FEDERAL RESERVE BANK *of* ATLANTA WORKING PAPER SERIES

Fitting a Distribution to Survey Data for the Half-Life of Deviations from PPP

Mark Fisher

Working Paper 2015-15
December 2015

Abstract: This note presents a nonparametric Bayesian approach to fitting a distribution to the survey data provided in Kilian and Zha (2002) regarding the prior for the half-life of deviations from purchasing power parity (PPP). A point mass at infinity is included. The unknown density is represented as an average of shape-restricted Bernstein polynomials, each of which has been skewed according to a preliminary parametric fit. A sparsity prior is adopted for regularization.

JEL classification: C11, C14, F31

Key words: nonparametric Bayesian estimation, Bernstein polynomials, simplex regression, importance sampling, PPP half-life deviations

The author thanks Tao Zha for bringing the paper to his attention. The views expressed here are the author's and not necessarily those of the Federal Reserve Bank of Atlanta or the Federal Reserve System. Any remaining errors are the author's responsibility.

Please address questions regarding content to Mark Fisher, Research Department, Federal Reserve Bank of Atlanta, 1000 Peachtree Street NE, Atlanta, GA 30309-4470, 404-498-8757, mark.fisher@atl.frb.org.

1. Introduction

Kilian and Zha (2002) present results from a survey of economists asking about prior beliefs for the half-life of deviations from purchasing power parity (PPP) for real exchange rates. The survey data are summarized in Table 1 and displayed in Figure 1. The numbers in the table are averages of the responses from 20 economists to a questionaire.[1] The data are composed of $n = 9$ pairs (h_i, y_i), where $y_i = \Pr[h \le h_i]$ and $h_i \in \{1, 2, 3, 4, 5, 6, 10, 20, 40\}$ (measured in years). Using the survey data, the authors estimate what they call a "consensus prior," which they compute through the lens a monthly autoregressive model with 12 lags.

In this note I provide an alternative approach to estimating a smooth distribution from the survey data. I treat the problem as an exercise in Bayesian inference.[2] In particular, I take a Bayesian approach that involves nonparametric regression using Bernstein polynomials subject to shape restrictions.[3] The procedure can be thought of as providing flexible variation around a preliminary parametric fit.

There are two additional novelties regarding the distribution I compute, both of which are related to my own research on PPP.[4] First, I allow for a point mass at infinity. Second, I transform the distribution into a prior for the first-order autoregressive coefficient for annual observations.

2. The model

The model I adopt for the unknown distribution for the half-life h is a mixture of an atom located at infinity and a density over over the positive real line:

$$p(h|\theta_k, w) = \begin{cases} w & h = \infty \\ (1-w)\, f(h|\theta_k) & h \in [0, \infty) \end{cases}, \tag{2.1}$$

where $\Pr[h = \infty] = w$. The density component in (2.1) is itself a mixture — a mixture of basis density functions:

$$f(h|\theta_k) := \sum_{j=1}^{k} \theta_{jk}\, f_{jk}(h), \tag{2.2}$$

where $\theta_k = (\theta_{1k}, \ldots, \theta_{kk})$ and $\theta_k \in \Delta^{k-1}$, the simplex of dimension $k - 1$.

The basis density functions are related to Bernstein polynomials. The idea can be found in Quintana et al. (2009), for example. Let $Q(x)$ denote the cumulative distribution funtion (CDF) for a continuous random variable defined on the real line. Thus $q(x) := Q'(x)$ is the probability density function (PDF). (For the half-life, $Q(x) = 0$ for $x \le 0$.) Define

$$f_{jk}(x) := \mathsf{Beta}\big(Q(x)|j, k-j+1\big)\, q(x), \tag{2.3}$$

[1]The paper refers to "a survey of 22 economists." However, one of the authors confirmed there were only 20 responses.

[2]An approach that is similar in spirit can be found in Gosling et al. (2007).

[3]Fisher (2015) places the approach taken here is the context of what he calls *simplex regression*.

[4]Dwyer and Fisher (2014).

Table 1. Survey prior probabilities for half-life.

	$h \le 1$	$h \le 2$	$h \le 3$	$h \le 4$	$h \le 5$	$h \le 6$	$h \le 10$	$h \le 20$	$h \le 40$	$h > 40$
Percent	4.6	14.1	31.4	49.6	64.0	75.8	83.9	91.0	94.1	5.9

Notes: [This table replicates of Table I in Kilian and Zha (2002).] Average probabilities based on a survey of [20] economists with a professional interest in the PPP question. The survey was conducted by the authors in July and August 1999.

where $1 \le j \le k \in \mathbb{N}$. Note

$$\mathsf{Beta}(x|a,b) = \frac{x^{a-1}\,(1-x)^{b-1}}{B(a,b)}, \tag{2.4}$$

where $B(a,b) = \int_0^1 x^{a-1}\,(1-x)^{b-1}\,dx$ is the beta function. Also note $f_{jk}(x) \ge 0$ for $x \in (-\infty,\infty)$ and

$$\int_{-\infty}^{\infty} f_{jk}(x)\,dx = 1. \tag{2.5}$$

Beta densities with integer coefficients can be interpreted as normalized Bernstein polynomial basis functions. With integer coefficients,

$$\mathsf{Beta}(x|j,k-j+1) = \frac{k!\,x^{j-1}\,(1-x)^{k-j}}{(k-j)!\,(j-1)!}, \tag{2.6}$$

which is a polynomial of degree $k-1$ in x. Bernstein polynomials have a number of useful properties that have led to their use in nonparametric estimations.[5] For example, the "adding-up" property of Bernstein polynomials amounts to

$$\sum_{j=1}^{k} \mathsf{Beta}(x|j,k-j+1) = k. \tag{2.7}$$

This property delivers the following result:

$$\sum_{j=1}^{k} \frac{1}{k}\,f_{jk}(x) = q(x). \tag{2.8}$$

In particular note $f_{11}(x) = q(x)$.

Cumulative distribution function. In order to make contact with the survey data, we will need the cumulative distribution function associated with (2.1). To that end define

$$F(x|\theta_k) := \sum_{j=1}^{k} \theta_{jk}\,F_{jk}(x), \tag{2.9}$$

[5]See, for example, `http://en.wikipedia.org/wiki/Bernstein_polynomial`.

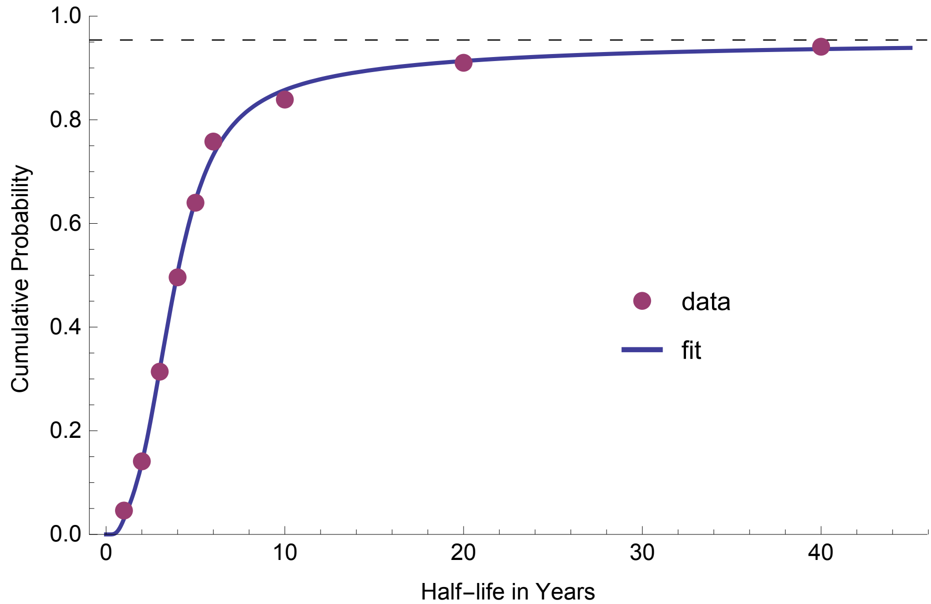

FIGURE 1. The survey data and the survey fit. The fit delivers a 4.6% chance that the half-life is infinite. The dashed line corresponds to the implied asymptote at 0.954.

where

$$\begin{aligned} F_{jk}(x) := \int_{-\infty}^{x} f_{jk}(t)\,dt &= \int_{-\infty}^{x} \mathsf{Beta}\big(Q(t)\,|\,j, k-j+1\big)\, q(t)\,dt \\ &= \int_{0}^{Q(x)} \mathsf{Beta}(t\,|\,j, k-j+1)\,dt \\ &= I_{Q(x)}(j, k-j+1), \end{aligned} \tag{2.10}$$

where $I_x(a,b)$ is the *regularized incomplete beta function*. The adding-up condition (2.8) implies

$$\sum_{j=1}^{k} \frac{1}{k}\, F_{jk}(x) = Q(x). \tag{2.11}$$

With (2.8) and (2.11) in mind, I refer to Q as the *centering function*. The centering function provides location and scale for the fit. Deviation of the weights θ_k from uniform (i.e., deviations from $\theta_{jk} = 1/k$) allow for variation around the centering function. Larger values of k provide greater flexibility.

Degree elevation. One of the properties of Bernstein polynomials is that of *degree elevation*, by which lower-degree polynomials can be represented exactly as higher degree polynomials. Degree elevation is useful for combing models with different values of k.

Applied to mixtures of Beta distributions, degree elevation implies that every mixture of order k_0 can be represented as a mixture of $k_1 > k_0$. Define the $k_1 \times k_0$ matrix

$$A^{k_1,k_0} := A^{k_1,k_1-1}\,A^{k_1-1,k_1-2}\cdots A^{k_0+1,k_0}, \tag{2.12}$$

where the $(k \times k-1)$ matrix $A^{k,k-1}$ is characterized by

$$A^{k,k-1}_{ij} = \begin{cases} 1-(j/k) & j=i \\ j/k & j=i-1 \\ 0 & \text{otherwise} \end{cases}. \tag{2.13}$$

In addition, define the row vector

$$f_k(x) := \big(f_{k1}(x),\ldots,f_{kk}(x)\big). \tag{2.14}$$

One may confirm that

$$f_{k_1}(x)A^{k_1,k_0} \equiv f_{k_0}(x). \tag{2.15}$$

As a consequence (and treating θ_k as a column vector),

$$\begin{aligned} f(x|\theta_{k_0}) = f_{k_0}(x)\,\theta_{k_0} &= \big(f_{k_1}(x)A^{k_1,k_0}\big)\theta_{k_0} \\ &= f_{k_1}(x)\big(A^{k_1,k_0}\theta_{k_0}\big) = f_{k_1}(x)\,\theta_{k_1} = f(x|\theta_{k_1}), \end{aligned} \tag{2.16}$$

where $\theta_{k_1} = A^{k_1,k_0}\theta_{k_0}$. For example, $A^{k,1}\,\theta_1 = (1/k,\ldots,1/k)^\top$.

Reparameterization. It is convenient to reparameterize the model as follows.

Fix $K \ge k$ and let

$$\phi = (1-w)\,A^{K,k}\,\theta_k. \tag{2.17}$$

The model [see (2.1)] can be reexpressed as

$$p(h|\phi) = \begin{cases} 1-\sum_{j=1}^K \phi_j & h=\infty \\ f(h|\phi) & h \in [0,\infty) \end{cases}, \tag{2.18}$$

since

$$1-\sum_{j=1}^K \phi_j = w \qquad \text{and} \qquad f(h|\phi) \equiv (1-w)\,f(h|\theta_k). \tag{2.19}$$

I will use (2.18) for estimation.

3. Bayesian approach to estimation

The goal is to compute the distribution $p(h|y)$ for h conditional on $y = (y_1,\ldots,y_n)$ where the uncertainty regarding the latent variable ϕ has been integrated out. Referring to (2.18), this distribution is given by

$$p(h|y) = \int p(h|\phi)\,p(\phi|y)\,d\phi = \begin{cases} 1-\sum_{j=1}^K \overline{\phi}_j & h=\infty \\ f(h|\overline{\phi}) & h\in[0,\infty) \end{cases}, \tag{3.1}$$

where

$$\overline{\phi} := E[\phi|y]. \tag{3.2}$$

Define

$$\overline{w} := 1 - \sum_{j=1}^{K} \overline{\phi}_j \qquad \text{and} \qquad \overline{\theta} := \frac{\overline{\phi}}{1-\overline{w}}. \tag{3.3}$$

Using (3.3), we can write

$$p(h|y) = \begin{cases} \overline{w} & h = \infty \\ (1-\overline{w})\, f(h|\overline{\theta}) & h \in [0,\infty) \end{cases}. \tag{3.4}$$

Note that $\overline{\phi}$ is computed from the posterior distribution for ϕ:

$$p(\phi|y) = \frac{p(y|\phi)\, p(\phi)}{p(y)}, \tag{3.5}$$

where

$$p(y) = \int p(y|\phi)\, p(\phi)\, d\phi. \tag{3.6}$$

For future reference let

$$L := p(y). \tag{3.7}$$

We can use L to compare models with different hyperparameter settings. For example, we can compare the base model to one with no point mass at infinity.

The likelihood $p(y|\phi)$ and the prior $p(\phi)$ are described next.

Likelihood. I assume the connection between the observations (i.e., the survey data) and the parameters is given by

$$y_i = F(h_i|\phi) + \varepsilon_i, \tag{3.8}$$

where $\varepsilon_i \overset{\text{iid}}{\sim} \mathsf{N}(0,\sigma^2)$. Note

$$F(h_i|\phi) = \sum_{j=1}^{K} \phi_j\, X_{ij}, \tag{3.9}$$

where

$$X_{ij} := F_{jK}(h_i) = I_{Q(h_i)}(j, K-j+1). \tag{3.10}$$

This setup delivers a linear regression:

$$y = X\phi + \varepsilon, \tag{3.11}$$

where X is an $n \times K$ design matrix. For $K > n$, X cannot have full column rank.

The likelihood including the nuisance parameter σ^2 is

$$p(y|\phi,\sigma^2) = \prod_{i=1}^{n} \mathsf{N}\big(y_i|F(h_i|\phi),\sigma^2\big), \tag{3.12}$$

where $\mathsf{N}(\,\cdot\,|\mu,\sigma^2)$ is the PDF of the normal distribution with mean μ and variance σ^2. We obtain the marginal likelihood for ϕ by integrating out σ^2, using $p(\sigma^2) \propto 1/\sigma^2$:

$$p(y|\phi) = \int p(y|\phi,\sigma^2)\, p(\sigma^2)\, d\sigma^2 \propto S(\phi)^{-n/2}, \tag{3.13}$$

where

$$S(\phi) := (y - X\phi)^\top (y - X\phi). \tag{3.14}$$

Prior. Recall $\phi = (1 - w)\, A^{K,k}\, \theta_k$. It is convenient to specify the prior for ϕ via the prior for k, θ_k, and w. Let $p(k, \theta_k, w) = p(\theta_k|k)\, p(k)\, p(w)$, where $p(w)$ and $p(k)$ will be specified later. For the time being, we note that we require $p(k) = 0$ for $k > K$.

Let the prior for θ_k be given by

$$p(\theta_k|k) = \mathsf{Dirichlet}\big(\theta_k|(\alpha/k)\, \iota_k\big), \tag{3.15}$$

where α (a fixed hyperparameter) is the concentration parameter and ι_k is a vector of k ones. The PDF of the Dirichlet distribution is given by

$$\mathsf{Dirichlet}(\theta_k|\lambda_k) = \frac{\Gamma(\lambda_{0k})}{\prod_{j=1}^k \Gamma(\lambda_{jk})} \prod_{j=1}^k \theta_{jk}^{\lambda_{jk}-1}, \tag{3.16}$$

where $\lambda_{jk} > 0$, $\lambda_{0k} := \sum_{j=1}^k \lambda_{jk}$, and $\Gamma(z) = \int_0^\infty t^{z-1}\, e^{-t}\, dt$. Note $E[\theta_{jk}|k] = \lambda_{jk}/\lambda_{0k}$. The prior variation around this expectation is inversely related to λ_{0k}, which is called the *concentration parameter*.

For the chosen prior, $\lambda_{jk} = \alpha/k$ and $\lambda_{0k} = \alpha$. Therefore the prior expectation of θ_{jk} is $1/k$ and consequently

$$E[F(x|\theta_k)|k] = \sum_{j=1}^k \frac{1}{k}\, F_{jk}(x) = Q(x). \tag{3.17}$$

In order to encourage sparsity, I set $\alpha = 1$.

Sampling scheme. Draws from the posterior are made via importance sampling. Let $\{\phi^{(r)}\}_{r=1}^R$ represent R draws of ϕ from its prior. These draws can be made by first drawing k and w from their priors, next drawing θ_k from its conditional prior (given the draw of k), and then setting

$$\phi^{(r)} = A^{K,k^{(r)}} \left((1 - w^{(r)})\, \theta^{(r)}_{k^{(r)}} \right). \tag{3.18}$$

Let

$$\zeta^{(r)} := S(\phi^{(r)})^{-n/2} \qquad \text{and} \qquad Z := \sum_{r=1}^R \zeta^{(r)}. \tag{3.19}$$

Then

$$\overline{\phi} \approx \widehat{\phi} := \frac{1}{Z} \sum_{r=1}^R \zeta^{(r)}\, \phi^{(r)} \qquad \text{and} \qquad L \approx \widehat{L} := Z/R. \tag{3.20}$$

Approximations to other quantities are $\overline{w} \approx \widehat{w} := 1 - \sum_{j=1}^K \widehat{\phi}_j$ and $\overline{\theta} \approx \widehat{\theta} := \widehat{\phi}/(1 - \widehat{w})$.

Computation reduction. We can reduce the amount of computation by not actually making draws of k and (more importantly) by delaying the elevation of $(1-w)\,\theta_k$. [When viewed from the perspective of Bayesian Model Averaging (as applied to a collection of models indexed by k), the organization of the computations described in this subsection is natural.]

Let $R_k \approx p(k)\,R$ denote the expected number of draws of k that would be made if k were drawn from its prior, where $\sum_{k=1}^K R_k = R$. For each k, make R_k draws of θ_k from its conditional prior along with R_k draws of w from its prior and set

$$\phi_k^{(r)} = (1 - w^{(r)})\,\theta_k^{(r)}. \tag{3.21}$$

The relevant draws now consist of $\{\phi_k^{(r)}\}_{r=1}^{R_k}$ for $k = 1, \ldots, K$.

Let

$$\zeta_k^{(r)} = S(A^{K,k}\,\phi_k^{(r)})^{-n/2}. \tag{3.22}$$

A significant reduction in computation comes from

$$S(A^{K,k}\,\phi_k^{(r)}) \equiv (y - X_k\,\phi_k^{(r)})^\top (y - X_k\,\phi_k^{(r)}), \tag{3.23}$$

where $X_k = X A^{K,k}$. Since X_k is computed once, $X_k\,\phi_k^{(r)}$ involves fewer operations than $X\,(A^{K,k}\,\phi_k^{(r)})$ as long as $k < K$.

Next define

$$Z_k := \sum_{r=1}^{R_k} \zeta_k^{(r)} \qquad \text{and} \qquad \widetilde{\phi}_k := \sum_{r=1}^{R_k} \zeta_k^{(r)}\,\phi_k^{(r)}. \tag{3.24}$$

Then $Z = \sum_{k=1}^K Z_k$ and

$$\widehat{\phi} = \frac{1}{Z} \sum_{k=1}^K A^{K,k}\,\widetilde{\phi}_k. \tag{3.25}$$

The total number of elevations is reduced from R to K.

We can give (3.25) a natural representation:

$$\widehat{\phi} = \sum_{k=1}^K \widehat{v}_k\,(A^{K,k}\,\widehat{\phi}_k), \tag{3.26}$$

where $\widehat{v}_k := Z_k/Z$ approximates the posterior probability of k and $\widehat{\phi}_k := \widetilde{\phi}_k/Z_k$ approximates the posterior conditional expectation $\overline{\phi}_k := E[\phi_k|z_{1:n}, k]$. Finally, define $\widehat{w}_k := 1 - \sum_{j=1}^k \widehat{\phi}_{jk}$ for future reference.

Adequacy of fit. The ability of the model to fit a prior depends on both the centering function Q and the maximum order of the polynomial K. The more closely the centering function is aligned to the data, the smaller is the required variation around it. In particular, if $F(h|\widehat{\theta})$ fits well, then using it as the centering function should obviate the need for $k > 1$. Thus an indication of the adequacy of fit can be obtained by setting $Q(h) = F(h|\widehat{\theta})$, estimating the model with $K' \gg 1$, and checking the posterior probabilities for $k' = 1, \ldots, K'$.

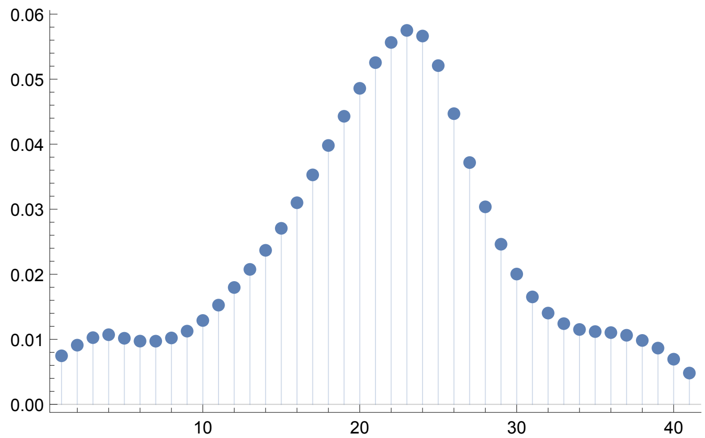

FIGURE 2. $\widehat{\phi}_{jK}$ for $j = 1, \ldots, K = 41$.

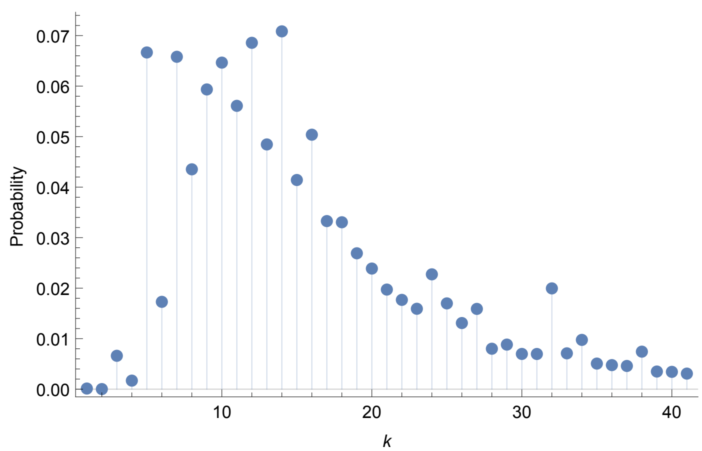

FIGURE 3. Posterior distribution for k.

4. Results

I chose $Q(x)$ by fitting a simple parametric distribution to the survey data: $Q(x) = 2^{-a^*/x}$ where

$$a^* = \underset{a}{\operatorname{argmin}} \sum_{i=1}^{n} \left(z_i - (1 - w^*)\, 2^{-a/h_i}\right)^2. \tag{4.1}$$

In particular, $a^* = 3.65$ given the chosen value of $w^* = 0.05$. Note

$$q(x) = \log(2)\, a^*\, 2^{-a^*/x}\, x^{-2}. \tag{4.2}$$

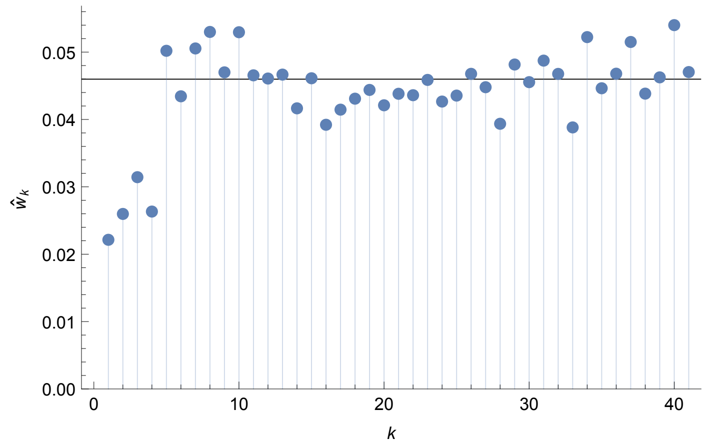

FIGURE 4. Posterior probabilities for the point mass, $\{\widehat{w}_k\}_{k=1}^{41}$ with $\widehat{w} = 0.046$ indicated.

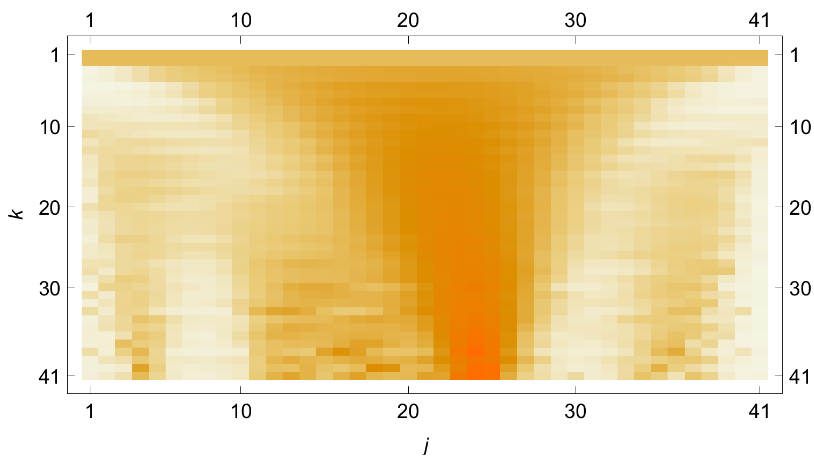

FIGURE 5. Row k shows $A^{K,k}\,\widehat{\phi}_k$ for $k = 1, \ldots, K = 41$.

I let $p(w) = \mathsf{Beta}(w|1, 19)$, which has a mean of 0.05. I chose $K = 41$ and let $p(k) = 1/K$ for $k = 1, \ldots, K$. I set $R = 41 \times 10^7$ for the number of draws from the prior so that $R_k = 10^7$.[6]

The central results are $\widehat{w} = 0.046$ and $\widehat{\phi}$ as shown in Figure 2. The posterior distribution for k is shown in Figure 3. Posterior probabilities $\widehat{w}_k$ for the point mass at infinity are shown in Figure 4 along with the model-averaged $\widehat{w} = 0.046$. The elevated vectors $A^{K,k}\,\widehat{\phi}_k$ for each k are shown row-by-row in Figure 5 and the corresponding weighted vectors $v_k\,A^{K,k}\,\widehat{\phi}_k$ are shown in Figure 6. See Figure 1 for a plot of $F(h|\widehat{\phi})$ and Figure 7 for a plot of $f(h|\widehat{\theta})$.

[6]The calculations were done on my MacBook Pro (circa 2014) using *Mathematica* (with pseudo-compiled code). The entire calculation, which involved generating close to 10^{10} gamma variates, took about 11 minutes using some parallelization.

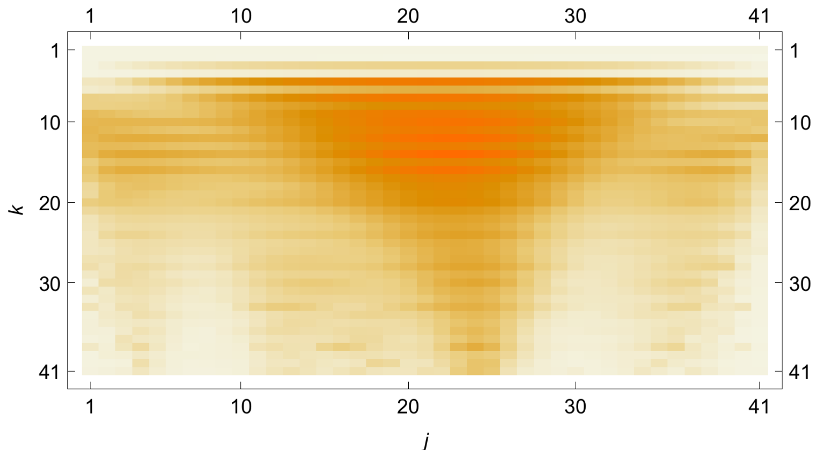

FIGURE 6. Row k shows $\widehat{v}_k\, A^{K,k}\, \widehat{\phi}_k$ for $k = 1, \ldots, K = 41$.

Adequacy of the fit. As a check on the adequacy of the fit, I redid the estimation using $F(h|\widehat{\theta})$ as the centering function, constructing the design matrix $\widehat{X}'$ via

$$\widehat{X}'_{ij} := I_{F(h_i|\widehat{\theta})}(j, K' - j + 1). \tag{4.3}$$

I chose $K' = 21$ and $R = 21 \times 10^6$. The posterior distribution for k is shown in Figure 8. The first two probabilities account for more than 99%. I found $F(h|\widehat{\phi}')$ to be indistinguishable from $F(h|\widehat{\phi})$. In summary, this check produced no evidence against the adequacy of the fit.

Evidence in favor of $w = 0$. I ran the model imposing $w = 0$. The centering function was refit under the assumption $w^* = 0$, producing $a^* = 3.96$ [see (4.1)]. The Bayes factor in favor of this restricted model relative to the unrestricted base model is $\widehat{L}'/\widehat{L} \approx 0.5$. In other words, there is very mild evidence in favor of $w > 0$.

5. First-order autoregressive coefficient

The first-order autoregressive model (for the log of the real exchange rate, m_t) can be expressed as

$$m_t = \gamma + \beta\, m_{t-1} + \varepsilon_t, \tag{5.1}$$

where β is the first-order autoregressive coefficient. According to (5.1), the half-life h is given by $\beta^h = 1/2$. This expression can be solved for

$$h(\beta) := \frac{-\log(2)}{\log(\beta)}. \tag{5.2}$$

Note

$$h'(\beta) = \frac{\log(2)}{\beta\, \log(\beta)^2}. \tag{5.3}$$

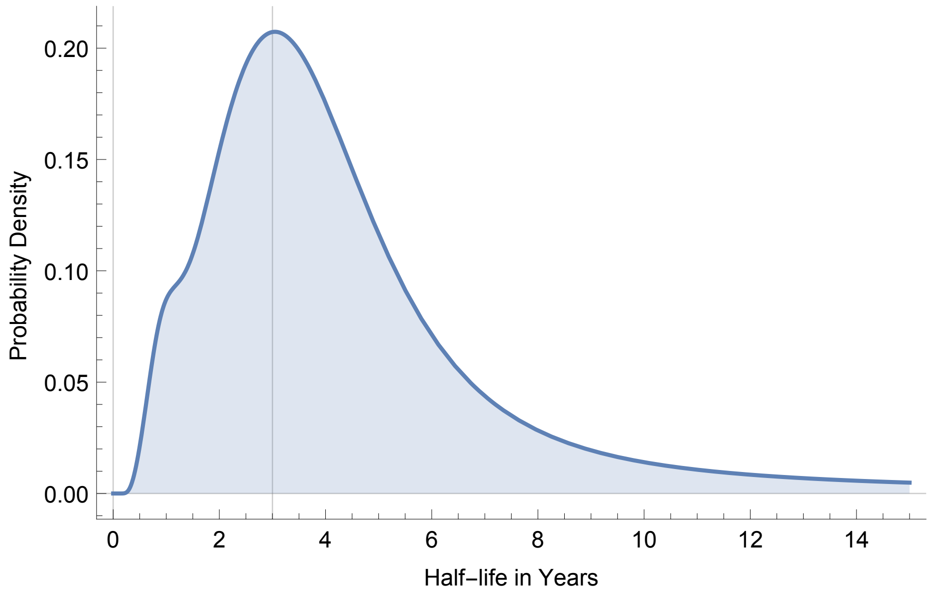

FIGURE 7. PDF for survey fit prior, $f(h|\widehat{\theta})$. The mode occurs at $h = 3.0$ years. The fit delivers $\Pr[h = \infty] = 0.046$.

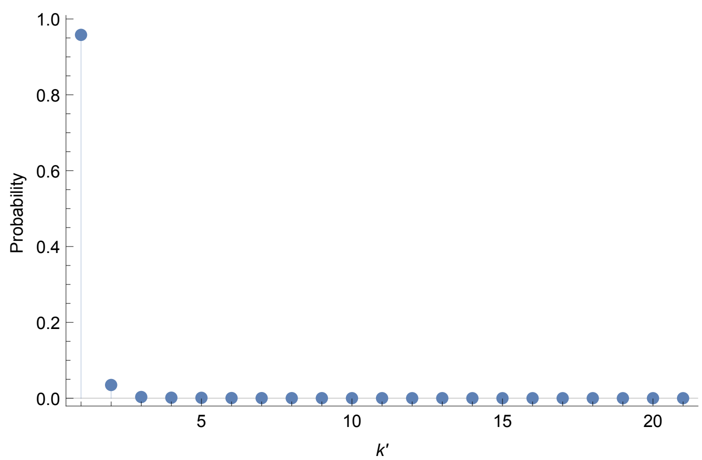

FIGURE 8. Posterior probabilities for $k = 1, \ldots, 21$, where $Q(h) = F(h|\widehat{\theta})$.

With these expressions, the model in (2.1) can be written in terms of β as follows:

$$p(\beta|\theta_k, w) = \begin{cases} w & \beta = 1 \\ (1-w)\, g(\beta|\theta_k) & \beta \in [0,1) \end{cases}, \tag{5.4}$$

where

$$g(\beta|\theta_k) := f(h(\beta)|\theta_k)\, h'(\beta). \tag{5.5}$$

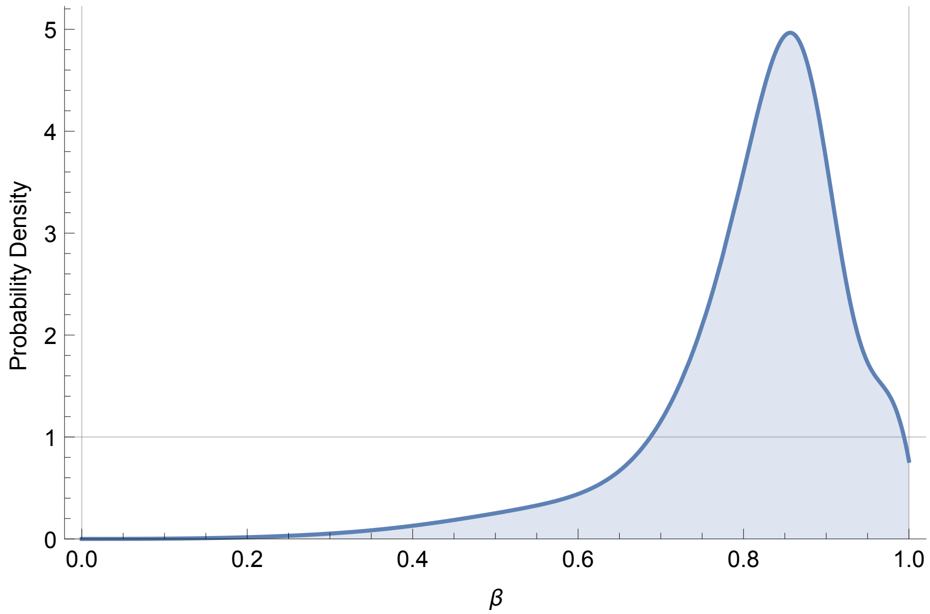

FIGURE 9. PDF for fit survey prior expressed in terms of β (with a uniform distribution for reference). This fit delivers $\Pr[\beta = 1] = 0.046$.

Consequently, the posterior probability of a unit root is approximated by $\widehat{w} = 0.046$ and the posterior density over the unit interval is given by $g(\beta|\widehat{\theta})$ as shown in Figure 9.

References

Dwyer, G. P. and M. Fisher (2014). Real exchange rates and unit roots: Learning about the distribution. Technical report, Federal Reserve Bank of Atlanta.

Fisher, M. (2015). Simplex regression. Technical report, Federal Reserve Bank of Atlanta.

Gosling, J. P., J. E. Oakley, and A. OHagan (2007). Nonparametric elicitation for heavy-tailed prior distributions. *Bayesian Analysis 2*(4), 693–718.

Kilian, L. and T. Zha (2002). Quantifying the uncertainty about the half-life of deviations from PPP. *Journal of Applied Econometrics 17*, 107–125.

Quintana, F. A., M. F. J. Steel, and J. T. A. S. Ferreira (2009). Flexible univariate continous distributions. *Bayesian Analysis 4*(4), 497–522.

FEDERAL RESERVE BANK OF ATLANTA, RESEARCH DEPARTMENT, 1000 PEACHTREE STREET N.E., ATLANTA, GA 30309–4470

E-mail address: mark.fisher@atl.frb.org

URL: http://www.markfisher.net