Journal of Atmospheric and Solar-Terrestrial Physics 70 (2008) 1797–1809

Contents lists available at ScienceDirect

Journal of Atmospheric and Solar-Terrestrial Physics

journal homepage: www.elsevier.com/locate/jastp

A state-space model of radiation belt electron flux dynamics

E.J. Rigler [a,*], D.N. Baker [b]

[a] High Altitude Observatory, National Center for Atmospheric Research, P.O. Box 3000, Boulder, CO 80307, USA
[b] Laboratory for Atmospheric and Space Physics, University of Colorado, Boulder, CO 80303, USA

ARTICLE INFO

Article history:
Accepted 29 January 2008
Available online 13 February 2008

Keywords:
Relativistic electrons
Radiation belt dynamics
State-space model
Data assimilation

ABSTRACT

A state-space model was optimized from data and used to characterize the linear dynamics governing variations observed in 2–6 MeV electron flux across a range of magnetic L shells. A correction term exploits correlated structure in previous one-step prediction errors, or innovations, to improve the current forecast. More importantly, this correction term helps reduce parameter estimation bias that arises when relevant inputs are ignored, or higher-order linear and nonlinear dynamical terms are left out of the model while it is being trained. Analyses of the L-dependent response functions, one-step predictions, and prediction error statistics, lead to several conclusions: (1) the direct effects of first-order solar wind perturbations only penetrate to $L{\sim}4R_E$, while linear feedback, which dominates flux dynamics throughout the radiation belts, accounts for over 80% of the observed variability below this location; (2) electron flux diffuses upward above $L{\sim}5R_E$, and downward below $L{\sim}5R_E$, except below $L{\sim}1.75R_E$, where the estimated model parameters are considered suboptimal anyway; (3) corrections to model output required above $L{\sim}4R_E$ suggest that modified or additional solar wind drivers may be required for a more complete physical description of solar wind-radiation belt coupling; while (4) corrections to model output required below $L{\sim}4R_E$ indicate episodic reconfigurations of the global electron radiation belt state, a type of variation that will never be captured with linear dynamics alone.

© 2008 Elsevier Ltd. All rights reserved.

1. Introduction

Earth's magnetosphere is a highly dynamic system that derives its energy predominately from the impinging solar wind (e.g., Arnoldy, 1971; Perreault and Akasofu, 1978; Akasofu, 1979; Vasyliunas et al., 1982). Strong correlations between measurements of various solar wind and magnetospheric parameters have led to the broad application of so-called linear prediction filters as proxies for the true, sometimes nonlinear, coupling between these two distinct space plasma regimes (Iyemori et al., 1979; Bargatze et al., 1985; Clauer, 1986; Nagai, 1988). For instance, because radiation belt electron flux has long been known to exhibit a strong linear relationship to solar wind speed (Paulikas and Blake, 1979; Baker et al., 1986), Baker et al. (1990) used finite impulse response (FIR) linear filters to predict the $\log_{10}$-flux of multi-MeV electrons at geostationary altitudes. A modified form of this model is used today to predict >2 MeV electron fluences in space weather forecasts at NOAA's Space Environment Center.

More recently, a series of electron radiation belt studies (Vassiliadis et al., 2002, 2005; Rigler et al., 2005, 2007) have applied FIR filters to a much broader range of the Earth's inner magnetosphere. Solar wind data from NASA's OMNI database (King and Papitashvili, 2005) were used to predict electron flux measured by the Solar, Anomalous, and Magnetospheric Particle Explorer (SAMPEX; Baker et al., 1993; Cook et al., 1993). These complimentary studies employed both single- and

* Corresponding author. Tel.: +1 303 497 2186; fax: +1 303 497 2180.
E-mail addresses: jrigler@ucar.edu (E.J. Rigler), daniel.baker@lasp.colorado.edu (D.N. Baker).

1364-6826/$ - see front matter © 2008 Elsevier Ltd. All rights reserved.
doi:10.1016/j.jastp.2008.01.019

multi-input FIR filters to ascertain which solar wind inputs affected flux variations most efficiently at different geomagnetic L shells.

One conclusion reached by these studies in common was that variations in solar wind electromagnetic properties, which are expected to enhance magnetospheric convection via increased dayside magnetic field reconnection, affect radiation belt electron fluxes strongly just outside the slot region ($L\sim3$–$4R_E$), and to a milder but broader extent just inside of geostationary orbital altitudes ($L\sim4.5$–$6.5R_E$). Multi-input filter results in particular show clearly that flux increases should be expected just outside the slot, while temporary decreases are expected just inside geostationary orbit, and are largely due to adiabatic losses related to ring current enhancements during magnetic storms (the Dst effect; e.g., Kim and Chan, 1997; Li et al., 1997).

Another common conclusion reached between these studies relates to solar wind parameters that are closely linked to mostly viscous interactions with the magnetosphere, especially the solar wind bulk speed. These tend to drive significant variations in electron flux beyond $L = 3\,R_E$. Results from the multi-input filters, which help separate the influence of simultaneous solar wind inputs, indicate that increases in the solar wind bulk speed are almost exclusively associated with increases in electron flux. If those multi-input filter components driven by electromagnetic solar wind inputs do indeed capture the adiabatic changes in electron flux, as suggested previously, a case might be made that the enhancements driven by increases in bulk speed are the result of a genuine increase in electron phase space density (PSD). However, it must be noted up front that these previous filter-based studies, as well as the present work, are all designed to model electron flux, so all physical interpretations must continue to be made in this context.

One apparent discrepancy arose from these studies when solar wind plasma density was considered: single-input FIR filter prediction error statistics suggested that density does have a significant impact on electron flux variations, while multi-input filters indicated that the role of plasma density is relatively insignificant, at least when daily averaged observations are used. To understand this discrepancy, we first reconsider results from Rigler et al. (2004), where it was shown that there exists a lack of time stationarity in adaptive single-input FIR filter coefficients used to predict SAMPEX electron fluxes. Similarly, Rigler et al. (2007) demonstrated a lack of stationarity in prediction error statistics from static single- and multi-input linear filter output. If instrument error can be discounted, such non-stationary behavior in a geophysical time series is usually indicative of some sort of "missing inputs" required by the model.

These "missing inputs" may indeed be additional drivers, like those considered by Rigler et al. (2007). They might also be time-lagged model output, which amounts to linear dynamic feedback. Finally, since it is always possible to approximate nonlinear system dynamics with a properly chosen polynomial expansion, neglecting the higher-order terms in this polynomial should also be considered a form of "missing input" (Pindyck and Rubinfeld, 1991). In all these cases, unless the "missing input" is completely uncorrelated with any of the known system variables, the end result will be biased model output, which generates biased prediction errors, and eventually leads to biased estimates of those empirical terms being optimized for the model.

An important objective of the present study is to minimize bias error associated with missing inputs so that physical insights might be more easily gained from the model parameters optimized on data. The first example above, or the "missing driver" problem, was already addressed by Rigler et al. (2007), so we just use the same solar wind drivers they did. Linear dynamic feedback has been neglected in most previous studies because the filters were designed with a sufficient number of time lags to statistically capture the system's entire dynamic response. However, if the filter is truncated before the full dynamic response of a system is able to decay to zero, it will be as if relevant inputs were neglected during the optimization process, leading again to biased parameters.

Clearly, this is a major problem in the radiation belts for L shells lower than $\sim4R_E$, where electron flux lifetimes can be on the order of weeks, months, or even years. To properly model such a response with an FIR filter requires dozens to hundreds of free parameters, and ultimately, similarly lengthy sequences of valid input data with which to convolve it. A solution to this problem is to just design a model that employs real dynamic feedback, and is therefore capable of reproducing a potentially infinite impulse response (e.g., Shynk, 1989). As we shall soon see, a linear state-space model makes this very easy. It also makes the inclusion of multiple inputs a fairly trivial task. Perhaps most important from the perspective of the current study, state-space models can accommodate dynamic feedback from multiple outputs.

The final source of bias error described above, nonlinear dynamics, is perhaps the most difficult type of "missing input" to add to any empirical model. As we noted above, a nonlinear system can be approximated, to an arbitrary precision, by using a properly chosen polynomial expansion of the system variables. Of course the most obvious problem with this approach is that all the relevant system variables may not even be known, or available as the observations required for parameter optimization. What's more, the degree of the polynomial required to approximate a nonlinear system may be so large as to discourage empirical optimization of the coefficients.

Fortunately, these problems can be largely mitigated by assuming that observation errors are negligible, or can at least be characterized as zero-mean, random sequences, and adding a non-deterministic correction term to the model that pushes model output toward observations. This approach is not particularly suited for long-term forecasts because future observations will not yet be available for calculating the necessary prediction errors. However, from the perspective of developing and refining our understanding of the dynamics of the system, it is bias error, more than any loss of predictive ability, that needs to be reduced in our empirical model. Besides, confidence in the terms that were explicitly included in the model makes it easier to ascertain which "missing inputs" should be

considered for inclusion in a future version of the model, thus eventually leading to better long-term forecasts.

One should not choose the correction term arbitrarily. If the corrector is unable to fully account for the predictable "noise" in the system, and correlated residuals persist, so will biased parameter estimates; only now these estimates also include effects from the correction term, which exacerbates the bias problem (Rigler et al., 2007). It is necessary for both the correction term and deterministic model coefficients to be estimated simultaneously, which inevitably requires a nonlinear, and computationally expensive, optimization algorithm. A recursive nonlinear estimator helps to redistribute this computational expense over time, and as we shall soon see, state-space models are well-suited to such algorithms.

Section 2 of this paper describes both the general linear state-space model, with an error (innovation) correction term, as well as the specific form used for this study. Section 3 presents and discusses the model output, which consists of 1-day predictions of the $\log_{10}$-flux of radiation belt electrons. Section 4, using techniques described in more detail in Appendix A, examines these results in a more statistically rigorous manner in order to determine exactly what fraction of the observed electron flux variability can be attributed to each component of the full state-space prediction. The paper ends with a summary of the preceding discussions, and a few concluding remarks.

2. Linear state-space model

Eq. (1) represents a generic linear state-space model designed to predict an arbitrary state vector, $\mathbf{x}$, one discrete time step into the future

$$\begin{aligned}\mathbf{x}(t+1) &= \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t) + \boldsymbol{\nu}(t),\\ \mathbf{y}(t) &= \mathbf{C}\mathbf{x}(t) + \mathbf{D}\mathbf{u}(t) + \boldsymbol{\varepsilon}(t).\end{aligned} \tag{1}$$

The state-space matrices are defined as $\mathbf{A} \in \mathbb{R}^{n\times n}$, $\mathbf{B} \in \mathbb{R}^{n\times p}$, $\mathbf{C} \in \mathbb{R}^{l\times n}$, and $\mathbf{D} \in \mathbb{R}^{l\times p}$, where n is the dynamic order of the system, p is the number of inputs, and l is the number of expected outputs or predictions. Simply put, the state is propagated forward in time as a linear function of its current self, $\mathbf{A}\mathbf{x}(t)$, plus a linear function of some external forcing or input, $\mathbf{B}\mathbf{u}(t)$. For generality, model uncertainty is represented by the process error term, $\boldsymbol{\nu}(t)$. Since the modeled state may not be directly comparable to observations of the real system output, $\mathbf{y}(t)$, another linear mapping function, $\mathbf{C}\mathbf{x}(t)$, can be used to interpolate, or otherwise modify the state, in order to best match output measurements. Again, for generality, the state-to-output mapping model considers measurement errors, $\boldsymbol{\varepsilon}(t)$, as well as potential statistical relationships that might exist between the input and output, but which cannot be captured by the system dynamics. This so-called "feed-through" term, $\mathbf{D}\mathbf{u}(t)$, is only required for discrete time models because certain dynamics will always be missed due to the inherent time lag.

Often, little is known a priori about either of the error terms, $\boldsymbol{\nu}(t)$ or $\boldsymbol{\varepsilon}(t)$. Therefore, it is not unreasonable to assume that one is related to the other in a predictable manner. In particular, within the context of model training and/or validation, it is assumed that a complete set of output observations is always available with which to compare model predictions. As a result, $\boldsymbol{\varepsilon}(t)$ would be known at each time step, so this information could be used to help specify $\boldsymbol{\nu}(t)$, and potentially improve predictions. In keeping with the linear state-space model concepts described above, a simple linear transformation is applied to the measurement error, and the process error becomes $\mathbf{K}\boldsymbol{\varepsilon}(t)$, which can be plugged back into Eq. (1) to improve predictions. A block diagram of this so-called innovations form of state-space model (ISS; Ljung, 1979, 1999) is shown in Fig. 1.

No unique state-space representation of a linear system exists (Priestley, 1988). This is because an equivalent state-space representation can always be constructed from the similarity transform $\mathbf{x}' = \mathbf{T}\mathbf{x}$ ($\mathbf{T}$ being a non-singular matrix of the same dimension as $\mathbf{A}$). If $\mathbf{T}$ exists, it can be used to modify the remaining state-space matrices according to: $\mathbf{A}' = \mathbf{T}\mathbf{A}\mathbf{T}^{-1}$, $\mathbf{B}' = \mathbf{T}\mathbf{B}$, $\mathbf{K}' = \mathbf{T}\mathbf{K}$, and $\mathbf{C}' = \mathbf{C}\mathbf{T}^{-1}$ (the feed-through term, $\mathbf{D}$, does not change). This being noted, there are certain forms that the matrices can take on that possess desirable numerical characteristics, or which allow for easier physical interpretation. For instance, consider a system in which only the diagonal of the $\mathbf{A}$ matrix possesses non-zero values. In such a system, each element of $\mathbf{x}$ might be considered a distinct physical state whose first-order linear dynamic is defined by a scaled version of its own value at the previous discrete time-step. If these values are all less than one, the model is considered stable, and represents multi-output exponential decay.

A slight modification to this state-space form might incorporate super and subdiagonals into the $\mathbf{A}$ matrix, implying that each state is related to the previous value of its two nearest-neighbor states, in addition to itself. Such dynamics represent discrete linear diffusion. Finally, if the mapping matrix is nothing more than a square identity matrix of dimension n, there will be a one-to-one mapping of the state to outputs, simplifying further our interpretation of the model dynamics.

For the purposes of this study, we considered a set of 27 distinct physical states that correspond to the $\log_{10}$-

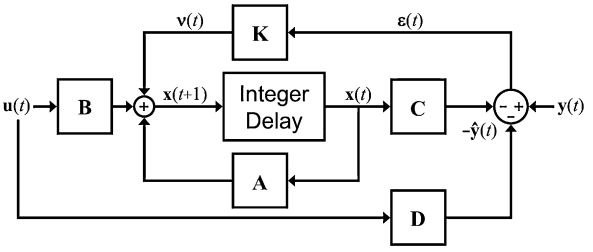

Fig. 1. This block diagram represents a generic state-space model, in a training or validation configuration, that uses innovations to adjust one-step predictions. A discrete, fixed length time delay allows both deterministic ($\mathbf{A}\mathbf{x}$) and stochastic ($\mathbf{K}\boldsymbol{\varepsilon}$) feedback loops. Filtered input can perturb the state ($\mathbf{B}\mathbf{u}$), or bypass the model's dynamics altogether ($\mathbf{D}\mathbf{u}$) to combine directly with the system output, thus exploiting correlations between the input and output not captured by the system dynamics due to the model's discretization.

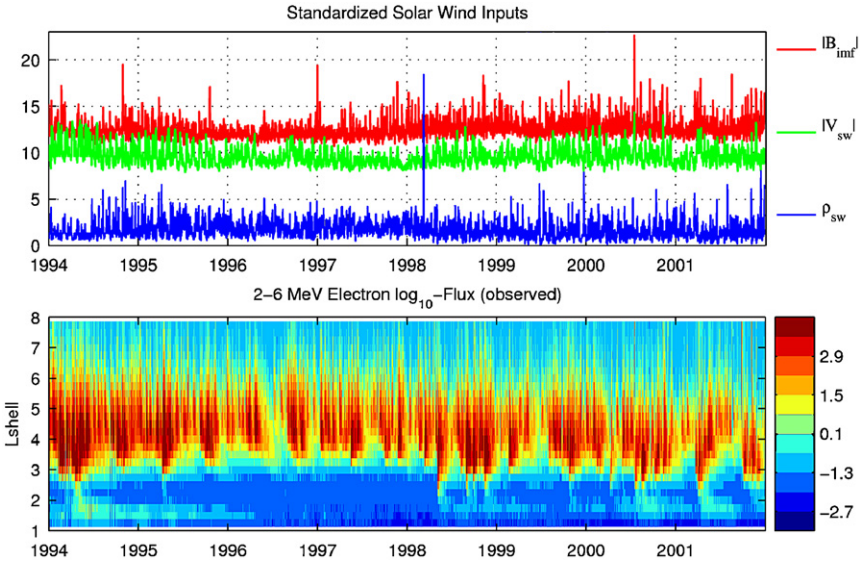

Fig. 2. Eight years of solar wind inputs (IMF magnitude, plasma bulk speed, and mass density) are shown in the top panel. All are normalized by their own standard deviations, then offset from zero by {10,5,0} respectively. The $\log_{10}$ of 2–6 MeV electron fluxes are shown in the bottom panel, as a function of magnetic L-shell, in units of Earth radii. The data shown here were used for both model estimation and validation.

flux of energetic electrons (2–6 MeV) measured by the SAMPEX satellite. Each state element represents daily averaged fluxes that have been resampled using quarter R_E magnetic L shell bins centered on the monotonically increasing span 1.25–7.75R_E. We also chose three daily averaged solar wind inputs as external perturbations to our system: the magnitude of the interplanetary magnetic field ($|\mathbf{B}_{imf}|$); the bulk speed of the solar wind ($|\mathbf{V}_{sw}|$); and the mass density of the solar wind plasma (ρ_{sw}). These data were either obtained directly, or derived, from the NASA's OMNI2 database (King and Papitashvili, 2005), and so have already been propagated to the Earth's position from their respective observation platforms. They were chosen to provide reasonably meaningful physical inputs for our system, but they are probably not optimal in any physical sense; mostly they were chosen so as to be comparable with Rigler et al. (2007). These training data are shown in Fig. 2.

Eq. (2) presents the input, state, output, and error vectors required at each time step. Eq. (3) presents the ISS model matrices in a manner that should make the adjustable and fixed coefficients apparent, as well as specify their respective dimensions:

$$\mathbf{u}(t)=\begin{bmatrix}1\\ u_1(t)\\ u_2(t)\\ u_3(t)\end{bmatrix},\quad \mathbf{x}(t)=\begin{bmatrix}x_1(t)\\ x_2(t)\\ \vdots\\ x_{27}(t)\end{bmatrix}$$

$$\mathbf{y}(t)=\begin{bmatrix}y_1(t)\\ y_2(t)\\ \vdots\\ y_{27}(t)\end{bmatrix},\quad \boldsymbol{\varepsilon}(t)=\begin{bmatrix}\varepsilon_1(t)\\ \varepsilon_2(t)\\ \vdots\\ \varepsilon_{27}(t)\end{bmatrix} \tag{2}$$

$$\mathbf{A}\equiv\begin{bmatrix}a_{1,1} & a_{1,2} & 0 & \cdots & 0\\ a_{2,1} & a_{2,2} & a_{2,3} & \ddots & \vdots\\ 0 & \ddots & \ddots & \ddots & 0\\ \vdots & \ddots & a_{26,25} & a_{26,26} & a_{26,27}\\ 0 & \cdots & 0 & a_{27,26} & a_{27,27}\end{bmatrix},$$

$$\mathbf{B}\equiv\begin{bmatrix}0 & b_{1,2} & b_{1,3} & b_{1,4}\\ 0 & b_{2,2} & b_{2,3} & b_{2,4}\\ \vdots & \vdots & \vdots & \vdots\\ 0 & b_{27,2} & b_{27,3} & b_{27,4}\end{bmatrix},$$

$$\mathbf{C}\equiv\begin{bmatrix}1 & 0 & \cdots & 0\\ 0 & 1 & \ddots & \vdots\\ \vdots & \ddots & \ddots & 0\\ 0 & \cdots & 0 & 1\end{bmatrix},$$

$$\mathbf{K}\equiv\begin{bmatrix}k_{1,1} & k_{1,2} & 0 & \cdots & 0\\ k_{2,1} & k_{2,2} & k_{2,3} & \ddots & \vdots\\ 0 & \ddots & \ddots & \ddots & 0\\ \vdots & \ddots & k_{26,25} & k_{26,26} & k_{26,27}\\ 0 & \cdots & 0 & k_{27,26} & k_{27,27}\end{bmatrix},$$

$$\mathbf{D}\equiv\begin{bmatrix}d_{1,1} & d_{1,2} & d_{1,3} & d_{1,4}\\ d_{2,1} & d_{2,2} & d_{2,3} & d_{2,4}\\ \vdots & \vdots & \vdots & \vdots\\ d_{27,1} & d_{27,2} & d_{27,3} & d_{27,4}\end{bmatrix}. \tag{3}$$

If these matrices are substituted back into Eq. (1), it is easy to see from the $\mathbf{Ax}(t)$ term how each state is a function of its previous value, as well as its two nearest neighbors (the upper and lower boundary conditions are assumed to be zero). The reader will also notice that the structure of $\mathbf{K}$ is very similar to that chosen for $\mathbf{A}$. This choice was not required, but it was not entirely arbitrary either. By imposing certain structural constraints on estimated coefficients, ISS model matrices can be converted into almost any well-known black-box linear time-series model. Coefficients in these particular matrices can be rearranged into a so-called multi-input/multi-output ARMAX model (Ljung, 1999), which may be more familiar to some readers than the linear state-space model. Finally, the $\mathbf{B}$ and $\mathbf{D}$ matrices are nearly identical in structure, despite their very distinct roles in the state-space model. However, $\mathbf{D}$ possesses an extra column of adjustable parameters, $d_{k=1\rightarrow l,1}$, that corresponds to the fixed singular value in the input vector $\mathbf{u}(t)$. This is analogous to how an intercept term is determined for simple linear regression using ordinary least squares (OLS). Since our system is assumed to be stable, the model's unperturbed response will always decay back to this static L shell-dependent baseline rather than zero.

If all that was required were optimal matrix coefficients for a standard state-space model (i.e., without corrections), a straight-forward least-squares estimation algorithm could be used. However, the ISS model described above introduces non-deterministic feedback via $\mathbf{K}\boldsymbol{\varepsilon}(t)$, and cannot be optimized using typical linear regression techniques since the regressors cannot all be known beforehand. It might be argued that improved predictions resulting from inclusion of the $\mathbf{K}\boldsymbol{\varepsilon}(t)$ term could be obtained just as readily from data assimilation. In this case, the standard state-space model would be optimized using more typical methods, then simply inserted into a Kalman filter, or some other well-known data assimilation algorithm, when making predictions.

However, if the $\mathbf{K}\boldsymbol{\varepsilon}(t)$ term is retained during parameter optimization, it helps account for missing inputs, non-linear feedback, and many other model deficiencies that can otherwise lead to biased estimates of the optimal model matrix coefficients. With this in mind, we employed a recursive prediction error estimator described by Ljung and Söderström (1983). This algorithm is very similar to the well-known extended Kalman filter (EKF) with a state vector augmented to include unknown model coefficients in addition to the dynamical state. The prediction error is minimized recursively in a manner that might one day be adapted to simulate non-stationary system behavior that manifests as time-varying model coefficients. For this study however, we assumed that the system dynamics were time-stationary, in which case the method can be proved to converge to a minimum in the error surface as $t\rightarrow\infty$.

Recent empirical studies of radiation belt dynamics using linear filters (Vassiliadis et al., 2002, 2005; Rigler et al., 2004, 2007) plotted the optimized FIR coefficients versus time lag and L shell, providing an L-dependent profile of the radiation belt's impulse response to various solar wind inputs. Fig. 3 is analogous to such plots, but

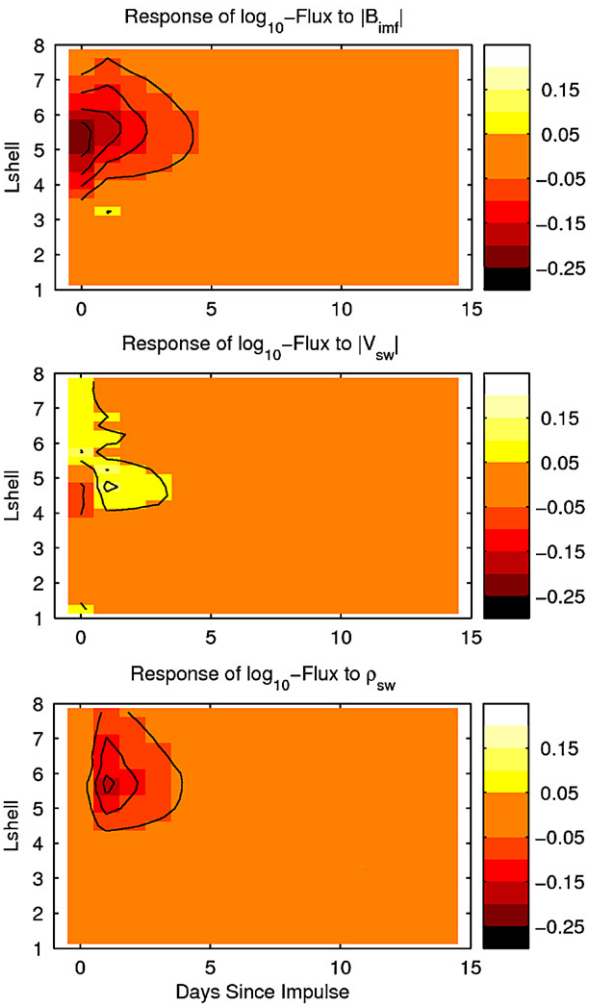

Fig. 3. A one-day unit impulse in each of the three solar wind inputs was processed by the state-space model to produce the profiles shown here. The model's output is discrete in time and space, with boxes centered on each day from 0 to 14, and each quarter L-shell from 1.25 to 7.75R_E. Discrete intensity bins are plotted for clarity, but the model's output is actually continuous in this regard.

because the profiles shown here are generated by a true dynamic model, and are not just the results of a large multi-variable linear regression, it should be understood that the state-space model will continue to generate a non-zero response beyond the largest time lag shown.

We actually start with the middle panel in Fig. 3, or the system's response to $|\mathbf{V}_{sw}|$, because it differs considerably from its FIR counterpart presented in Rigler et al. (2007). This difference arises mostly because enhancements in observed electron flux tend to occur 1–2 days after peaks in solar wind speed, and so they do not get incorporated into the estimation concurrently with the peak in $|\mathbf{V}_{sw}|$. This time lag is easily incorporated into an FIR filter by inserting additional time-lagged regression coefficients, but a state-space model that is first order in time like ours may very well miss many of these perturbations, leaving it up to the correction term to try to reproduce the flux variations instead. In contrast, the top and bottom panels

are similar to their multi-input counterparts for the same solar wind inputs, at least above $L{\sim}4R_E$. This is because $|\mathbf{B}_{imf}|$ and ρ_{sw} tend to exhibit marked increases less than a day before observed enhancements in electron flux, so their perturbation of the state is better captured by our first-order model. All of the input response profiles shown in Fig. 3 fail to register much of a change below $L{\sim}4R_E$, a significant deviation from multi-input FIR filter results that will be explained shortly.

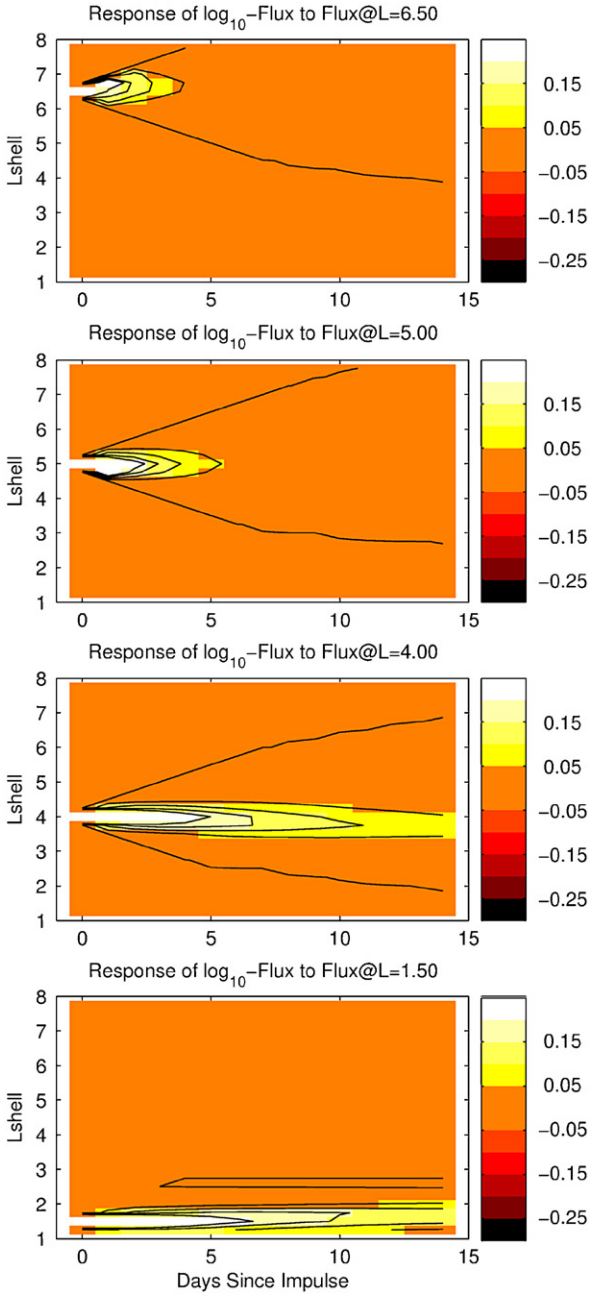

Fig. 4. A one-day unit impulse in each of its 27 states was processed by the state-space model. Representative response profiles are shown here. The model's output is discrete in time and space, with boxes centered on each day from 0 to 14, and each quarter L-shell from 1.25 to $7.75R_E$. Discrete intensity bins are plotted for clarity, but the model's output is actually continuous in this regard.

Input impulse response profiles are not the only useful way to represent the dynamic capabilities of a state-space model without examining closely every single coefficient. For instance, if one were to remove the inputs from the state-space equations temporarily, and artificially change each state from 0 to 1 for a single time step, it becomes possible to see how flux decays and diffuses across L shells over time when the system is free from external perturbations. All 27 model states were tweaked in this manner, and their respective response profiles were recorded. Decay times at lower L shells are extremely long, showing little change in state from the initial condition after 14 days. Decay times gradually become shorter as L increases, but there is a very noticeable drop between $L = 4$–$5R_E$. At geostationary altitudes, initial conditions decay to zero within 3–4 days. As for cross-L shell coupling, there is a slight tendency for electron flux to diffuse upward below $L = 1.75$. Otherwise, there is generally downward diffusion between $L = 1.75$ and 4.75, peaking near $L = 4.0$. There is no discernible preference for flux to diffuse up or down at $L = 5.0$, and there is generally upward diffusion of flux at all locations above $L = 5.0$. Fig. 4 presents four representative response profiles demonstrating what was just stated in words, and should be examined only while keeping in mind the caveats regarding flux and PSD mentioned previously.

3. Model component predictions

We begin this section by noting that a sequence of full ISS predictions is not easily distinguishable from the SAMPEX observations plotted previously if the intercept terms are retained. Therefore, we will not inundate the reader with such redundant information here. More important for physical interpretations, we will present and discuss the various components of our ISS model predictions. These include each of the solar wind variables' contributions to changes in electron flux modeled by the $\mathbf{Bu}(t)$ term, the potentially coupled linear feedback described by $\mathbf{Ax}(t)$, and the coupled error, or innovations, corrections described by $\mathbf{K\varepsilon}(t)$. For reference, the first panel in Fig. 5 is the full state-space model prediction with the intercept terms removed, effectively disregarding the background radiation belt state. This panel is somewhat redundant because it is the sum of the different component predictions plotted below it, but it still offers an interesting global view of relative flux variations with respect to the background state. Similarly, the column of small plots on the right are just averages of the component predictions to their left.

The first three rows of L-dependent component predictions offer a compelling visual representation of how different solar wind inputs contribute to different dynamic responses in the electron radiation belts consistently over a long time span. $|\mathbf{B}_{imf}|$, considered here as a proxy for electromagnetic coupling between the solar wind and terrestrial magnetosphere, clearly drives reductions in electron flux in the outer belt, in addition to a

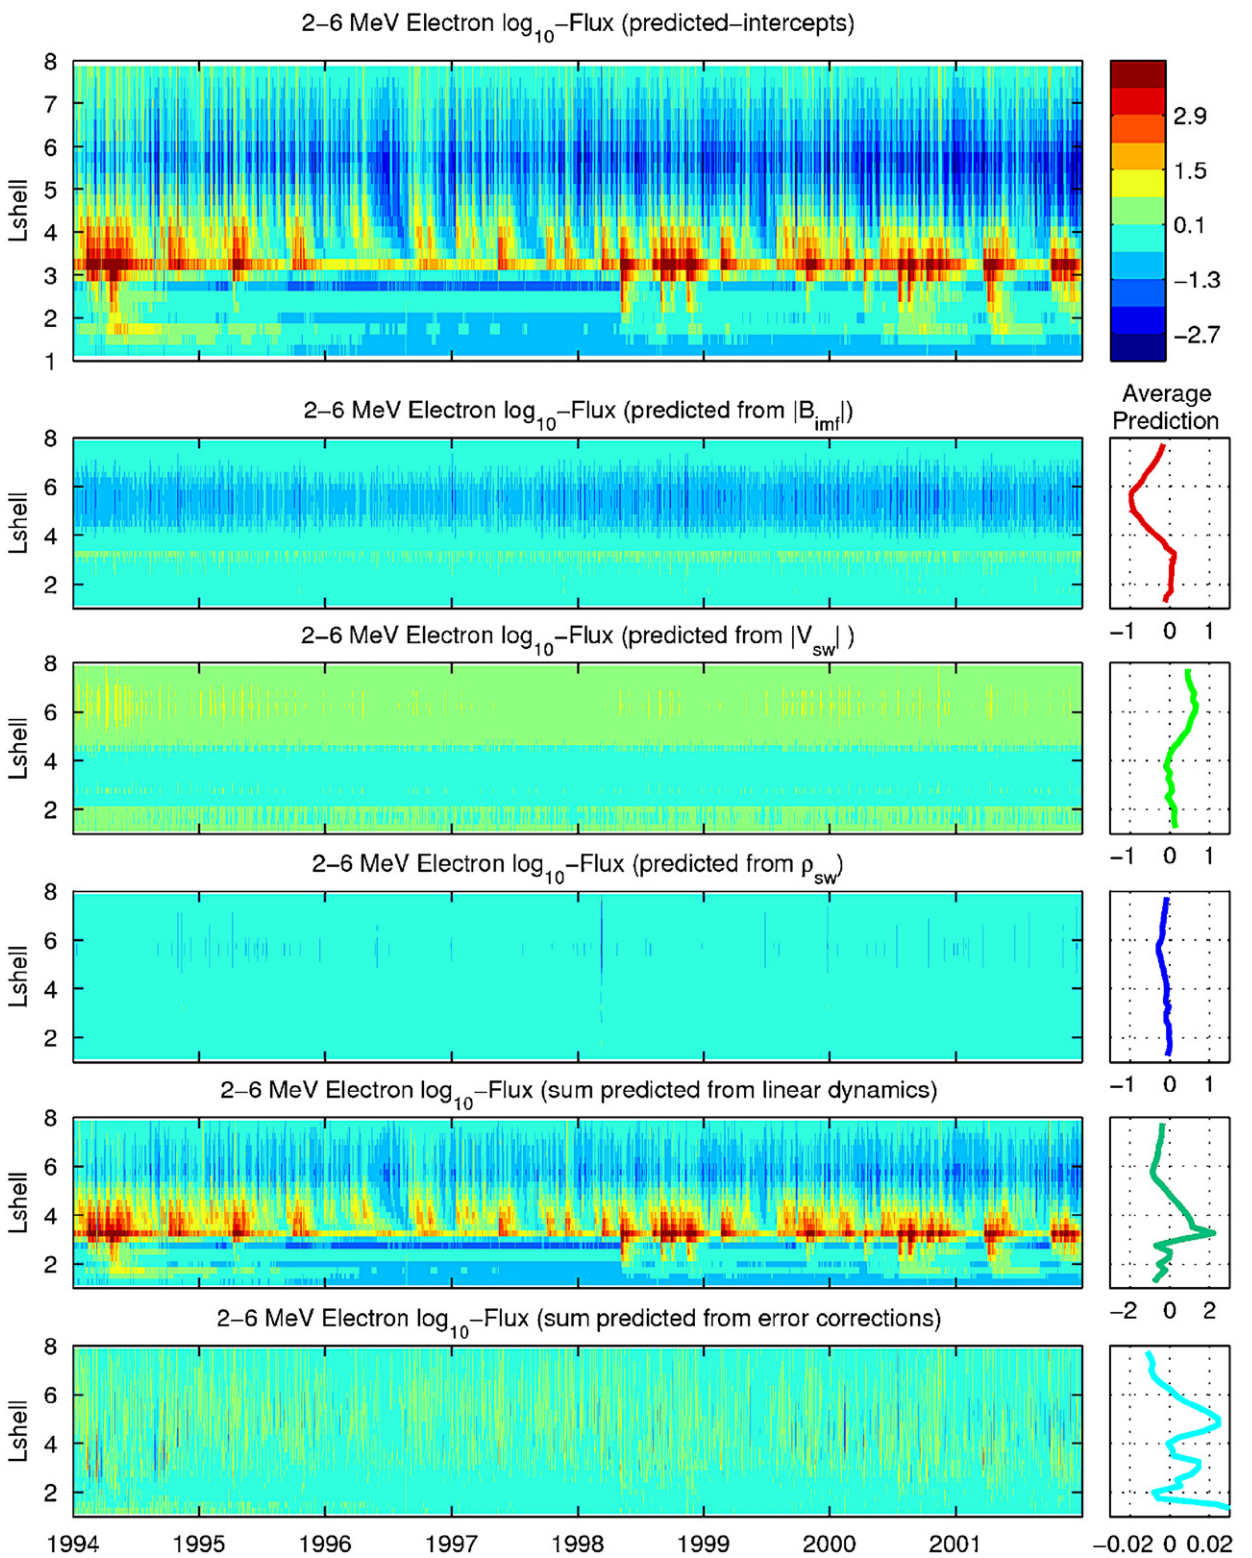

Fig. 5. The top panel is a full ISS model prediction with its static intercepts removed. It is actually the sum of the remaining panels, which describe contributions from perturbations caused by each of the three solar wind inputs, coupled linear feedback, and coupled innovations corrections, respectively.

slight enhancement just outside of $L = 3$. $|\mathbf{V}_{sw}|$, representing viscous interactions between these two plasma regimes, is clearly associated with flux enhancements in the outer belt, although perhaps not to the extent that we might have anticipated had the model's input response captured the 1–2 days lagged peak of the solar wind perturbation. Mass density is a proxy for plasma pressure effects, and appears to have only a minor influence on radiation belt electron fluxes measured by SAMPEX. Physically, we might suspect that the latter may not hold true when higher time cadences are considered, but based on these results, as well as results presented in Rigler et al. (2007), we conclude that extreme increases in ρ_{sw} are required to significantly affect fluxes over periods longer than a day.

Predictions plotted in the fourth row of component predictions are by far the dominant component of the full ISS prediction, so it is perhaps prudent at this juncture to remind the reader that these predictions are relative to the background radiation belt state for each L shell. With this in mind, it appears as though a large fraction of the time is spent recovering from relative flux dropouts at L shells higher than $5R_E$, while flux at L shells between 3 and $5R_E$ is typically decaying back to the baseline following strong and rapid enhancements. Much of this signal comes from exponential decay at any given L shell, especially inside $5R_E$, but a significant portion can be attributed to diffusion of electron flux across L shells due to the off-diagonal terms in the $\mathbf{A}$ matrix. The data shown here provide some visual evidence for this phenomenon, but more robust statistics presented in the following section will help make this argument more definitively.

Finally, the last row of component predictions shows flux variations captured by the $\mathbf{K}\boldsymbol{\varepsilon}(t)$ term in the ISS model, accounting for many of the dynamic processes that were not considered explicitly in our state-space model. First, note that the absolute value of the average $\mathbf{K}\boldsymbol{\varepsilon}(t)$ component prediction is very nearly zero, in spite of the fact that there are very obvious, sometimes quite large, corrections in the time-series plot to its left. This is because corrections, whether large or small, tend to oscillate about zero, thus canceling each other out in the end. There is a negative trend in the amplitude of these oscillations as L increases, similar to the negative trend in flux contributed by linear feedback. This is evident early on in the time series shown in the figure, and much more obvious upon closer inspection of the raw data. This is more consistent with the idea of missing internal dynamics in the model than missing external drivers. Since the corrections are usually quite large, especially at middle to lower L shells, these represent significant reconfigurations of the radiation belt state, and are most likely due to some highly nonlinear dynamics that our state-space model structure is incapable of reproducing in any deterministic way.

4. Prediction statistics

We now take the model predictions obtained in the previous section and determine robust prediction metrics designed to ascertain what fraction of the variability observed in the actual measurements can be attributed to changes in the different component predictions. Previous studies have used various correlation coefficients, prediction efficiencies, and other measures of "skill", all typically designed to provide a metric that falls between -1 and 1, or between 0 and 1. This has much value when comparing the predictive ability of one model with another model, or even a future incarnation of itself, but it does not really describe the model's ability to capture variations in the quantity it is attempting to emulate. The metric we chose might best be described as a fractional covariance, and does not necessarily fall within these limits

$$\sigma'_{y\hat{y}} = \frac{\sigma_{y\hat{y}}}{\sigma_{yy}}. \tag{4}$$

Eq. (4) represents the ratio of the variance shared between observations and predictions (i.e., "covariance"), and the simple variance of the observations (covariance of the observations with themselves). The advantage of this metric over the others mentioned above becomes clear if one considers a simple model that predicts $\sin(x)$, but which is mistakenly used to simulate a system that resembles $2\sin(x)$. An observation–prediction correlation coefficient would give a value of 1, which is clearly not appropriate since the predictions did not capture all of the true system's variability. So-called prediction efficiency (i.e., $1 - \sigma_{\varepsilon\varepsilon}/\sigma_{yy}$) provides a metric equal to 0.75, which is incorrect because there remain correlations between the residuals and the predictions that have not been considered (see below, and Appendix A). The fractional covariance is equal to 0.5 for this observation–prediction comparison. This equals the ratio of integrals of the functions that describe our predictions and observations (also equal to the ratio of their respective standard deviations), scaled by the simple correlation between the two, which is in this case equal to 1. Furthermore, a fractional covariance larger than 1, or even less than zero, are both perfectly valid; it simply means that the model tends to either overshoot, and/or is anticorrelated with observations (e.g., a model that predicts $-2\sin(x)$, used to simulate observations that go as $\sin(x)$, will produce a fractional covariance equal to -2, or the model captures -200% of the observed variability).

Perhaps most relevant to our current purposes, fractional covariance is a perfectly additive property when there are additive component predictions to consider:

$$\sigma'_{y\hat{y}_{cum}} = \sum_{k=1}^{K} \sigma'_{y\hat{y}_k}. \tag{5}$$

This relationship holds regardless of any correlations that might exist between the component predictions, which means that we can easily compare the relative contributions of each component prediction described in the previous section to the variations in the observed output. Finally, Eq. (5) is only one part of a convenient relationship that describes not only the variance shared between component predictions and observed output, but the covariance with prediction errors, and the variance of

the errors themselves

$$1 = \sum_{k=1}^{K} \sigma'_{y\hat{y}_k} + \sum_{k=1}^{K} \sigma'_{\varepsilon\hat{y}_k} + \sigma'_{\varepsilon\varepsilon}. \quad (6)$$

The last term in Eq. (6) is the fractional variance of the prediction errors, noted previously when discussing prediction efficiency. The middle term can often be assumed to equal zero because a least-squares-optimized model will decorrelate predictions and residuals. It must be kept in mind that this only holds true for in-sample validation. If the model is validated using a set of data other than the one on which it was trained, there is no guarantee that the predictions and residuals will be uncorrelated. More relevant to this study is the fact that our recursive estimation algorithm may not actually converge to an optimal solution given the finite set of training data available. However, as long as this middle term is reasonably small it may be assumed that the model is optimized enough for real-world predictions and to make physical interpretations of the results (see Appendix A for proof of these relationships).

Fig. 6 presents the L-dependent profiles of each of the three terms in Eq. (6). Our ISS model does a very good job predicting electron fluxes at the lower, if not the lowest, L shells, capturing well over 80% of the observed electron flux variability. The quality of these predictions drops smoothly as L increases beyond $4R_E$, but the model still manages to capture at least 50% of the observed variance, even near the outer limits of the electron radiation belts. The fact that the shared variance between observations and prediction errors is near zero for all but the lowest L shells indicates that our model is nearly optimal. With this in mind, we proceed to look more closely at the individual components of our model predictions in an attempt to gain some physical insights.

Fig. 7 illustrates the existence of at least two distinct dynamic regimes. Even though internal linear dynamics dominate radiation belt electron flux variations at all L shells, this contribution drops rapidly between $L = 4$–$5R_E$, after which it continues to drop somewhat less precipitously out to the highest L shells. Likewise, the impact of solar wind perturbations drops very quickly inside $L \sim 5$, falling to near-zero for $L < 4$. Concurrently, the variance captured by corrective component predictions appears to be relatively minor at the lower L shells, somewhat larger at higher L shells, and undergoes a perceptible, if not very sharp, transition between the two regions just described. This would seem to be at odds with results discussed at the end of the previous section, which showed that the amplitude of corrections tends to follow a profile more similar to internal linear feedback. However, we now recall that most of the corrections at lower L shells constituted significant reconfigurations of the radiation belt state. Once the state adjusts itself to the current conditions, linear feedback once again takes over as the dominant dynamics in this region.

It is also important to remember that corrective contributions can account for more than just missing internal dynamics in our model, and may in fact correct for the absence of relevant external drivers. Assuming it was relatively persistent, a missing driver should manifest itself as a constant source of error, requiring more regular than episodic corrections by $\mathbf{K}\boldsymbol{\varepsilon}(t)$. This helps explain the tendency for the fractional covariance of the corrective predictions to increase along with the relative influence of those external drivers that were included, in spite of the fact that the average amplitude of these corrections

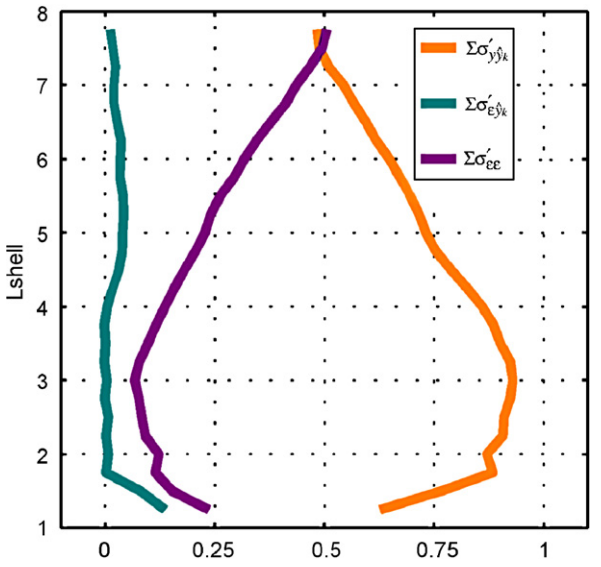

Fig. 6. Fractional covariances between observations and the cumulative ISS prediction ($\sum \sigma'_{y\hat{y}_k}$), between prediction errors and ISS predictions ($\sum \sigma'_{\varepsilon\hat{y}_k}$), and the fractional variance of the prediction errors themselves ($\sum \sigma'_{\varepsilon\varepsilon}$), are all shown as functions of L shell. The sum of all three curves equals unity.

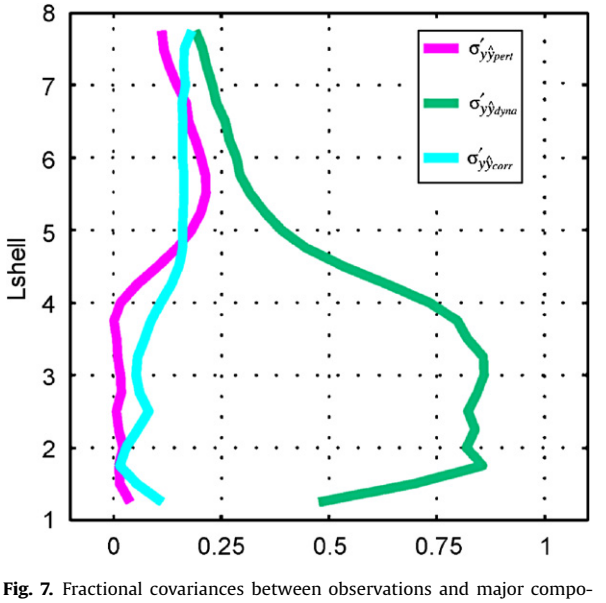

Fig. 7. Fractional covariances between observations and major components of the ISS model predictions are shown as functions of L shell. These include contributions from the solar wind perturbations ($\sigma'_{y\hat{y}_{pert}}$), linear dynamic feedback ($\sigma'_{y\hat{y}_{dyna}}$), and corrections made based on previous prediction errors ($\sigma'_{y\hat{y}_{corr}}$). The sum of all three curves equals $\sum \sigma'_{y\hat{y}_k}$.

decreases. If, on the other hand, a missing external driver is not very persistent, but rather constitutes a rare but significant perturbation to the system (a shock in the solar wind, for example), the correction term will probably not capture all of its dynamic impact to the system. Still, it does correct for some fraction of the variance, and the bias that would otherwise be introduced into the estimated model parameters is reduced.

We noted in Section 2 that both the **A** and **K** matrices contained off-diagonal terms that allowed neighboring states to influence the local state of the radiation belt. We could have plotted a fractional covariance profile for the contribution of every single off-diagonal term, similar to Fig. 7, but this would be difficult to decipher, even for a researcher who was intimately familiar with the physics, data, and ISS model structure being used here. We chose instead to analyze the relative influence of nearest neighbors on each L shell bin. Our approach used a simple algorithm to define an index describing whether the local state was influenced from above (+1), below (-1), or if its dynamics were overwhelmingly dominated by its own previous state (0), or if there was no significant imbalance in the direction of influence (also 0):

1. If the fractional covariance between the flux observed at each L shell and flux predicted from linear feedback/error corrections at the same L shell is more than five times the fractional covariance between observed flux and predictions for either nearest neighbor, a value of 0 is assigned to that L shell.
2. If this is not the case, but the ratio of fractional covariances of the lower/higher neighbors' predictions with local observations is between 80% and 125%, it is assumed that neither neighbor had a significantly stronger impact on the local state, so a value of 0 is once again assigned to that L shell.
3. If the ratio of fractional covariances of the lower/higher neighbors' predictions with local observations is greater than l25%, it is assumed that the lower neighbor had a stronger impact on the local state than the higher neighbor, and a value of -1 is assigned to that L shell.
4. If the ratio of fractional covariances of the lower/higher neighbors' predictions with local observations is less than 80%, it is assumed that the higher neighbor had a stronger impact on the local state than the lower neighbor, and a value of +1 is assigned to that L shell.

The resulting indices might be considered an indication of the direction from which information flows in the system being modeled, and are presented in Fig. 8 as a function of L shell. In general, these results indicate that changes in electron flux at L shells below $5R_E$ are influenced more by what is happening above them than below. Changes in electron flux at L shells above $5R_E$ are influenced more by what is happening below them than above. For the most part, this rule holds true for both linear dynamic feedback and error corrections. While it is true that this "diffusion" of information might just as easily have arisen from anticorrelations between observations and predictions, a closer inspection of the **A** and **K** matrix coefficients showed that they were almost all positive. This implies that we are reproducing a genuine diffusion of $\log_{10}$-flux with the ISS model, if not necessarily diffusion of PSD. One clear exception to this rule occurs between just below $L = 2$, and just above $L = 3$, where linear feedback has an index value equal to 0, implying that simple persistence dominates. Isolated deviations to the rule that occur at $L = 4.5$ and $6R_E$ are probably anomalous.

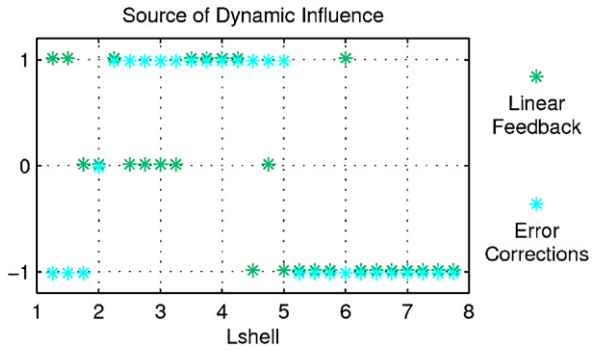

Fig. 8. An L-dependent index equal to +1 indicates that the local electron flux is more influenced by variations in flux occurring above than below itself. An index equal to -1 indicates that variations in local electron flux are more influenced by variations occurring below than above. An index equal to 0 indicates either that recurrence completely dominates the dynamics, or that there is no significant difference between the influence coming from above or below.

It is also worth noting that, while non-persistent effects of linear feedback come mostly from above at the very lowest L shells, the influence of error corrections seem to diffuse upward for $L < 1.75R_E$. If we look back at Fig. 7, a small but significant enhancement in the fractional covariance associated with error corrections can also be seen at these L shells, while a decrease in the fractional covariance associated with linear feedback is observed. One might be inclined to interpret this as some sort of nonlinear dynamical phenomenon in this region. However, if we look even further back, to Fig. 6, there is a clear spike in the covariance between predictions and residuals in this region, which implies that our model is not optimal here. Any physical interpretations in this region should be made with caution until more data become available with which to refine our estimates of the corresponding ISS matrix coefficients.

5. Summary and conclusions

In this study, we optimized a linear state-space model in order to reproduce that portion of the electron radiation belt flux dynamics that resembles linear diffusion, and which is perturbed directly by our three representative solar wind quantities: $|\mathbf{B}_{imf}|$, $|\mathbf{V}_{sw}|$, and ρ_{sw}. Even though our understanding of radiation belt dynamics is far from complete, we can be certain that this model is missing many relevant "inputs", including additional solar wind drivers and higher order, even non-linear, dynamic feedback terms. Rigler et al. (2007) showed how such missing

inputs can lead to biased estimates of those parameters that are explicitly included in empirical models. We therefore introduced a correction term that is optimized concurrently with the deterministic state-space model parameters in order to exploit non-random signals in the prediction error, or innovations, sequence. A perfect corrector (i.e., residuals constitute a zero-mean, time stationary, random sequence) should account for any missing inputs/physics in the system, and the strictly linear dynamics described by our state-space model would be unbiased.

For various reasons a perfect corrector is impossible to attain, so biased model parameters will always exist. Still, the linear state-space model presented in this study is substantially less biased than a similar state-space model that was developed in parallel, but not presented here, and whose coefficients were estimated without innovations-based corrections. As such, we proceeded to analyze the ISS model, as well as its in-sample predictions and prediction error statistics, to better understand what role linear dynamics play in governing the evolution of electron flux in the radiation belts. Input impulse response functions were presented as functions of L shell, and can be considered analogous to the FIR filter profiles used in prior studies. Considering that the solar wind inputs used here served as rough proxies for a variety of solar wind–magnetosphere coupling mechanisms, these showed a clear separation between electromagnetic, represented by $|\mathbf{B}_{\mathrm{imf}}|$, and viscous interactions, represented by $|\mathbf{V}_{\mathrm{sw}}|$. The former leads primarily to reductions in electron flux between $L{\sim}4$–$6.5R_{\mathrm{E}}$, while the latter leads to electron flux enhancements between $L{\sim}5$–$8R_{\mathrm{E}}$. There also appear to be losses related to pressure, represented by ρ_{sw}, but a cursory examination of the actual time-dependent component predictions reveals that these are very episodic, and require extreme changes in the solar wind density to have any significant impact on radiation belt fluxes.

Given the state-space representation of our model, it was fairly straight-forward to conduct simple experiments wherein the state at each L shell was perturbed directly to see how flux might evolve in the absence of solar wind perturbations or flux from neighboring L shells. We found that flux diffuses upward beyond $L{\sim}5R_{\mathrm{E}}$, downward between L $\sim$1.75–$5R_{\mathrm{E}}$, and surprisingly, upward again at the very lowest L shells. This last result is somewhat suspect though, since there are significant correlations between the prediction errors and component predictions at these altitudes, indicating that the model was not fully optimized here.

Component predictions corresponding to each of the solar wind perturbations, linear dynamic feedback, and the innovations correction term were presented next. These more clearly demonstrated the largely separable effects of electromagnetic and viscous interactions between the solar wind and radiation belts, as well as the relative non-impact of changes in solar wind pressure at daily time scales. More importantly, these perturbation effects were almost trivial when compared to the relative influence of linear feedback, which tends to account for the recovery of lost electrons in the heart of the outer belt, and the decay of electron flux after strong enhancements near the inner edge of the outer belt. Contributions from the corrective term tended to be less coherent in both space and time, manifesting as strong but infrequent changes in flux. This is not surprising because once a major correction was made to the model output, the remaining model dynamics take over to propagate those changes forward in time in a linear fashion.

The predictions were subjected to a more robust analysis than simple visual inspection using a metric we refer to as "fractional covariance". Unlike metrics used in previous studies, the fractional covariance is not arbitrarily scaled to a fixed interval (e.g., $\{-11\}$). Rather it very deliberately represents a fraction of the total observed variability in the state variable of interest that is captured by each component prediction. By design, it is perfectly additive, meaning that the fractional covariance of each component prediction will add up to the cumulative fractional covariance. The ISS state-space model captures at least 50% of the observed $\log_{10}$-flux variance at the highest L shells, and increases smoothly downward until it plateaus above 80% between $L{\sim}2$–$3R_{\mathrm{E}}$, and drops rapidly below $L = 2$.

The component prediction statistics offer substantially more insight into the actual physics of the system than do the combined prediction statistics. While linear feedback dominates the flux dynamics at all L shells, there is a rapid transition between $L{\sim}4$–5 where the share of variance that can be attributed to linear feedback drops from $\sim$80% to 30% and less. There is also a corresponding increase in the combined effects of solar wind perturbations across this same region. The flux contribution from model corrections appears to be fairly low inside of $L = 4$, and grows steadily beyond this point until it surpasses linear solar wind perturbations, and is nearly comparable to recurrence at the highest L shells.

The fact that corrections tend to be smaller at higher L shells, and yet still account for a fairly large fraction of the total variance of observed fluxes, suggests that these probably account for missing, and potentially nonlinear, combinations of solar wind drivers. Corrections made at lower altitudes tend to be episodic, but quite substantial when they do occur, and because linear feedback already dominates the overall radiation belt dynamics here, we believe they account more for major reconfigurations in the radiation belts that occur during intense solar wind events. It is unlikely that such events could ever be captured by our state-space model in any deterministic sense without expanding the current state dimension significantly, so the correction term is critical for reducing model bias in this region.

The fractional covariances were further processed using a simple algorithm to generate an index for each L shell that indicated whether the local state was mostly persistent, or whether it was influenced more strongly by variations in flux occurring above or below itself. Aside from a few anomalous data points, it is clear that radiation belt fluxes above $L = 5$ are influenced from below, while fluxes below $L = 5$ are influenced from above, or are just so persistent as to be uninfluenced at all by what is happening at their neighboring states. These results are

consistent with the state impulse response functions described earlier, and indicate that flux diffuses downward below $L = 5$, upward above $L = 5$, and is influenced equally from both directions near $L = 5$, which corresponds to the heart of the electron radiation belts.

We conclude by reminding the reader that all the results presented so far relate to electron flux, not PSD, a more physically meaningful quantity that is difficult to extract accurately from the SAMPEX data. However, relatively long duration time-averaging, a recognized global coherence in the electron radiation belts (Kanekal et al., 2001), and perhaps most importantly, the fact that all dynamics discussed up to this point describe variations about an estimated baseline, and not any absolute values, gives us reason to believe that the relative relationships between states and inputs, if not the exact model parameters, will remain unchanged if such a transformation from flux to PSD were to be performed. If this is true, the work presented here strongly suggests that radial diffusion/transport does play a significant role in radiation belt dynamics, but not necessarily as an acceleration mechanism, especially above $L = 5R_E$. In addition, there are several localized un-modeled dynamical processes being reproduced, at least in part, by the innovations correction term, indicating non-diffusive and potentially nonlinear electron flux dynamics in the Earth's radiation belts.

Acknowledgments

The authors wish to thank D. Vassiliadis, R.S. Weigel, T.P. O'brien, and M. Wiltberger for valuable feedback pertaining to this study. OMNI2 solar wind data came from the National Space Science Data Center at http://omniweb.gsfc.nasa.gov. Electron flux data were generously provided by the SAMPEX data team, via S. Kanekal. This material is based on work supported by the Center for Integrated Space weather Modeling, an NSF Science and Technology Center, award number ATM-0120950.

Appendix A. : Fractional covariance as a prediction metric

The covariance between a single observed dependent variable (Y), and potentially several observed independent variables (X_k), can be interpreted as a fraction of the observed variability in Y that can be attributed to variations in X_k if it is properly scaled.

We start with a simple two-variable linear regression, complete with an intercept (α) and prediction error term (ε) for generality, but the forthcoming conclusions will be scalable to an arbitrary number of independent variables:

$$Y = \alpha + \beta_1 X_1 + \beta_2 X_2 + \varepsilon. \tag{7}$$

Multiply both sides of Eq. (7) by Y to get:

$$YY = \alpha Y + \beta_1 X_1 Y + \beta_2 X_2 Y + \varepsilon Y. \tag{8}$$

Next, apply the expectation operator E to this relationship, accounting for its distributive properties, and the fact that α, β_1, and β_2 are all constants, while the remaining terms are not:

$$E\{YY\} = \alpha E\{Y\} + \beta_1 E\{X_1 Y\} + \beta_2 E\{X_2 Y\} + E\{\varepsilon Y\}. \tag{9}$$

Finally, substitute Eq. (7) into Eq. (9)'s final term to give:

$$\begin{aligned} E\{YY\} &= \alpha E\{Y\} + \beta_1 E\{X_1 Y\} + \beta_2 E\{X_2 Y\} \\ &\quad + E\{\varepsilon(\alpha + \beta_1 X_1 + \beta_2 X_2 + \varepsilon)\} \\ &= \alpha E\{Y\} + \beta_1 E\{X_1 Y\} + \beta_2 E\{X_2 Y\} \\ &\quad + \alpha E\{\varepsilon\} + \beta_1 E\{X_1 \varepsilon\} + \beta_2 E\{X_2 \varepsilon\} + E\{\varepsilon\varepsilon\}. \end{aligned} \tag{10}$$

At this point, the expectation terms in Eq. (10) are becoming somewhat unwieldy, and to be perfectly accurate, they only apply to entire populations of random variables anyway. In reality, only limited samples of a population are available to any statistical study. The expected value of any sampled variable is simply its mean, μ. With this in mind, we can rewrite Eq. (10) for a sample instead of the entire population:

$$\begin{aligned} \mu_{YY} &= \alpha\mu_Y + \beta_1 \mu_{X_1 Y} + \beta_2 \mu_{X_2 Y} + \alpha\mu_\varepsilon \\ &\quad + \beta_1 \mu_{X_1 \varepsilon} + \beta_2 \mu_{X_1 \varepsilon} + \mu_{\varepsilon\varepsilon}. \end{aligned} \tag{11}$$

Now, the sample covariance between any two random variables (σ_{XY}) can be defined as the mean of the product of two random variables, which have in turn had their own means removed.

$$\begin{aligned} \sigma_{XY} &= \overline{(X - \mu_X)(Y - \mu_Y)} \quad \ldots \text{or} \ldots \\ &= \mu_{XY} - \mu_X \mu_Y. \end{aligned} \tag{12}$$

Rearranging the second form of Eq. (12) so μ_{XY} can be substituted back into Eq. (11) gives

$$\begin{aligned} \sigma_{YY} + \mu_Y \mu_Y &= \alpha\mu_Y + \beta_1 \sigma_{X_1 Y} + \beta_1 \mu_{X_1} \mu_Y + \beta_2 \sigma_{X_2 Y} \\ &\quad + \beta_2 \mu_{X_2} \mu_Y + \alpha\mu_\varepsilon + \beta_1 \sigma_{X_1 \varepsilon} + \beta_1 \mu_{X_1} \mu_\varepsilon \\ &\quad + \beta_2 \sigma_{X_2 \varepsilon} + \beta_2 \mu_{X_2} \mu_\varepsilon + \sigma_{\varepsilon\varepsilon} + \mu_\varepsilon \mu_\varepsilon. \end{aligned} \tag{13}$$

Fortuitously, the product of means on the LHS of Eq. (13) equals exactly the sum of every product on the RHS that contains a mean. No proof is provided, but this is easy to confirm numerically. This allows us to write

$$\begin{aligned} \sigma_{YY} &= \beta_1 \sigma_{X_1 Y} + \beta_2 \sigma_{X_2 Y} + \beta_1 \sigma_{X_1 \varepsilon} \\ &\quad + \beta_2 \sigma_{X_2 \varepsilon} + \sigma_{\varepsilon\varepsilon}. \end{aligned} \tag{14}$$

Noting that the covariance of any variable with itself is just the variance of that single variable, we simply divide both sides by the variance of the observed output, σ_{YY}, to arrive at

$$1 = \beta_1 \frac{\sigma_{X_1 Y}}{\sigma_{YY}} + \beta_2 \frac{\sigma_{X_2 Y}}{\sigma_{YY}} + \beta_1 \frac{\sigma_{X_1 \varepsilon}}{\sigma_{YY}} + \beta_2 \frac{\sigma_{X_2 \varepsilon}}{\sigma_{YY}} + \frac{\sigma_{\varepsilon\varepsilon}}{\sigma_{YY}}. \tag{15}$$

If β_1 and β_2 were estimated such that there existed no correlation between the prediction residuals and the inputs, the third and fourth terms in Eq. (15) could be neglected. This is not generally the case, so these terms are left in place. These relationships hold regardless of the number of observed input variables, so we can write:

$$1 = \sum_{k=1}^{K} \beta_k \frac{\sigma_{X_k Y}}{\sigma_{YY}} + \sum_{k=1}^{K} \beta_k \frac{\sigma_{X_k \varepsilon}}{\sigma_{YY}} + \frac{\sigma_{\varepsilon\varepsilon}}{\sigma_{YY}}. \tag{16}$$

Simply stated, Eq. (16) reflects the fact that the sum of the β-weighted fractional covariances of the independent

variables X_k with the dependent variable Y, added to the sum of β-weighted fractional covariances of X_k with the prediction residuals ε, added to the fractional variance of the residuals, equals unity. Finally, if all of the β_k's happen to equal 1, implying that the independent variables in Eq. (7) are actually just additive components of a cumulative prediction (i.e., $\hat{Y} = \alpha + \sum_{k=1}^{K} \hat{Y}_k$), we can write each component's fractional contribution to Y's variance as

$$1 = \sum_{k=1}^{K} \frac{\sigma_{\hat{Y}_k Y}}{\sigma_{YY}} + \sum_{k=1}^{K} \frac{\sigma_{\hat{Y}_k \varepsilon}}{\sigma_{YY}} + \frac{\sigma_{\varepsilon\varepsilon}}{\sigma_{YY}}. \quad (17)$$

If so-called standardized regression coefficients, or the ratio of the standard deviation of each component prediction to the standard deviation of the observed dependent variable (i.e., $\beta_k' = \sigma_{\hat{Y}_k}/\sigma_Y$), are multiplied by corresponding correlation coefficients $r_{\hat{Y}_k Y}$, one also gets a fractional covariance. This was as demonstrated almost accidentally by Rigler et al. (2007) in their efforts to determine R^2 (i.e., the first RHS term in Eq. (18)) from the regression coefficients of multi-input FIR linear prediction filter models.

$$1 = \sum_{k=1}^{K} \beta_k' r_{\hat{Y}_k Y} + \sum_{k=1}^{K} \beta_k' r_{\hat{Y}_k \varepsilon} + \frac{\sigma_{\varepsilon\varepsilon}}{\sigma_{YY}}. \quad (18)$$

It should be noted that Rigler et al. (2007) did not actually calculate β_k in this manner, but instead regressed the standardized dependent variable (Y') on the standardized component predictions ($\hat{Y}_k'$). Unless correlations between these component predictions and the residuals happen to be zero, the relative significance of each β_k must change because OLS regression re-optimizes the parameters in order to de-correlate predictions and residuals. Plugging these re-optimized β_k values into Eq. (18) will lead to a value other than unity on the LHS, and the term "fractional (co)variance" loses all meaning. Rigler et al. (2007) were not aware of this nuance at the time of publication, and so were lucky to have chosen an estimation algorithm (i.e., OLS) for their FIR models that guaranteed zero correlation between residuals and component predictions.

References

Akasofu, S.-I., 1979. Interplanetary energy flux associated with magnetospheric substorms. Planetary and Space Science 27, 425–431.

Arnoldy, R.L., 1971. Signature in the interplanetary medium for substorms. Journal of Geophysical Research 76, 5189–5200.

Baker, D.N., Blake, J.B., Klebesadel, R.W., Higbie, P.R., 1986. Highly relativistic electrons in the earth's outer magnetosphere 1: lifetimes and temporal history 1979–1984. Journal of Geophysical Research 91, 4265–4276.

Baker, D.N., McPherron, R.L., Cayton, T.E., Klebesadel, R.W., 1990. Linear prediction filter analysis of relativistic electron properties at 6.6R_E. Journal of Geophysical Research 95, 15,133–15,140.

Baker, D.N., Mason, G.M., Figueroa, O., Colon, G., Watzin, J.G., Aleman, R.M., 1993. An overview of the solar, anomalous, and magnetospheric particle explorer (SAMPEX) mission. IEEE Transactions on Geoscience and Remote Sensing 31, 531–541.

Bargatze, L.F., Baker, D.N., McPherron, R.L., Hones Jr., E.W., 1985. Magnetospheric impulse response for many levels of geomagnetic activity. Journal of Geophysical Research 90, 6387–6394.

Clauer, C.R., 1986. The technique of linear prediction filters applied to studies of solar wind–magnetosphere coupling. In: Kamide, S. (Eds.), Solar Wind–Magnetosphere Coupling, pp. 39–57.

Cook, W.R., Cummings, A.C., Cummings, J.R., Garrard, T.L., Kecman, B., Mewaldt, R.A., Selesnick, R.S., Stone, E.C., Baker, D.N., von Rosenvinge, T.T., Blake, J.B., Callis, L.B., 1993. PET: A proton/electron telescope for studies of magnetospheric, solar, and galactic particles. IEEE Transactions on Geoscience and Remote Sensing 31, 565–571.

Iyemori, T., Maeda, H., Kamei, T., 1979. Impulse response of geomagnetic indices to interplanetary magnetic field. Journal of Geomagnetism and Geoelectricity 31, 1–9.

Kanekal, S.G., Baker, D.N., Blake, J.B., 2001. Multisatellite measurements of relativistic electrons: global coherence. Journal of Geophysical Research 106, 29721–29732.

Kim, H.J., Chan, A.A., 1997. Fully adiabatic changes in storm time relativistic electron fluxes. Journal of Geophysical Research 102, 22107–22116.

King, J.H., Papitashvili, N.E., 2005. Solar wind spatial scales in and comparisons of hourly wind and ACE plasma and magnetic field data. Journal of Geophysical Research 110.

Li, X., Baker, D.N., Temerin, M., Cayton, T.E., Reeves, G.D., Christensen, R.A., Blake, J.B., Looper, M.D., Nakamura, R., Kanekal, S.G., 1997. Multisatellite observations of the outer zone electron variation during the November 3–4, 1993, magnetic storm. Journal of Geophysical Research 102, 14123–14140.

Ljung, L., 1979. Asymptotic behavior of the extended Kalman filter as a parameter estimator for linear systems. IEEE Transactions on Automatic Control ac-24, 36–50.

Ljung, L., 1999. System Identification: Theory for the User. PTF, Prentice Hall.

Ljung, L., Söderström, T., 1983. Theory and Practice of Recursive Identification. MIT Press, Cambridge, MA.

Nagai, T., 1988. "Space weather forecast": prediction of relativistic electron intensity at synchronous orbit. Geophysical Research Letters 15, 425–428.

Paulikas, G.A., Blake, J.B., 1979. Effects of the solar wind on magnetospheric dynamics: energetic electrons at the synchronous orbit. In: Qualitative Modeling of Magnetospheric Processes, Geophysical Monograph Series, pp. 180–202.

Perreault, P., Akasofu, S.-I., 1978. A study of geomagnetic storms. Geophysical Journal of the Royal Astronomical Society 54, 547–573.

Pindyck, R.S., Rubinfeld, D.L., 1991. Econometric Models and Economic Forecasts. McGraw-Hill, New York.

Priestley, M.B., 1988. Non-linear and Non-Stationary Time Series Analysis. Academic Press.

Rigler, E.J., Baker, D.N., Weigel, R.S., Vassiliadis, D., Klimas, A.J., 2004. Adaptive linear prediction of radiation belt electrons using the Kalman filter. Space Weather 2.

Rigler, E.J., Baker, D.N., Weigel, R.S., Vassiliadis, D., 2005. Solar wind-driven electron radiation belt response functions at 100-min time scales. Advances in Space Research 36.

Rigler, E.J., Wiltberger, M., Baker, D.N., 2007. Radiation belt electrons respond to multiple solar wind inputs. Journal of Geophysical Research 112.

Shynk, J.J., 1989. Adaptive IIR filtering. IEEE ASSP Magazine, doi:10.1109/53.29644.

Vassiliadis, D., Klimas, A.J., Kanekal, S.G., Baker, D.N., Weigel, R.S., 2002. Long-term-average solar cycle, and seasonal response of magnetospheric energetic electrons to the solar wind speed. Journal of Geophysical Research 107.

Vassiliadis, D., Fung, S.F., Klimas, A.J., 2005. Solar, interplanetary, and magnetospheric parameters for the radiation belt energetic electron flux. Journal of Geophysical Research 110.

Vasyliunas, V.M., Kan, J.R., Siscoe, G.L., Akasofu, S.-I., 1982. Scaling relations governing magnetospheric energy transfer. Planetary and Space Science 30, 359–365.