix}, \tag{B.4}$$

where we can take the form of $X_{\text{HP}} = [\mathbf{x}_1, \mathbf{x}_2]$ as suggested by Schlicht (2005), with the n -vectors \mathbf{x}_1 and \mathbf{x}_2 as

$$\mathbf{x}_1 = [1, \dots, 1]' / c_1, \quad \text{and} \quad \mathbf{x}_2 = [(1 - (n + 1)/2), \dots, (n - (n + 1)/2)]' / c_2,$$

where the normalizing constants $c_1 = \sqrt{n}$ and $c_2 = \sqrt{3n^3 - 3n}/6$ ensure that $\mathbf{x}_i' \mathbf{x}_i = 1$, $i = 1, 2$. Now note that if $C = 0_{n-2,2}$ and $D = I_{n-2}$ of L^{-1} are as specified in (i) and (ii), then the form of L will also have the same structure for the lower blocks, and vice-versa, i.e.

$$L^{-1} = \begin{bmatrix} A_{2,2} & B_{2,n-2} \\ C_{n-2,2} & D_{n-2,n-2} \end{bmatrix} \iff L = \begin{bmatrix} \tilde{A}_{2,2} & \tilde{B}_{2,n-2} \\ C_{n-2,2} & D_{n-2,n-2} \end{bmatrix},$$

where $\tilde{A}_{2,2}$ and $\tilde{B}_{2,n-2}$ can easily be deduced from standard matrix results (but are not needed here). Consequently, it suffices to show (i) and (ii) for L . The form of $L = B_{\text{PS}}^{-1} B_{\text{HP}}$ is easier to deal with, since B_{PS}^{-1} has the sparse structure noted in (B.4). Therefore, block-multiplying B_{PS}^{-1} and B_{HP} , gives

$$L = B_{\text{PS}}^{-1} B_{\text{HP}} = \begin{bmatrix} E \\ \Delta_2 \end{bmatrix} [X_{\text{HP}}, B_{\text{HP}}] = \begin{bmatrix} EX_{\text{HP}} & EZ_{\text{HP}} \\ \Delta_2 X_{\text{HP}} & \Delta_2 Z_{\text{HP}} \end{bmatrix}.$$

Now it is easy to see that $C = \Delta_2 X_{\text{HP}}$ and $D = \Delta_2 Z_{\text{HP}} = \Delta_2 \Delta_2' (\Delta_2 \Delta_2')^{-1} = I_{n-2}$, so that (ii) follows immediately. To show (i), note that the $(i, 1)$ and $(i, 2)$ entries of C , $i = 1, \dots, n - 2$, are

$$C(i, 1) = (1 - 2 + 1) / c_1 = 0$$

$$C(i, 2) = [2i - (n + 1) - 4(i + 1) + 2(n + 1) + 2(i + 2) - (n + 1)] / (2c_2) = 0.$$

Hence, (i) and (ii) follow, which establishes (B.3), and this proves that $\hat{\mathbf{y}}_{\text{HP}} = \hat{\mathbf{y}}_{\text{PS}}$.

Appendix C: R code

With the vector of responses assigned to \mathbf{y} , the HPF fitted penalized spline for quarterly data ($\alpha = 1,600$) using the `SemiPar` package is obtained with the following code. Note that the `spar` parameter corresponds to $\lambda = \alpha^{1/(2p)}$, as defined in (2.8).

```
library(SemiPar)
n=length(y); x=seq(1,n);
hpf1=spm(y~f(x, basis="trunc.poly", degree=1, knots=seq(2,n-1),
           spar=sqrt(1600)))
```

With the package `mFilter` that is specifically designed for HPF, the corresponding code is as follows.

```
library(mFilter)
hpf2=hpfiler(y, type="lambda", freq=1600)
```

The fitted values from the two fits, `hpf1` and `hpf2`, will be identical.

Acknowledgements

The authors would like to acknowledge the following forms of assistance in this work.

- *Refereeing.* We are indebted to the Editor and Associate Editor for a careful review that led to a much improved paper.
- *Computing.* Dr. P.W. Smith, Professor, Department of Mathematics & Statistics, and Director, High Performance Computing Center, Texas Tech University, for access to additional computing facilities without which the simulations results reported in Section 5 would have taken much longer to obtain. Dr. Sukanta Basu, Department of Geosciences, Texas Tech University, for access to additional software resources that made possible the use of the additional computing facilities. Dr. Ciprian Crainiceanu, Department of Biostatistics, Johns Hopkins University, for sharing his Matlab codes, which permitted a speedier (and more accurate) implementation of the exact method.
- *Funding.* Project sponsored by the National Security Agency (NSA) under Grant Numbers H98230-09-1-0071 (Paige) and H98230-08-1-0071 (Trindade). The United States Government is authorized to reproduce and distribute reprints notwithstanding any copyright notation herein. The authors gratefully acknowledge the support of NSA.

References

- BRUMBACK, B. A., RUPPERT, D. and WAND, M. P. (1999). Comment on “Variable selection and function estimation in additive nonparametric regression using a data-based prior”. *Journal of the American Statistical Association* **94** 794-797. [MR1723272](#)
- CHRISTENSEN, R. (1996). *Plane Answers to Complex Questions: The theory of Linear Models*. Springer, New York. [MR1402692](#)
- CRAINICEANU, C. and RUPPERT, D. (2004). Likelihood ratio tests in linear mixed models with one variance component. *Journal of the Royal Statistical Society, Series B* **66** 165-185. [MR2035765](#)
- CRAINICEANU, C., RUPPERT, D., CLAESKENS, G. and WAND, M. (2005). Exact likelihood ratio tests for penalized splines. *Biometrika* **92** 91-103. [MR2158612](#)
- DERMOUNE, A., DJEHICHE, B. and RAHMANIA, N. (2008). A consistent estimator of the smoothing parameter in the Hodrick-Prescott filter. *Journal of the Japanese Statistical Society* **38** 225-241. [MR2458929](#)
- EILERS, P. H. and MARX, B. D. (1996). Flexible smoothing with B-splines and penalties (with discussion). *Statistical Science* **11** 89-121. [MR1435485](#)
- EUBANK, R. L. (1999). *Nonparametric Regression and Spline Smoothing*. Marcel Dekker, New York. [MR1680784](#)
- GREINER, A. (2009). Estimating penalized spline regressions: Theory and application to economics. *Applied Economics Letters* **16** 1831-1835.

- HARVEY, A. C. and JAEGER, A. (1993). Detrending, stylized facts and the business cycle. *Journal of Applied Econometrics* **8** 231-247.
- HARVEY, A. C. and TRIMBUR, T. M. (2008). Trend estimation and the Hodrick-Prescott filter. *Journal of the Japanese Statistical Society* **38** 41-49. [MR2458316](#)
- HARVILLE, D. A. (1977). Maximum likelihood approaches to variance component estimation and to related problems (with discussion). **72** 320-340. [MR0451550](#)
- HASTIE, T. J., TIBSHIRANI, R. and FRIEDMAN, J. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed. Springer, New York. [MR1851606](#)
- HODRICK, R. J. and PRESCOTT, E. C. (1981). Postwar U.S. business cycles: an empirical investigation Technical Report No. 451, Northwestern University, Center for Mathematical Studies in Economics & Management Science.
- HODRICK, R. J. and PRESCOTT, E. C. (1997). Postwar U.S. business cycles: an empirical investigation. *Journal of Money, Credit, and Banking* **29** 1-16.
- KIM, S. J., KOH, K., BOYD, S. and GORINEVSKY, D. (2009). ℓ_1 trend filtering. *SIAM Review* **51** 339-360. [MR2505584](#)
- KING, R. G. and REBELO, S. T. (1993). Low frequency filtering and real business cycles. *Journal of Economic Dynamics and Control* **17** 207-231.
- KRIVOBOKOVA, T. and KAUEMANN, G. (2007). A Note on penalized spline smoothing with correlated errors. *Journal of the American Statistical Association* **102** 1328-1337. [MR2412553](#)
- LI, Y. and RUPPERT, D. (2008). On the asymptotics of penalized splines. *Biometrika* **95** 415-436. [MR2521591](#)
- OPSOMER, J., WANG, Y. and YANG, Y. (2001). Nonparametric regression with correlated errors. *Statistical Science* **16** 134-153. [MR1861070](#)
- PAIGE, R. L., TRINDADE, A. A. and FERNANDO, P. H. (2009). Saddlepoint-based bootstrap inference for quadratic estimating equations. *Scandinavian Journal of Statistics* **36** 98-111. [MR2508333](#)
- PAIGE, R. L. and TRINDADE, A. A. (2010). Fast and accurate inference for the smoothing parameter in semiparametric models Technical Report, Missouri University of Science and Technology/Texas Tech University.
- PARKER, R. L. and RICE, J. A. (1985). Comments on "Some aspects of the spline smoothing approach to non-parametric regression curve fitting". *Journal of the Royal Statistical Society, Series B* **47** 40-42.
- RUPPERT, D., WAND, M. P. and CARROLL, R. J. (2003). *Semiparametric Regression*. Cambridge, London. [MR1998720](#)
- RUPPERT, D., WAND, M. P. and CARROLL, R. J. (2009). Semiparametric regression during 2003-2007. *Electronic Journal of Statistics* **3** 1193-1256. [MR2566186](#)
- SCHLICHT, E. (2005). Estimating the smoothing parameter in the so-called Hodrick-Prescott filter. *Journal of the Japanese Statistical Society* **35** 99-119. [MR2183502](#)

- SCHLICHT, E. (2008). Trend extraction from time series with structural breaks and missing observations. *Journal of the Japanese Statistical Society* **38** 285-292. [MR2458932](#)
- TRIMBUR, T. M. (2006). Detrending economic time series: A Bayesian generalization of the Hodrick-Prescott filter. *Journal of Forecasting* **25** 247-273. [MR2242141](#)
- WANG, X., SHEN, J. and RUPPERT, D. (2010). Local asymptotics of P-spline smoothing Technical Report, <http://arxiv.org/abs/0912.1824>.
- WELHAM, S. J., CULLIS, B. R., KENWARD, M. G. and THOMPSON, R. (2007). A comparison of mixed model splines for curve fitting. *Australian and New Zealand Journal of Statistics* **49** 1-23. [MR2345406](#)