Improved electricity demand forecasts using state-space, time-varying parameter models

Evan Bagley and Bill Shobe
University of Virginia
10-14-18

Dominion’s Demand Forecast

- Like other regulated utilities, Dominion Energy is paid a guaranteed ROI of 10% on capital
- And Dominion does its own demand forecast
- At least since the Great Recession, Dom has been way over-forecasting demand growth

Unsurprising result

- In 2013, the Virginia State Corporation Commission granted a Certificate of Public Convenience and Necessity (CPCN) for the Brunswick Power Station.
- Dom’s CPCN application expressly relied on its 2012 forecast demand.
- By 2017, the Company’s forecast was already too high by the entire annual output of the Brunswick plant.

What’s Dom getting so wrong

- Electricity demand is considered strongly linked to general economic growth
  - Residential, commercial and industrial segments
- It also has a strong seasonal component related to heating and cooling demand
  - In Virginia, heating and cooling is by heat pump
- Dom’s model fits a standard OLS regression model using gross state product as the trend variable
  - The model is fit to 30 years or more of data

Dominion IRP Annual Sales Forecasts: 2012-2018

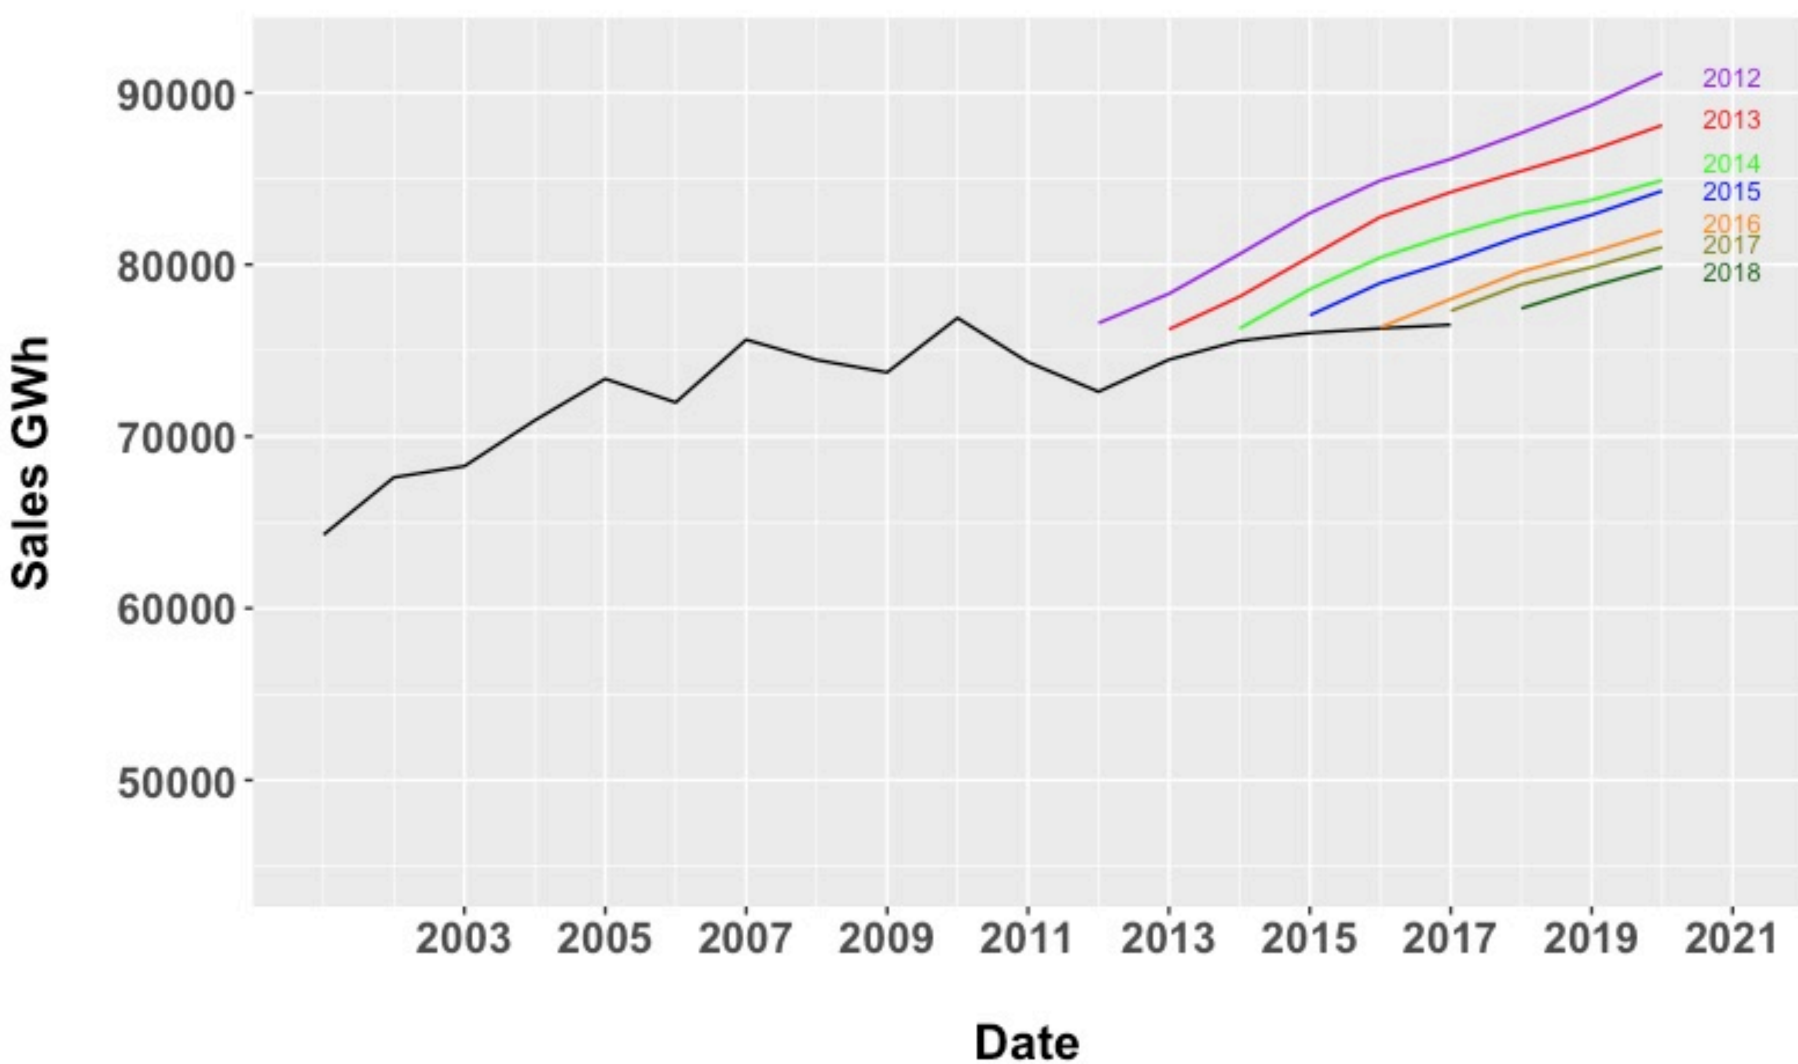

Pattern of over-forecasts

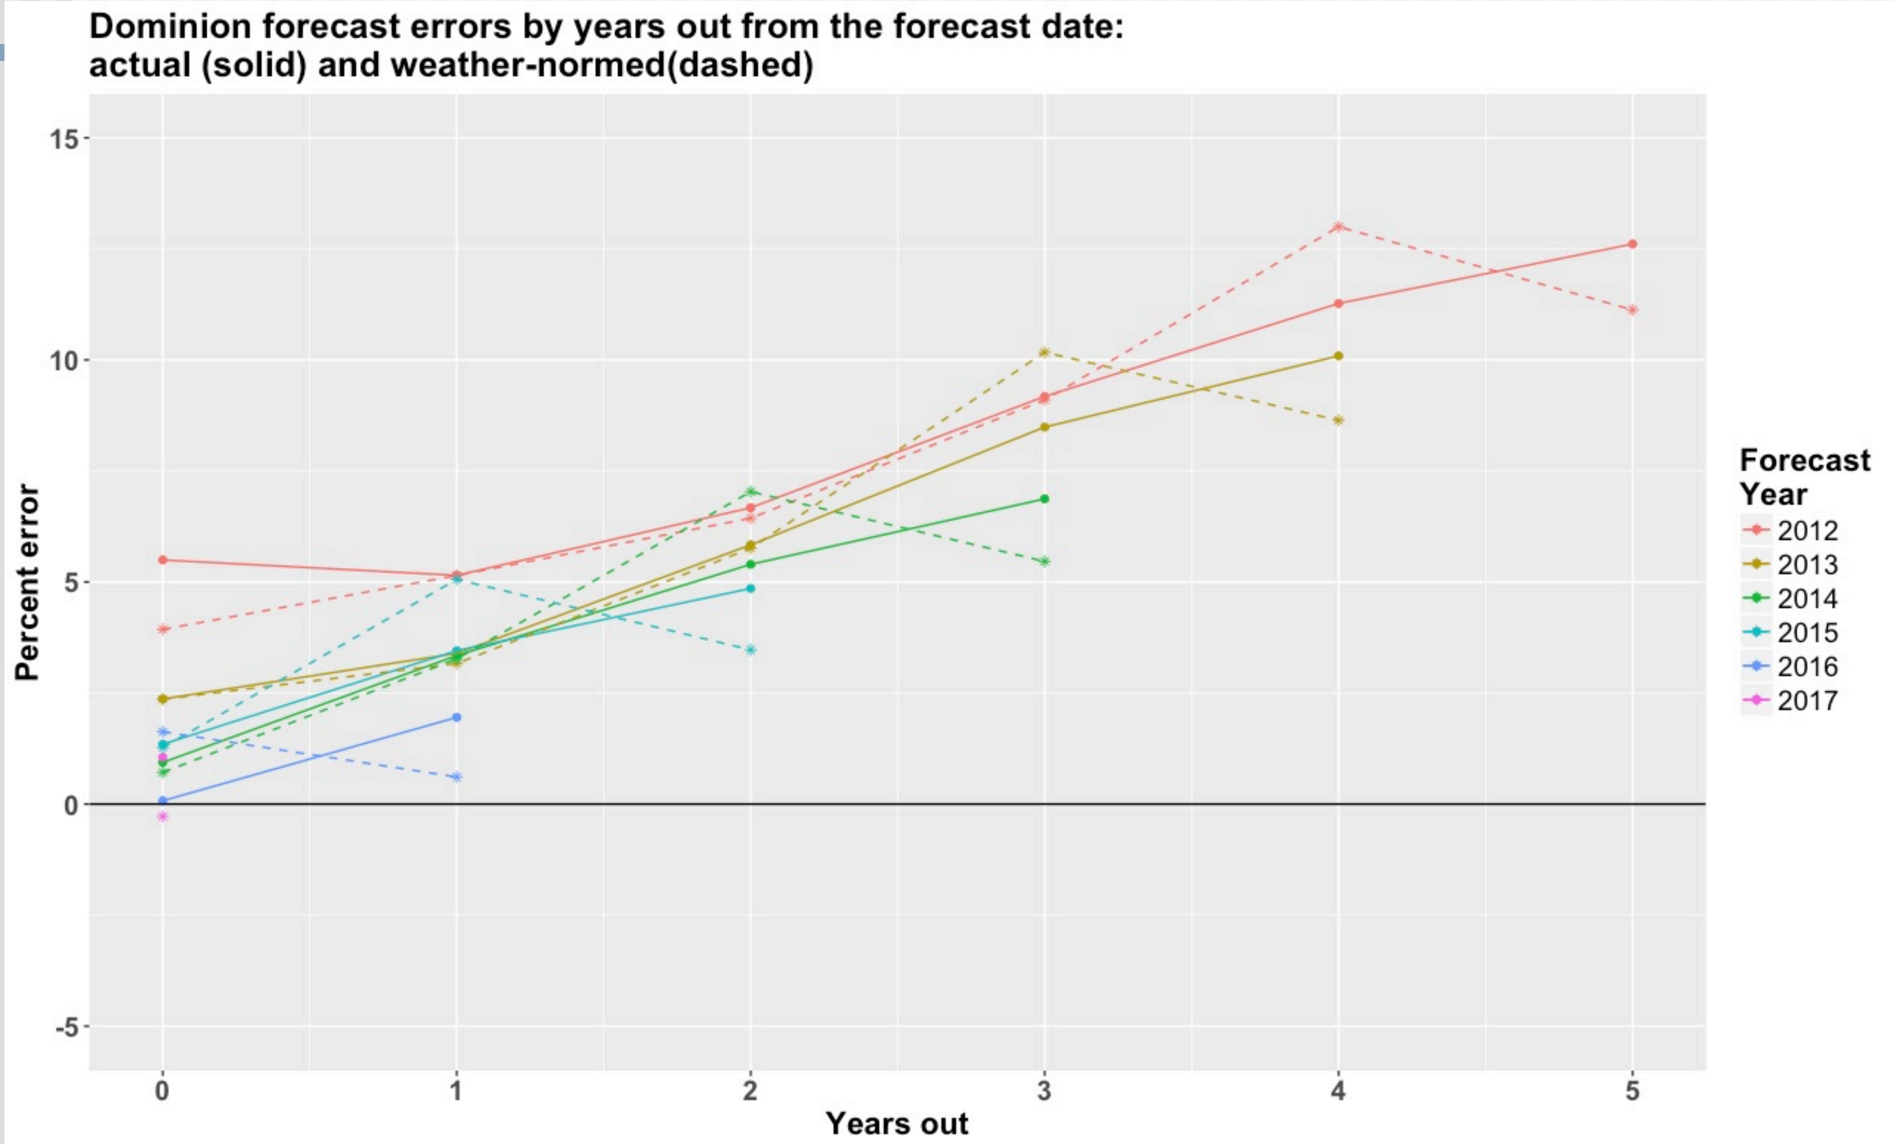

Our challenge: develop a cheap alternative forecast using public data

- We do not have access to customer data on customer appliance and usage data
- An Andrews structural stability test showed a strong peak in August of 2008
- Rolling regressions showed declining sensitivity of demand to GSP from 2008 to 2016
- This suggests one possible approach: truncate the data at the structural break

Improved forecast performance, but…

- We employed a very simple model:

Daily Demand = f(VCI, heating and cooling
degree days, monthly dummies)
where VCI = Philly Fed Coincident Index for VA

- Using post-break data, we had superior fit and out-of-sample forecast performance
  - MAPE cut in half over Dom's complicated but error-prone model

Unsatisfying, ad hoc procedure

- While the Andrews test did not show additional structural breaks, rolling regressions indicated continued parameter instability
- But choosing a shorter data window is ad hoc and unsatisfying
- The continued parameter instability lead us to explore time-varying coefficient models
- We wanted to explicitly account for our lack of information on important demand drivers

The State Space Specification

The model

$$y_t = X_t \beta_t + e_t$$

$$\beta_t = \bar{u} + F\beta_{t-1} + v_t$$

where:

$$e_t \sim i.i.d.N(0, R)$$

$$v_t \sim i.i.d.N(0, Q)$$

If $\bar{u} = 0$ and F is I_k then elements of β follows a random walk.
If F is diagonal with values less than 1, elements of β follows a stationary AR(1) process.

Estimation: the Kalman filter

- This model can be estimated as a sequence of GLS estimations, but the Kalman filter is much easier and more computationally efficient
- The idea is quite simple:
- With each new observation, you ask:
  1. How far is the new y_t from what we forecast in t-1?
  2. How much of this error is due to e_t and to v_t?
  3. What does this tell us about our estimate of the current coefficient vector and, hence, the forecast of y_{t+1}?

Kalman Updating: the magic of K_t

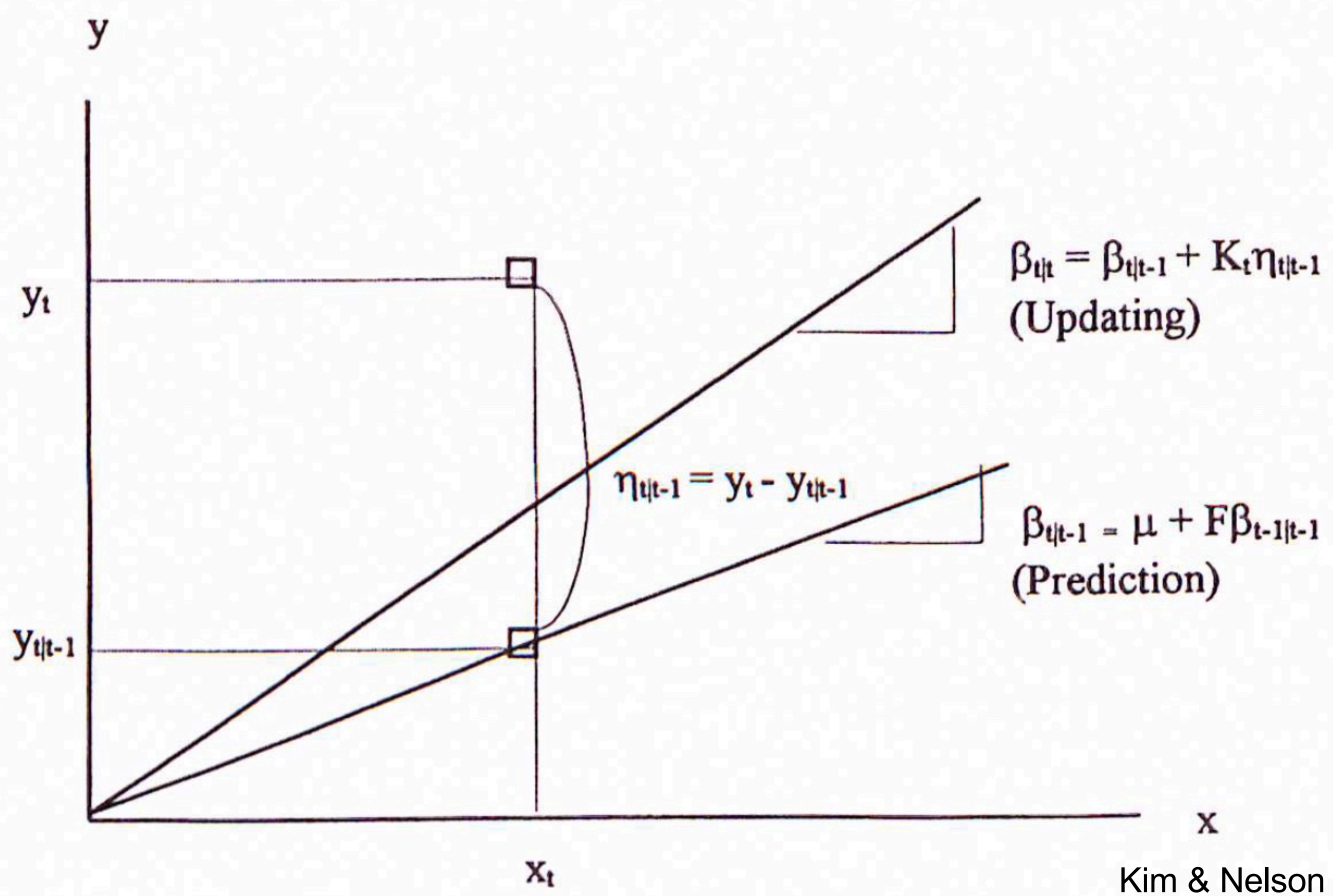

Kim & Nelson (1999)

Where K is a weight based on the share of the total error due to uncertainty over beta (scaled by X_t) << the Kalman magic

Easy to implement in RATS

- Regression Analysis of Time-Series (RATS)
- One command, DLM does all the hard work
  - Read in the data
  - Run an OLS to get some initial variance estimates
  - Run DLM
  - Graph results
- No other package makes state space models this easy

```
open data sales_tvp_data.xls
calendar(m) 1990:1
data(format=xls,org=columns) 1990:1 2033:12 {var list}

* Normal Regression
linreg res_sales
# coinc hpd cpd Jan Feb Mar Apr May Jun Aug Jul Sep Oct Nov Dec

equation(lastreg) mdeq
compute swbase=%seesq*%xx
dec real swscale
dec real sigmae
compute sigmae=.5*sqrt(%seesq)
compute swscale = .01
nonlin sigmae swscale

dlm(y=res_sales,c=%eqnxvector(mdeq,t),sw=swscale*swbase,sv
=sigmae^2,presample=diffuse,method=bfgs,yhat=yhat) 1990:1
2033:12 xstates vstates
```

The Coefficients ARE Changing

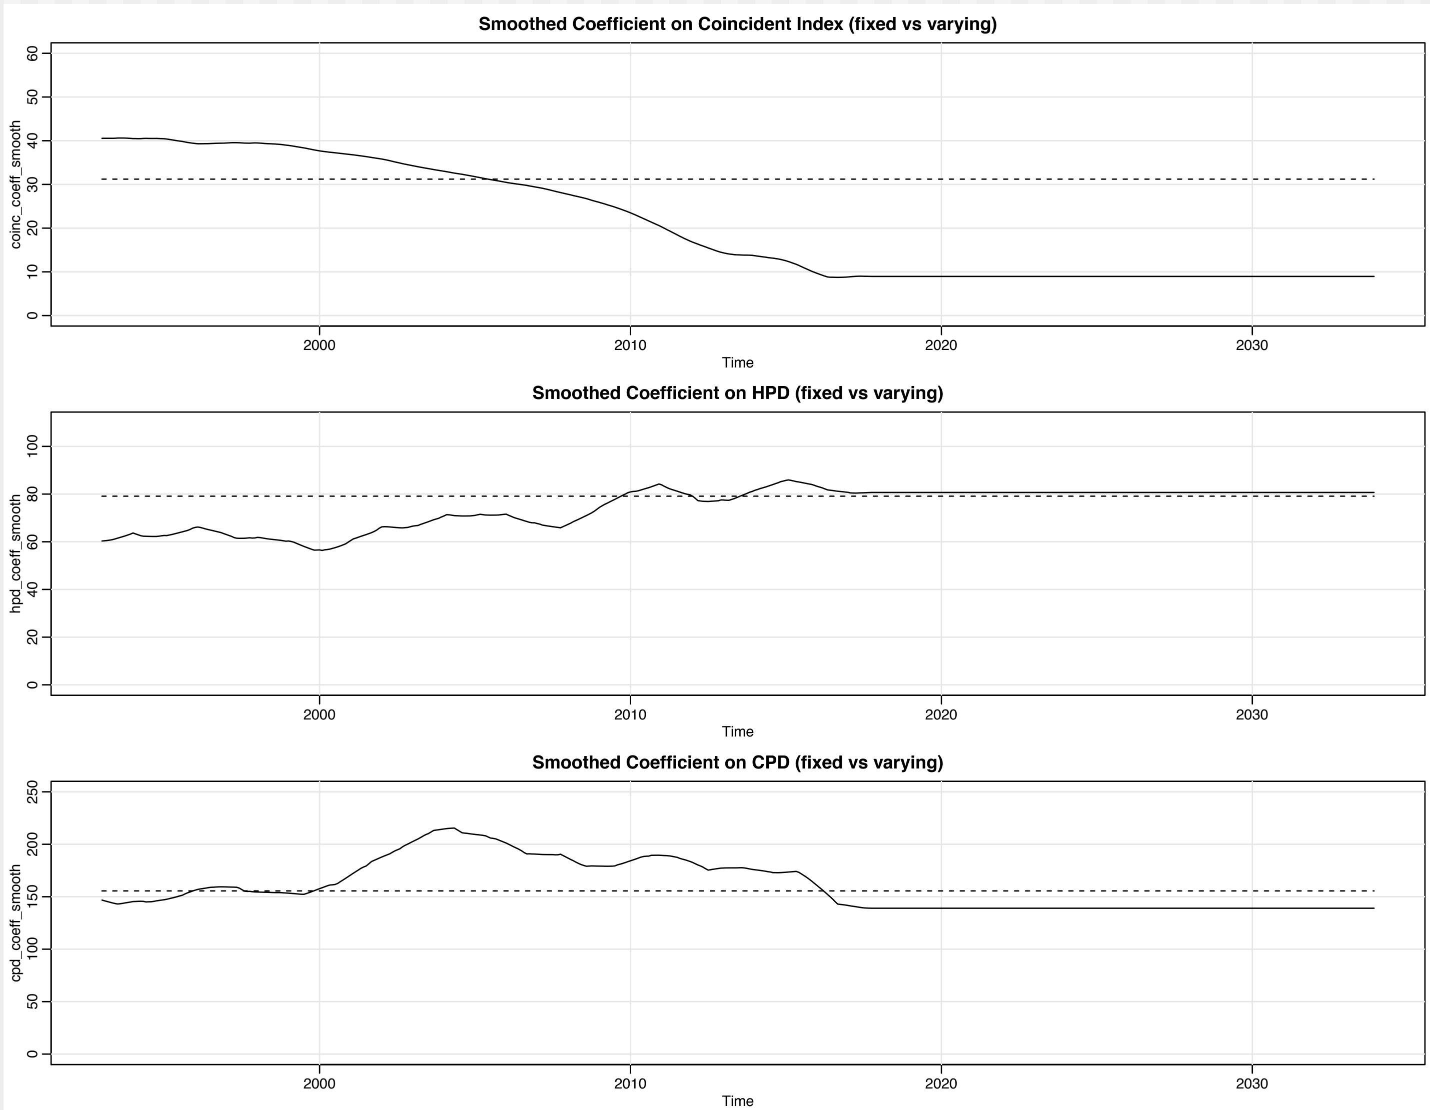

But A Worrisome Result: Why the difference?

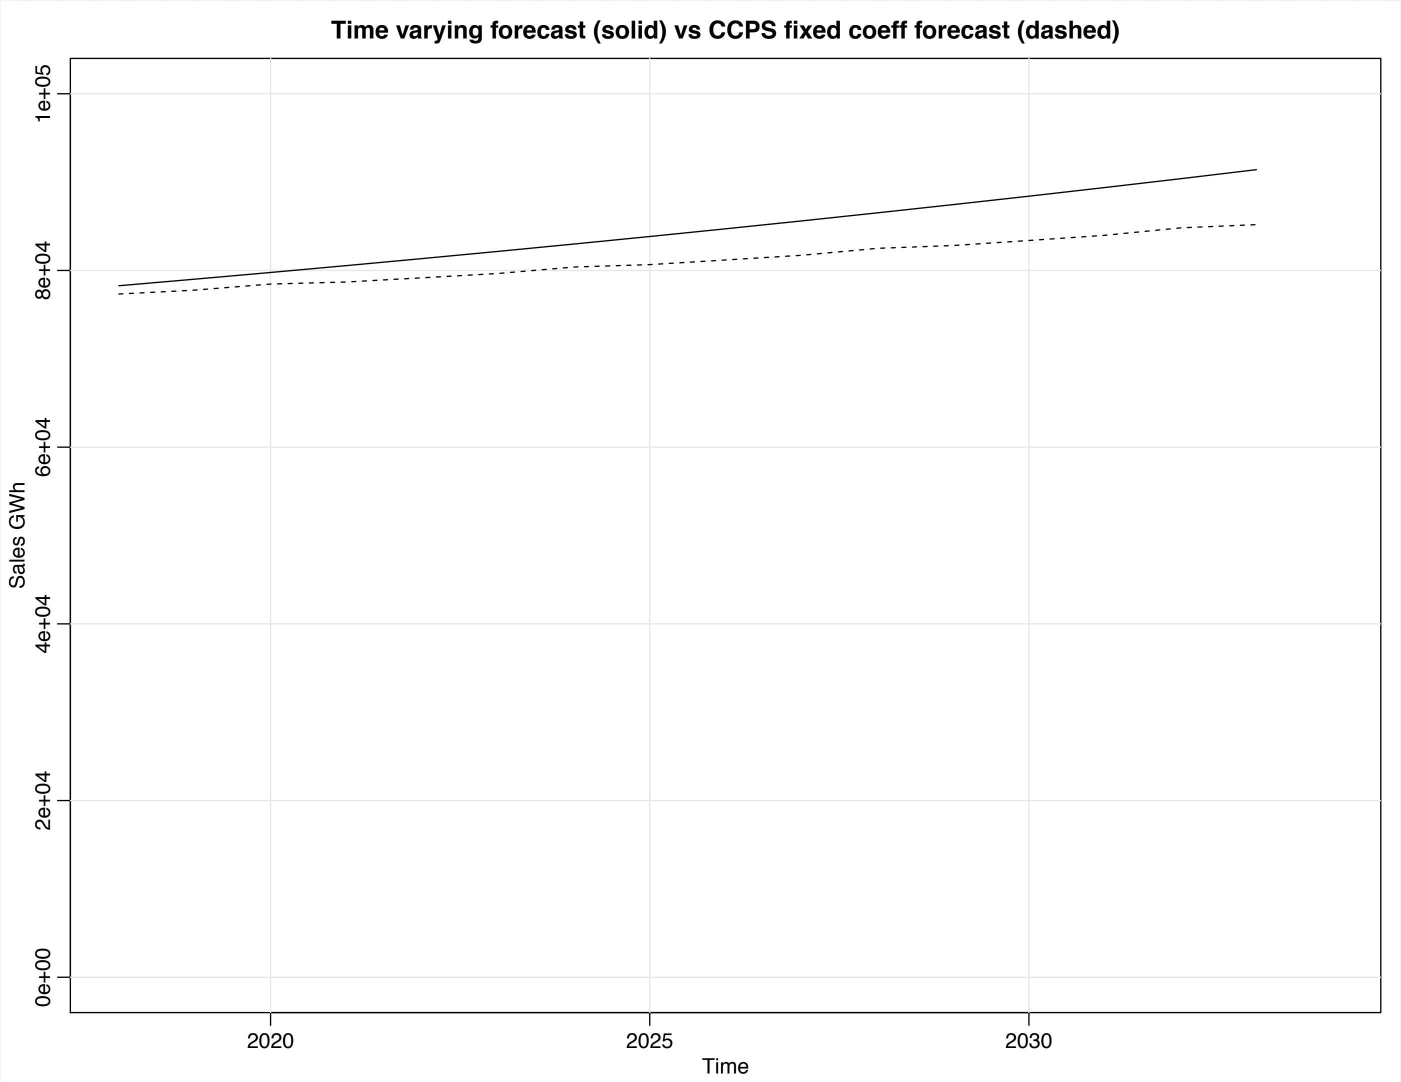

Breaking Demand into Components

- To discover why our post structural break model was lower, we broke demand down into residential, commercial and other
- We then applied our model separately to the components
- We found that the residential demand forecast is flat (due to retrofits and more efficient new units)
  - This forecast very closely matched our linear, truncated data estimation.

Commercial Demand

- Commercial is growing but for a very specific reason: data center sales
- Since data centers and other commercial have different underlying growth processes, we sub-divided again
- We estimated commercial demand with and without data centers

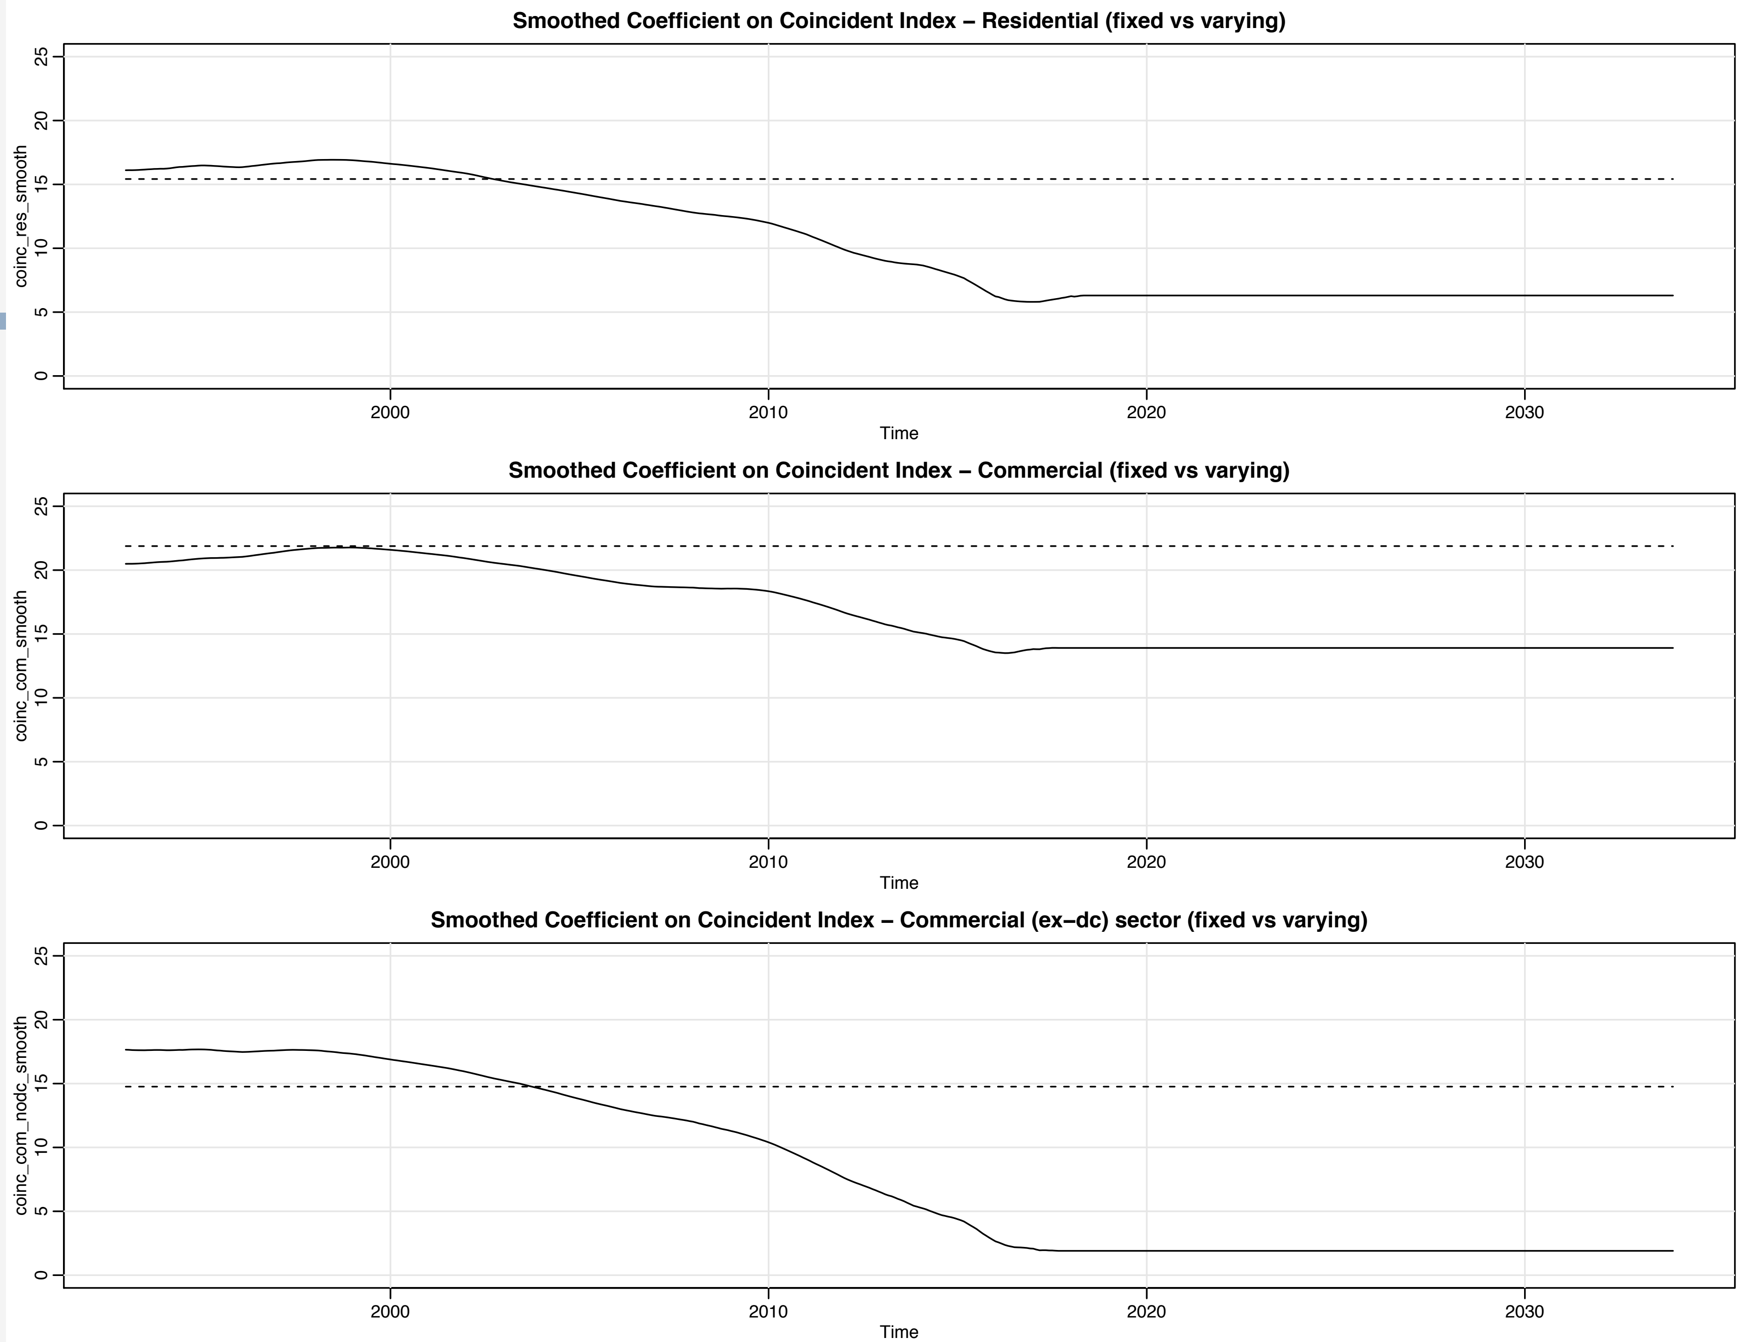
Smoothed Coefficient on Coincident Index – Residential (fixed vs varying)
coinc_res_smooth
0
5
10
15
20
25
2000
2010
2020
2030
Time
Smoothed Coefficient on Coincident Index – Commercial (fixed vs varying)
coinc_com_smooth
0
5
10
15
20
25
2000
2010
2020
2030
Time
Smoothed Coefficient on Coincident Index – Commercial (ex–dc) sector (fixed vs varying)
coinc_com_nodc_smooth
0
5
10
15
20
25
2000
2010
2020
2030
Time

All Commercial Demand Growth is in Data Centers

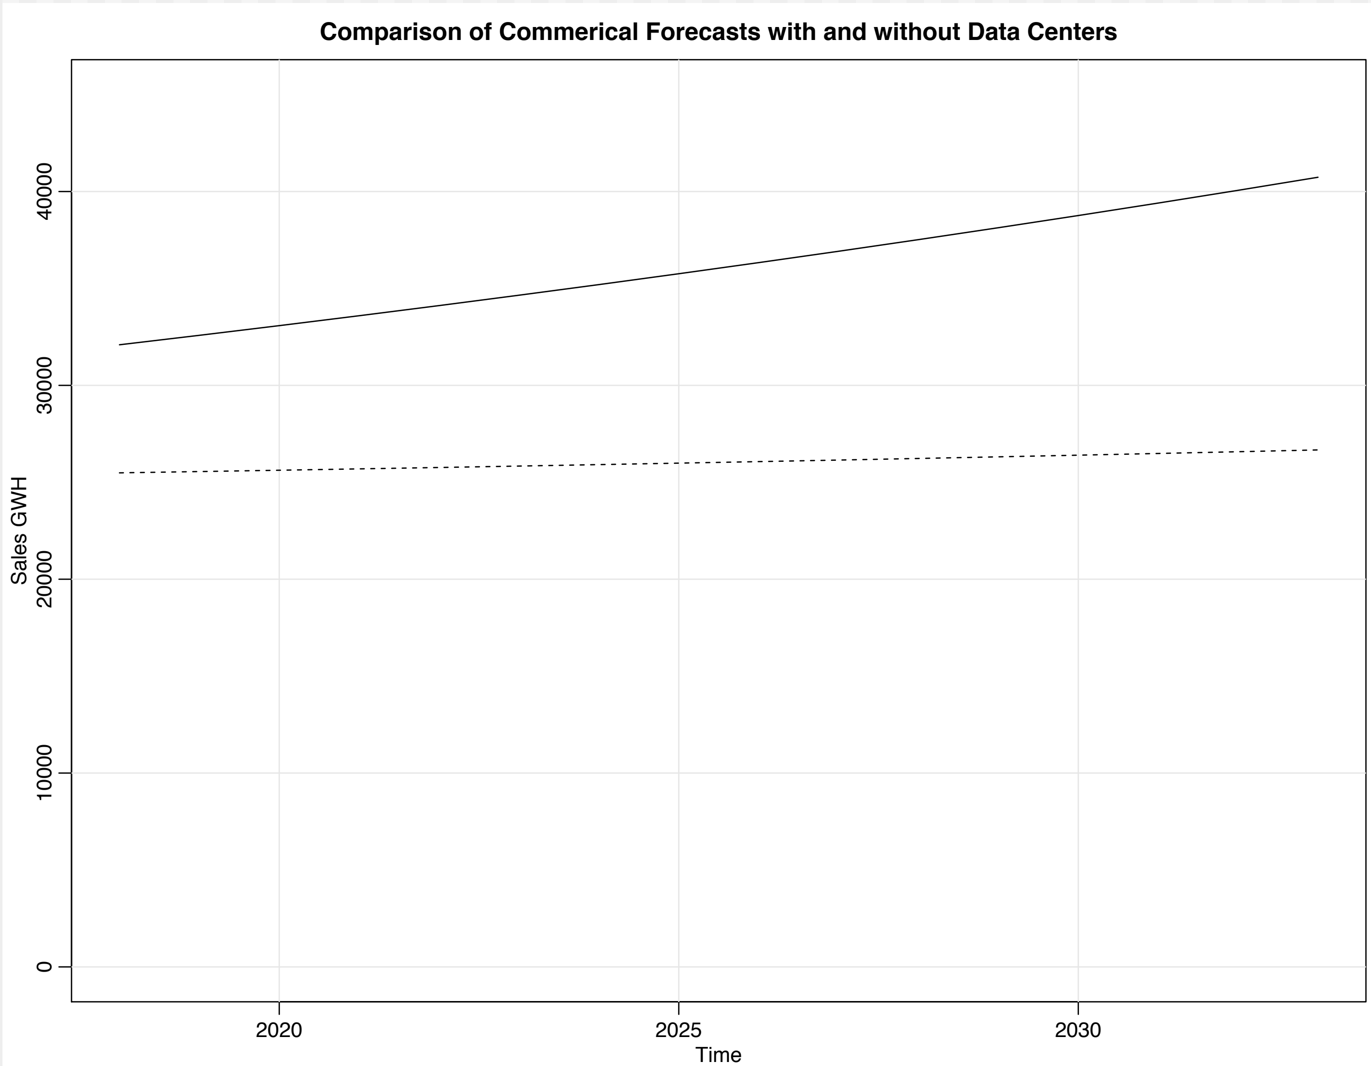

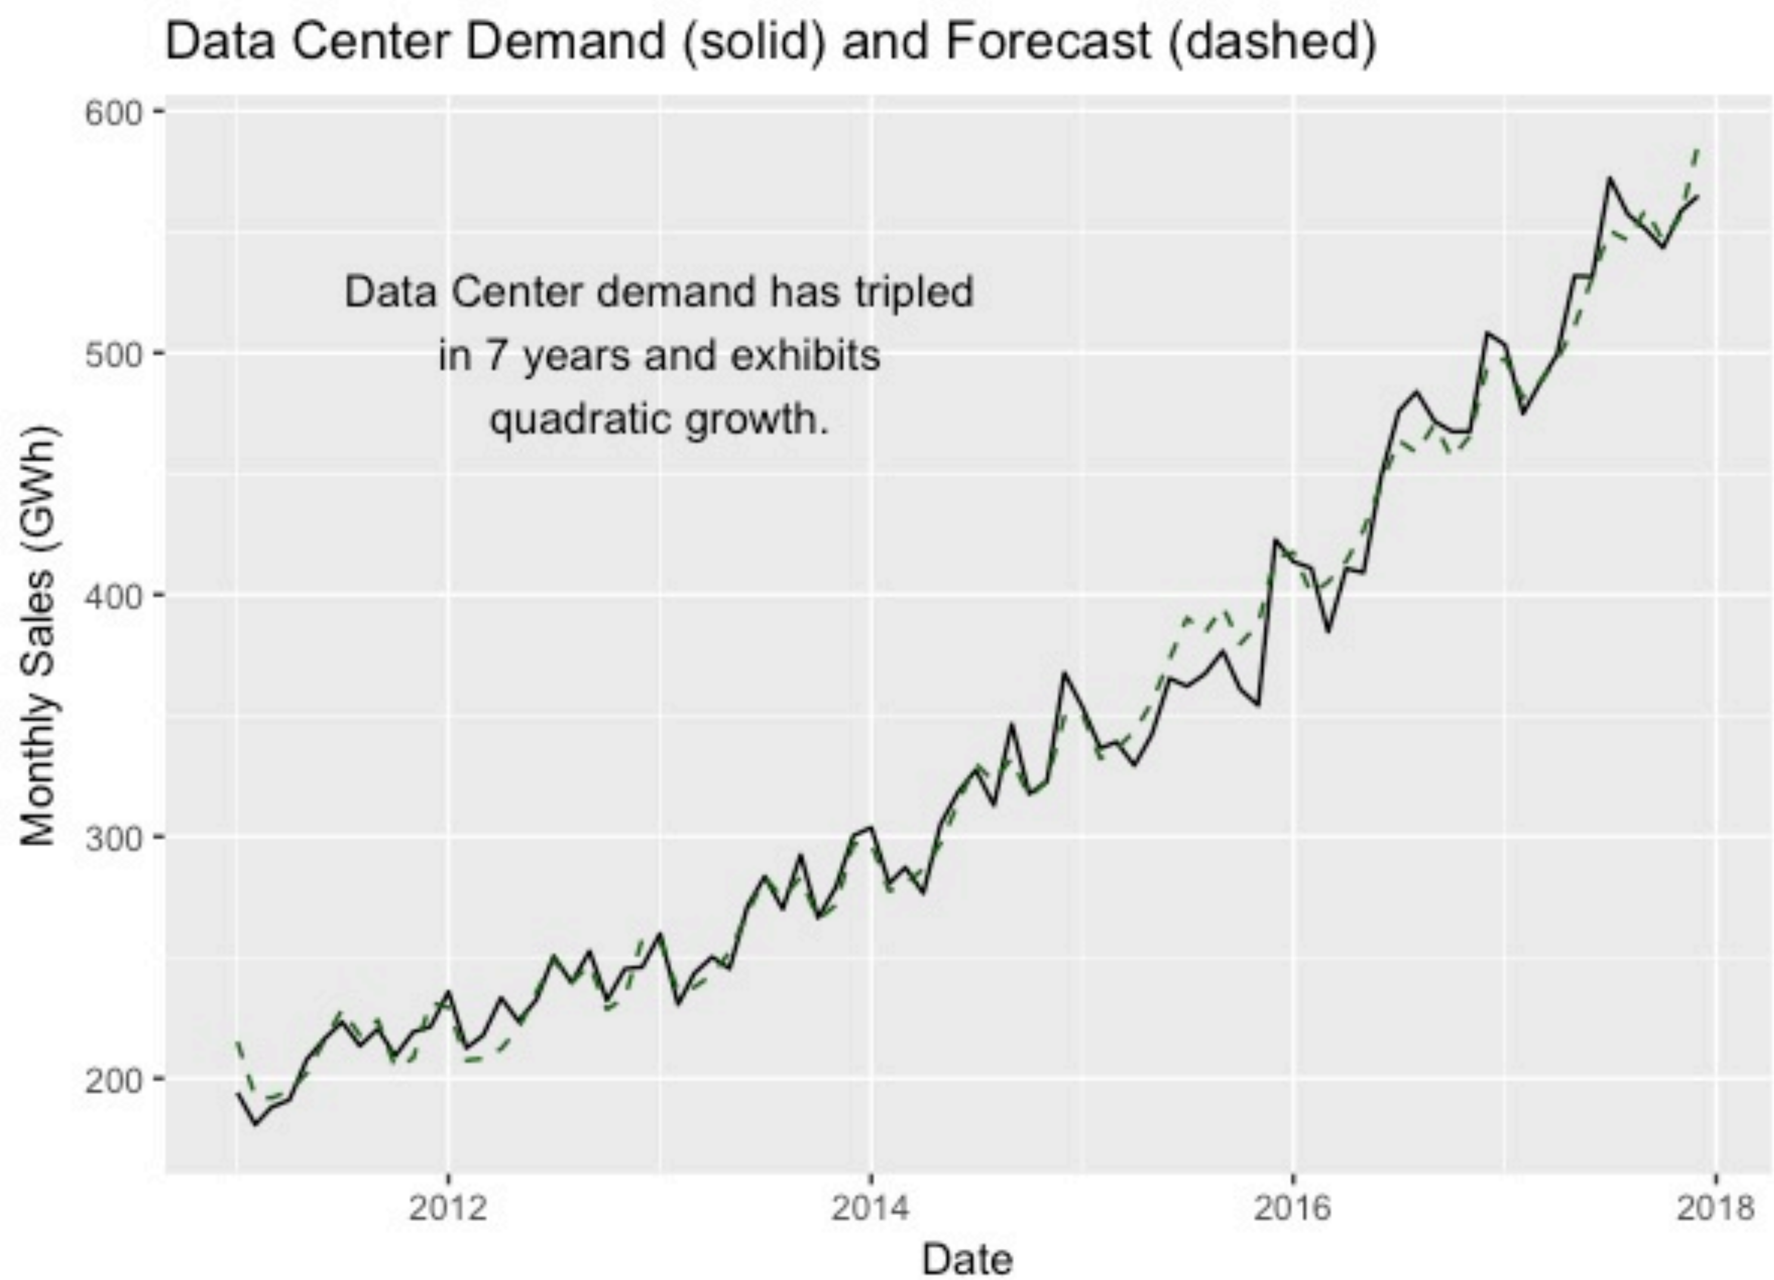
Data Center Demand (solid) and Forecast (dashed)
Data Center demand has tripled
in 7 years and exhibits
quadratic growth.
Monthly Sales (GWh)
600
500
400
300
200
2012
2014
2016
2018
Date

Quirky End of the Story

- Dominion uses a Bass Diffusion Model to forecast data centers
  - **This forces a logistic growth path on the future**
  - By this model, data center growth slows in near term
  - But this is almost certainly wrong
  - Wrong model, wrong forecast
- If Dominion persists in using this model, it will shift from persistent over-forecasting to persistent under-forecasting within a few years

Total Demand Forecasts Compared

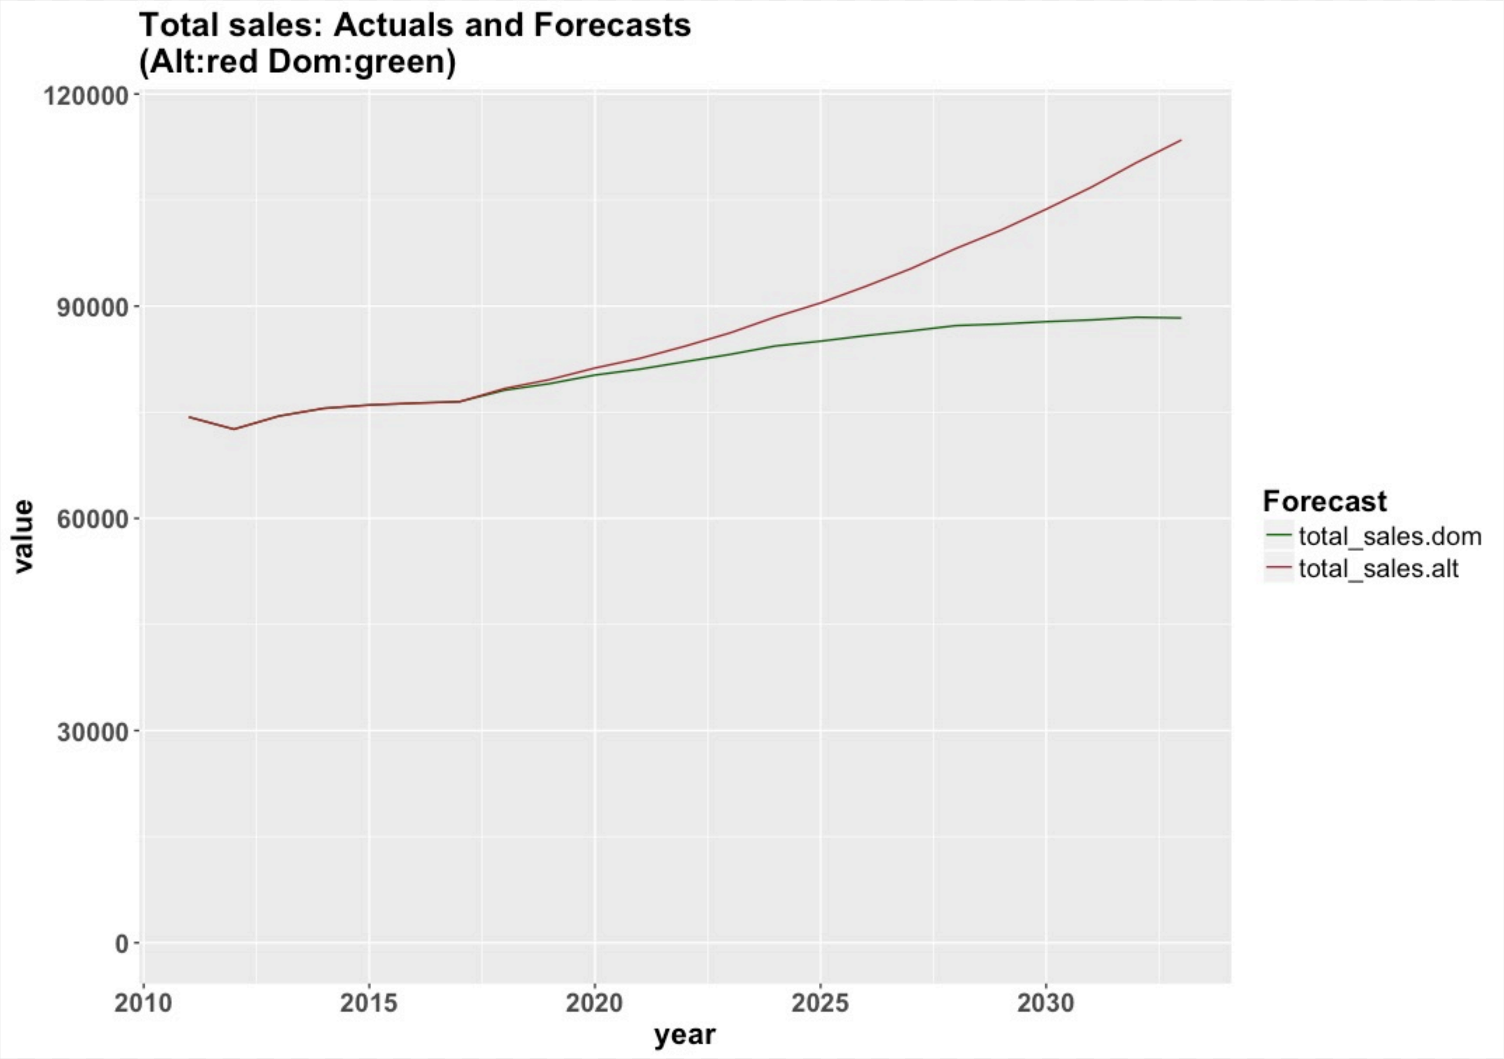

Key Lessons

- State Space Models are less restrictive and provide additional information to the modeler
  - Comparison to other specifications may be instructive
  - Explicit recognition of the presence of unmeasured state variables avoids some misspecification problems
- Disaggregating data into components with distinctive data generating processes is important
- Time-varying parameter models provide a simple improvement over linear regression

- Thank you.
- shobe@virginia.edu