Visualization of complex seasonal patterns in time series

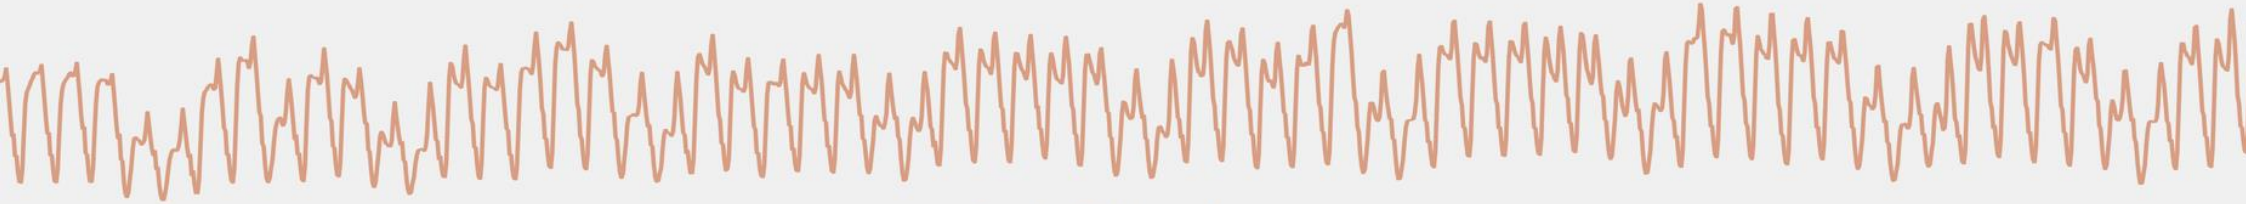

Rob J Hyndman

MONASH University

23 September 2022

Complex seasonalities

Victorian half-hourly electricity demand

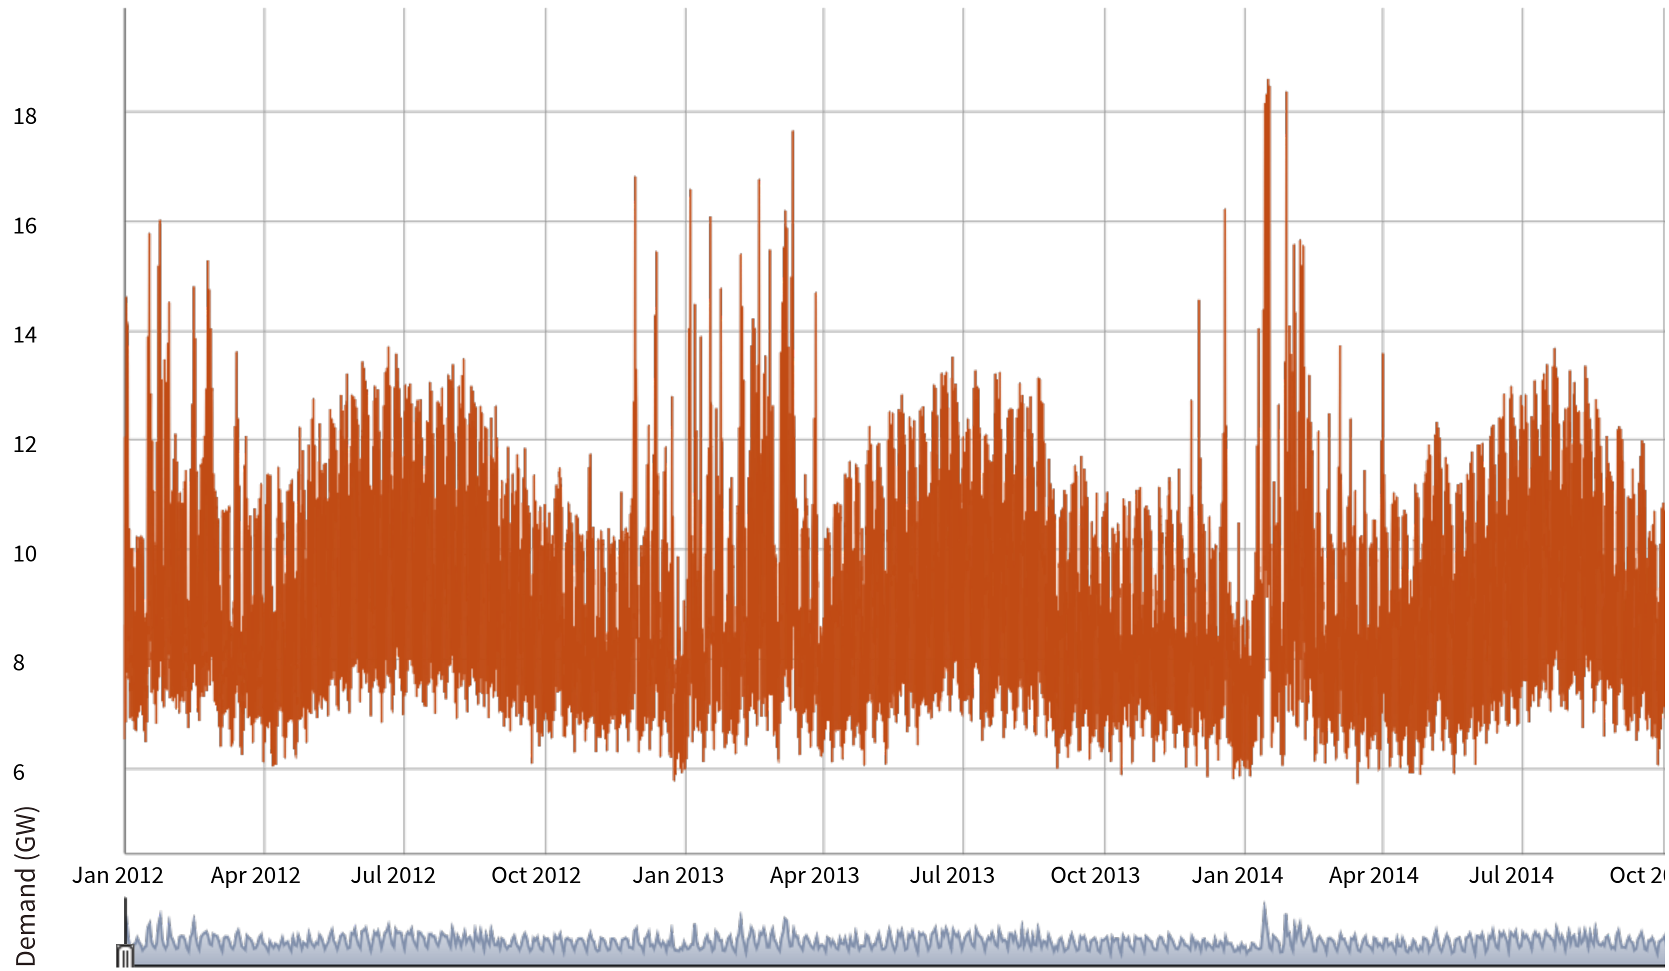

Complex seasonalities

Turkish daily electricity demand

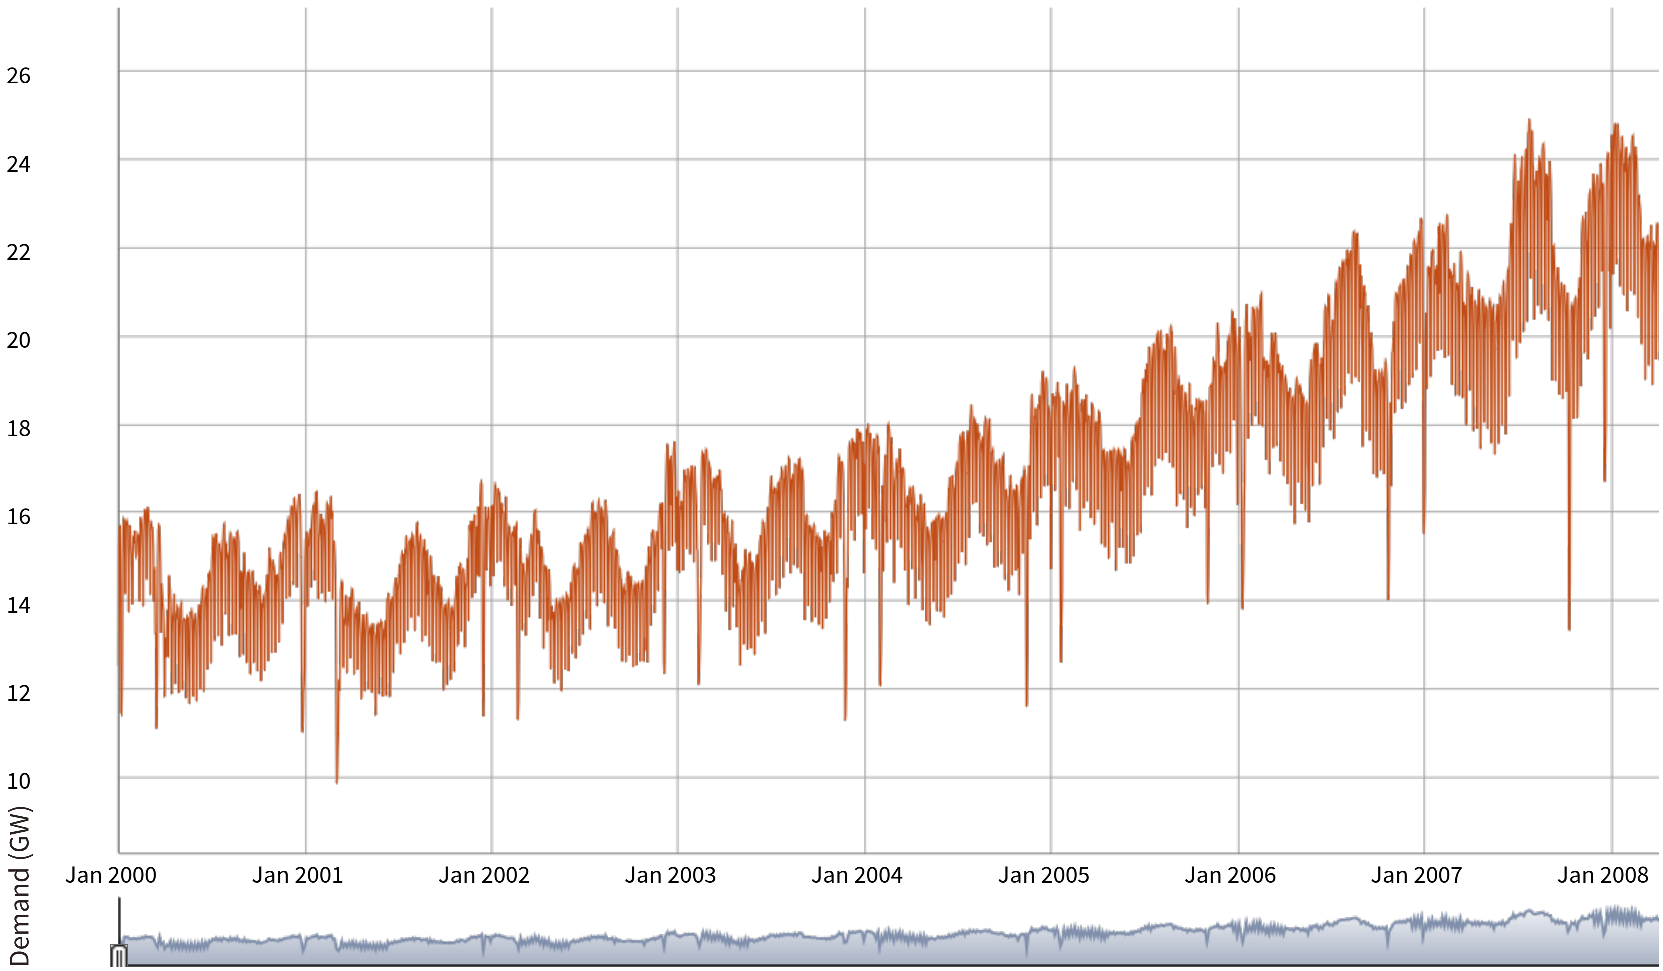

Complex seasonalities

US finished motor gasoline product supplied

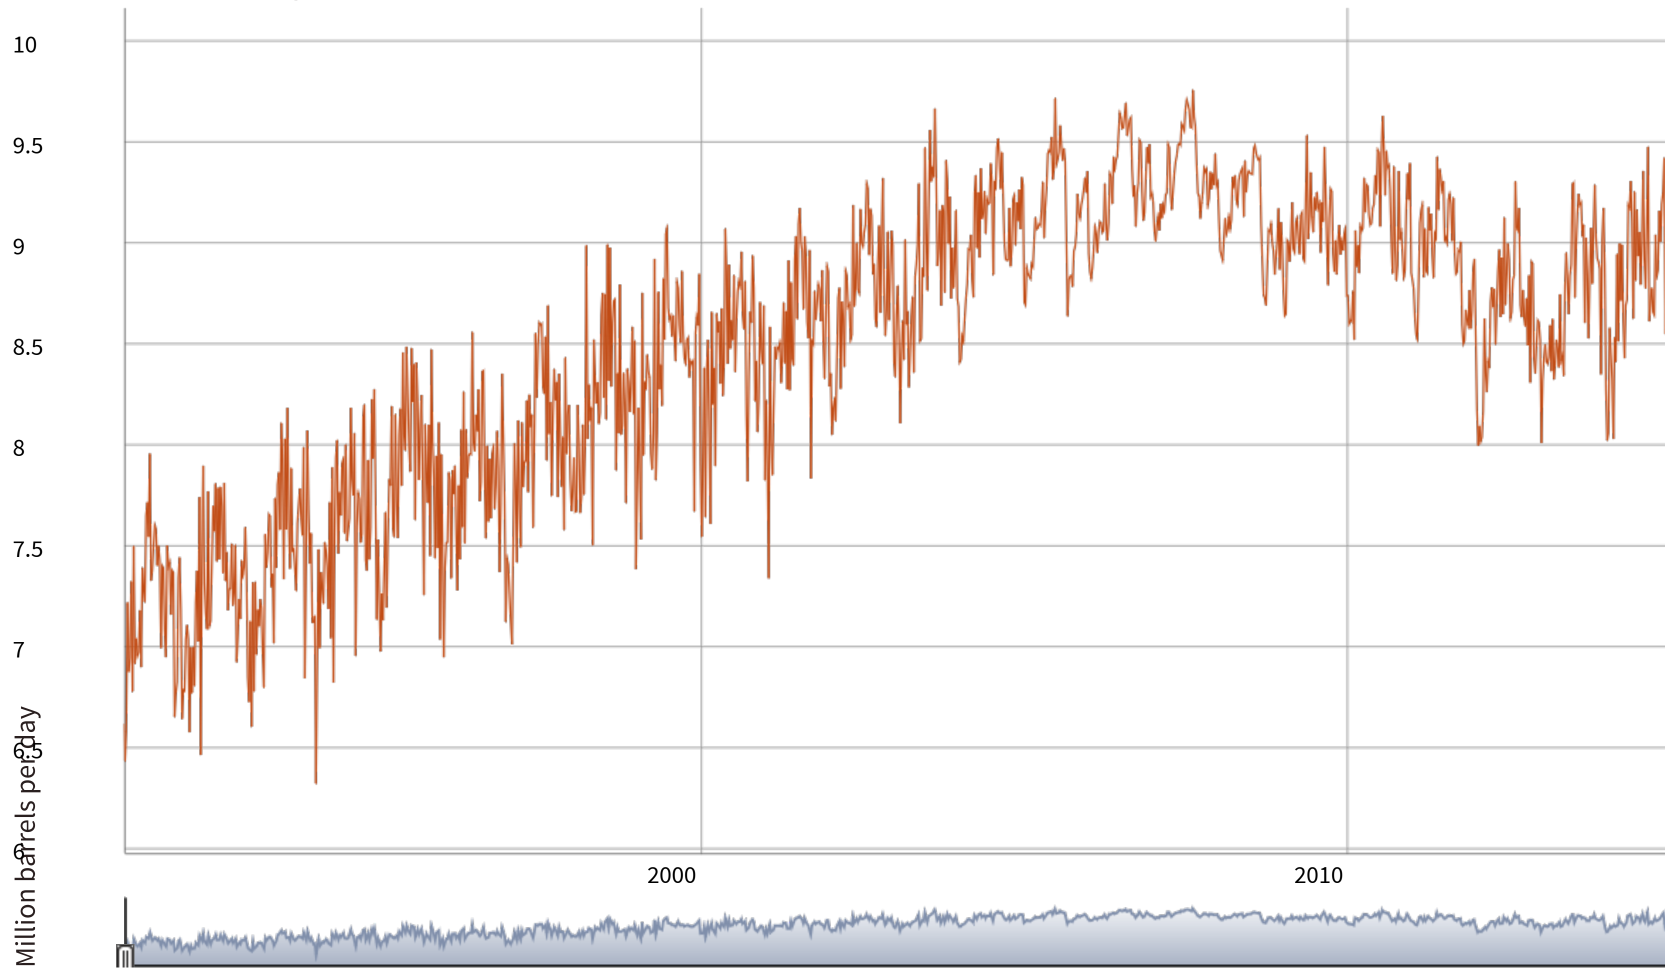

Complex seasonal topology

Example: hourly data

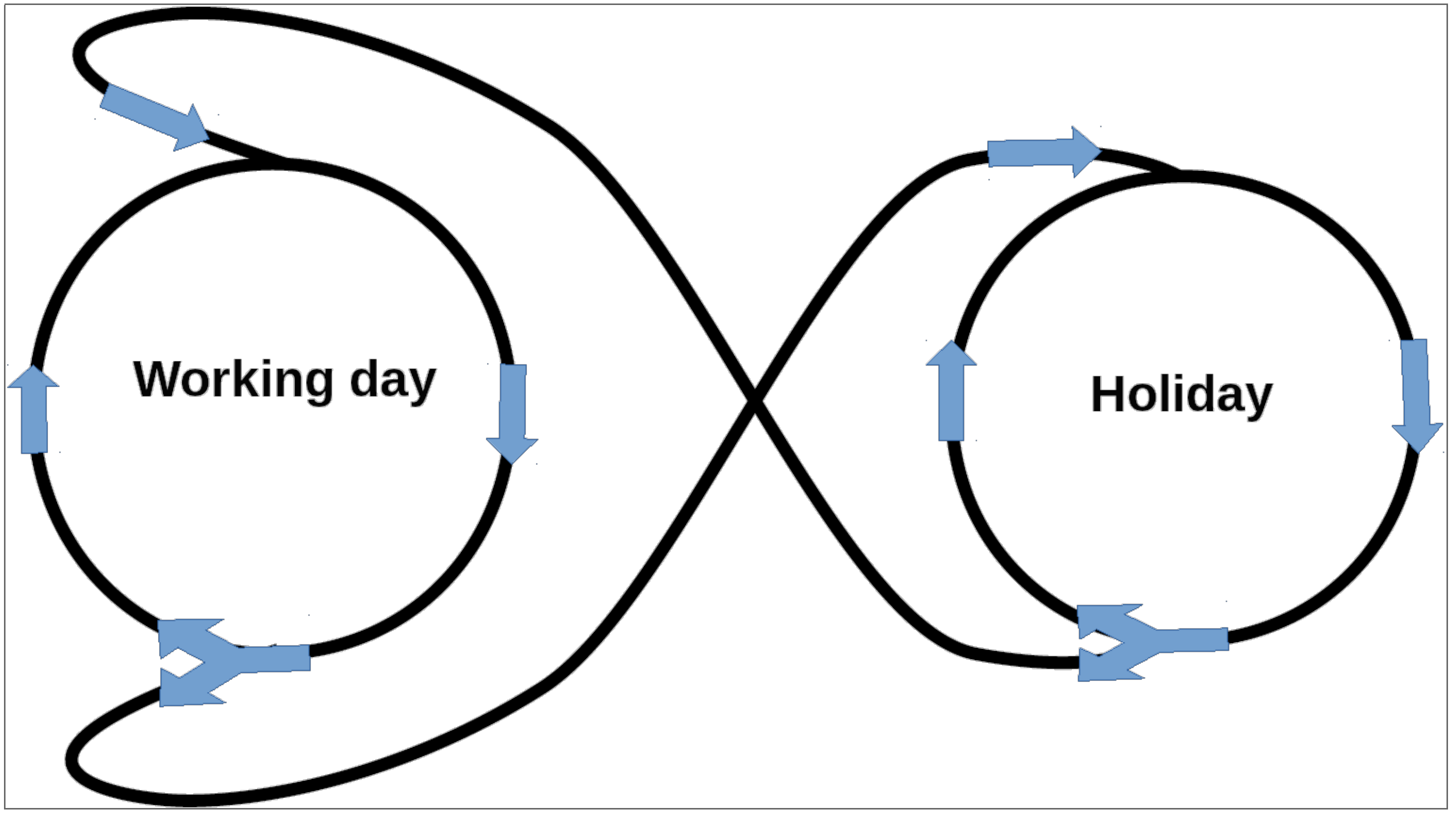

Complex seasonalities

- Multiple seasonal periods, not necessarily nested
- Non-integer seasonality
- Irregular seasonal topography
- Seasonality that depends on covariates
- Complex seasonal topology

- How to effectively visualise the underlying seasonalities?
- How to decompose such time series into trend and multiple season components?

Visualizing complex seasonalities

- Omitted granularities `hour`, `fortnight`, `quarter`, `semester`

Granularity plots

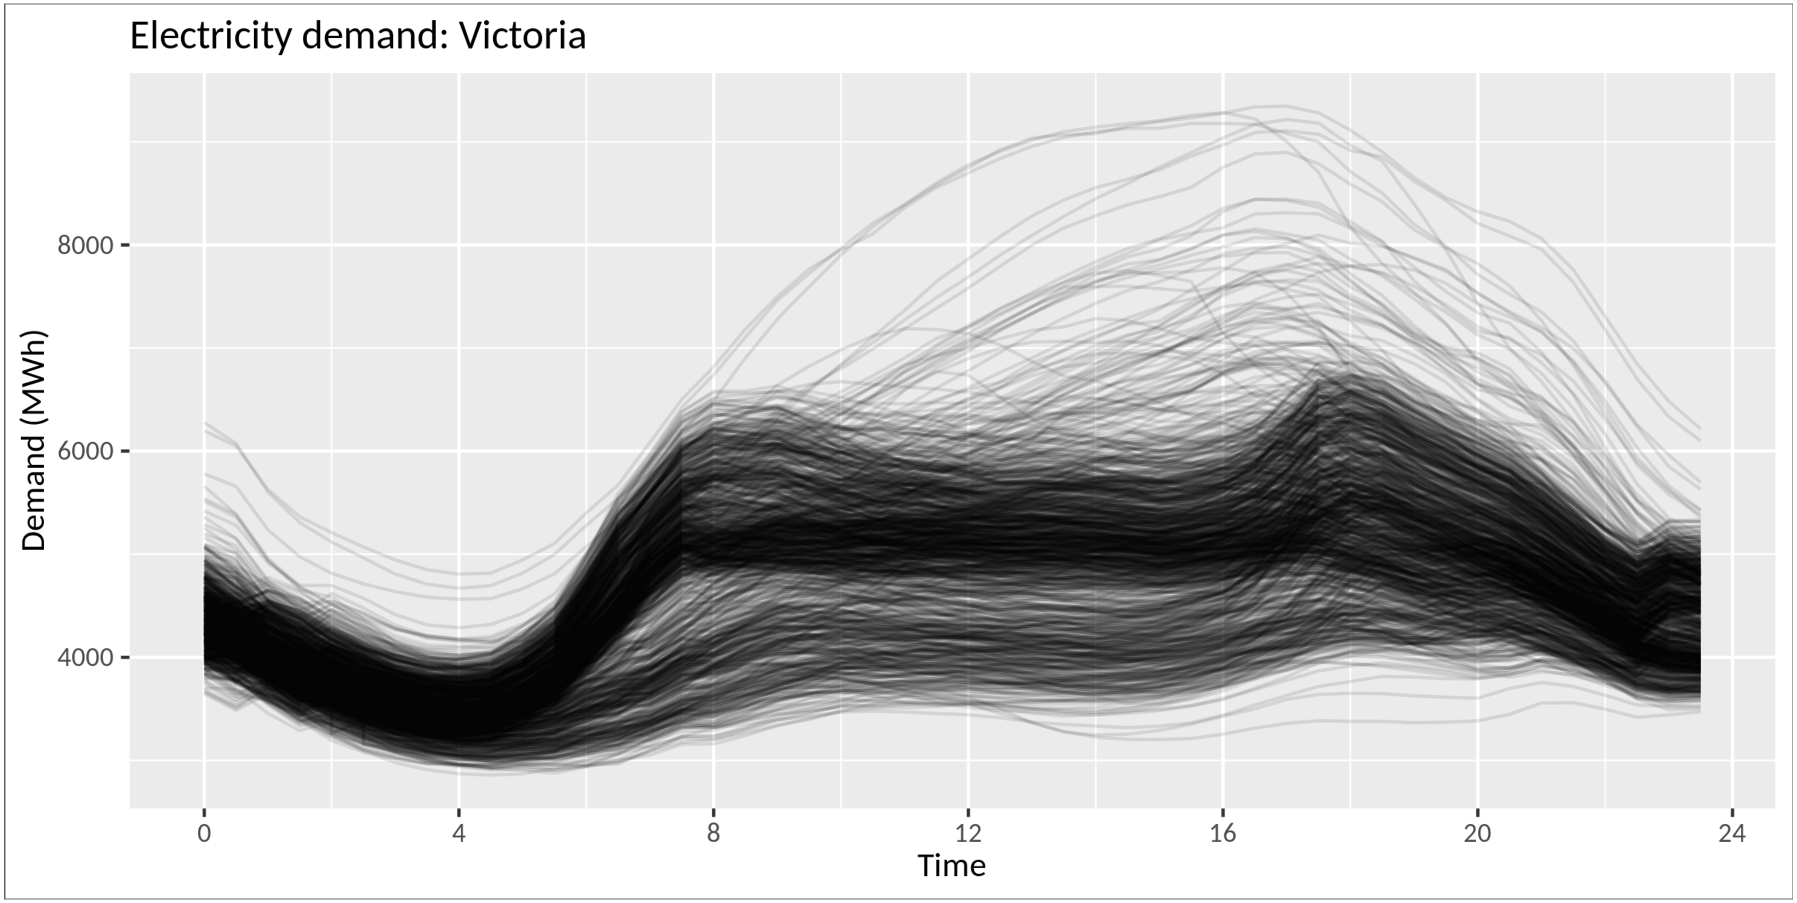

Granularity plots

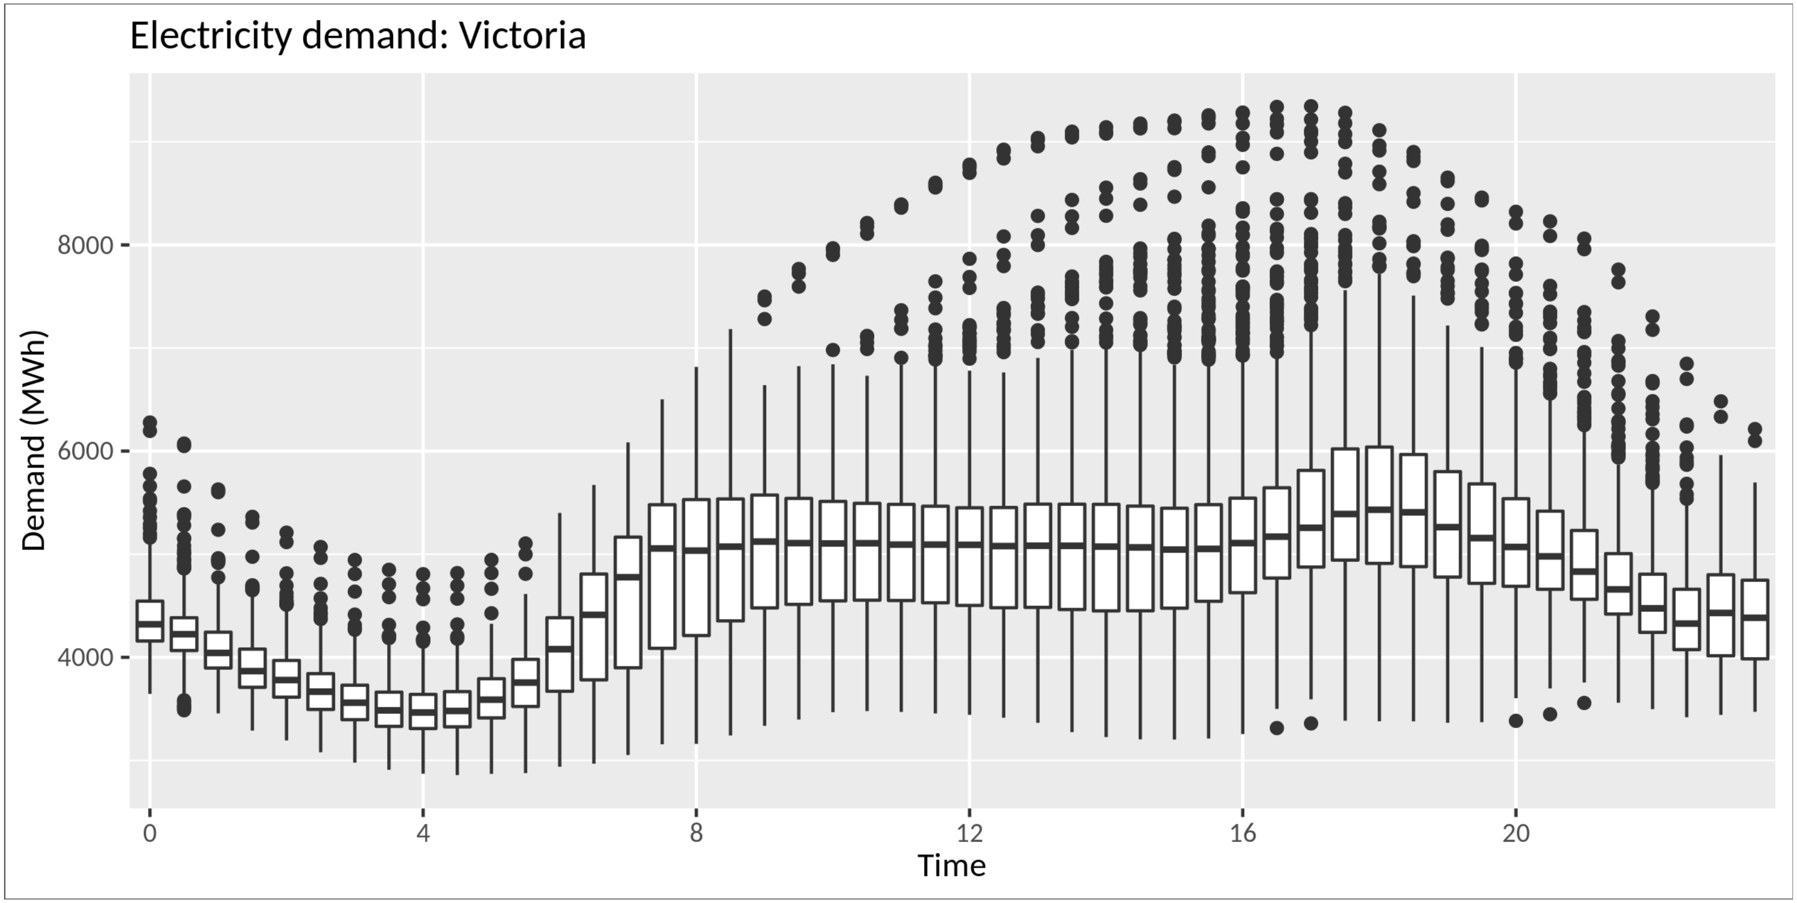

Granularity plots

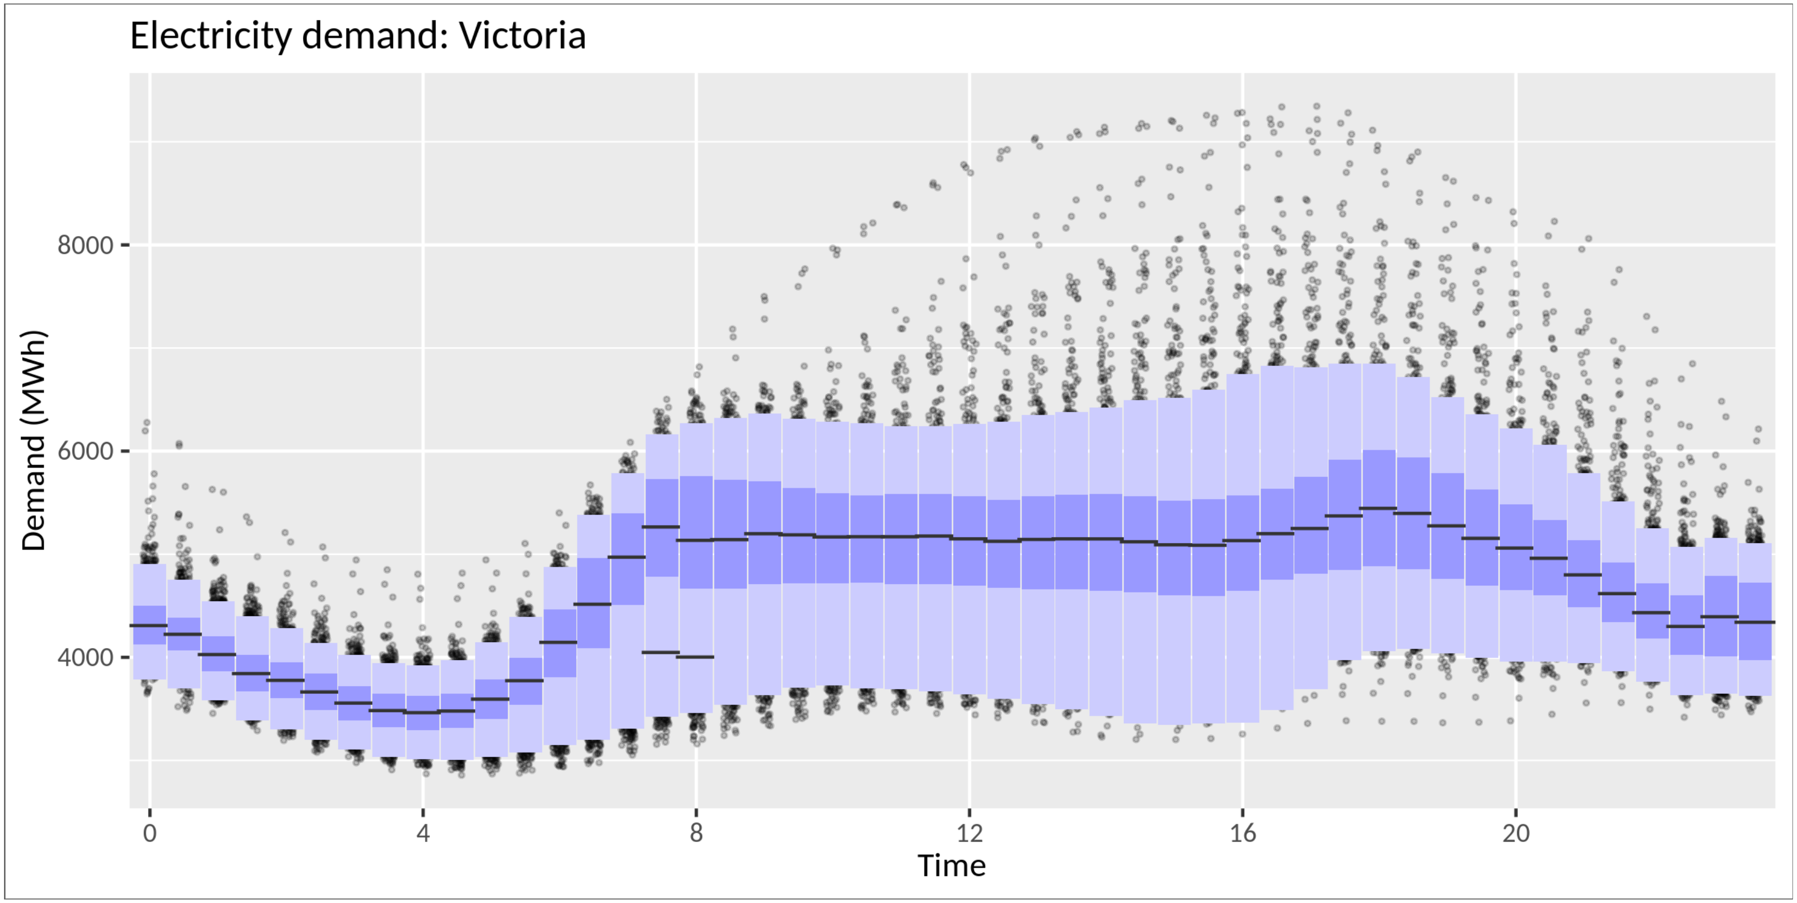

Granularity plots

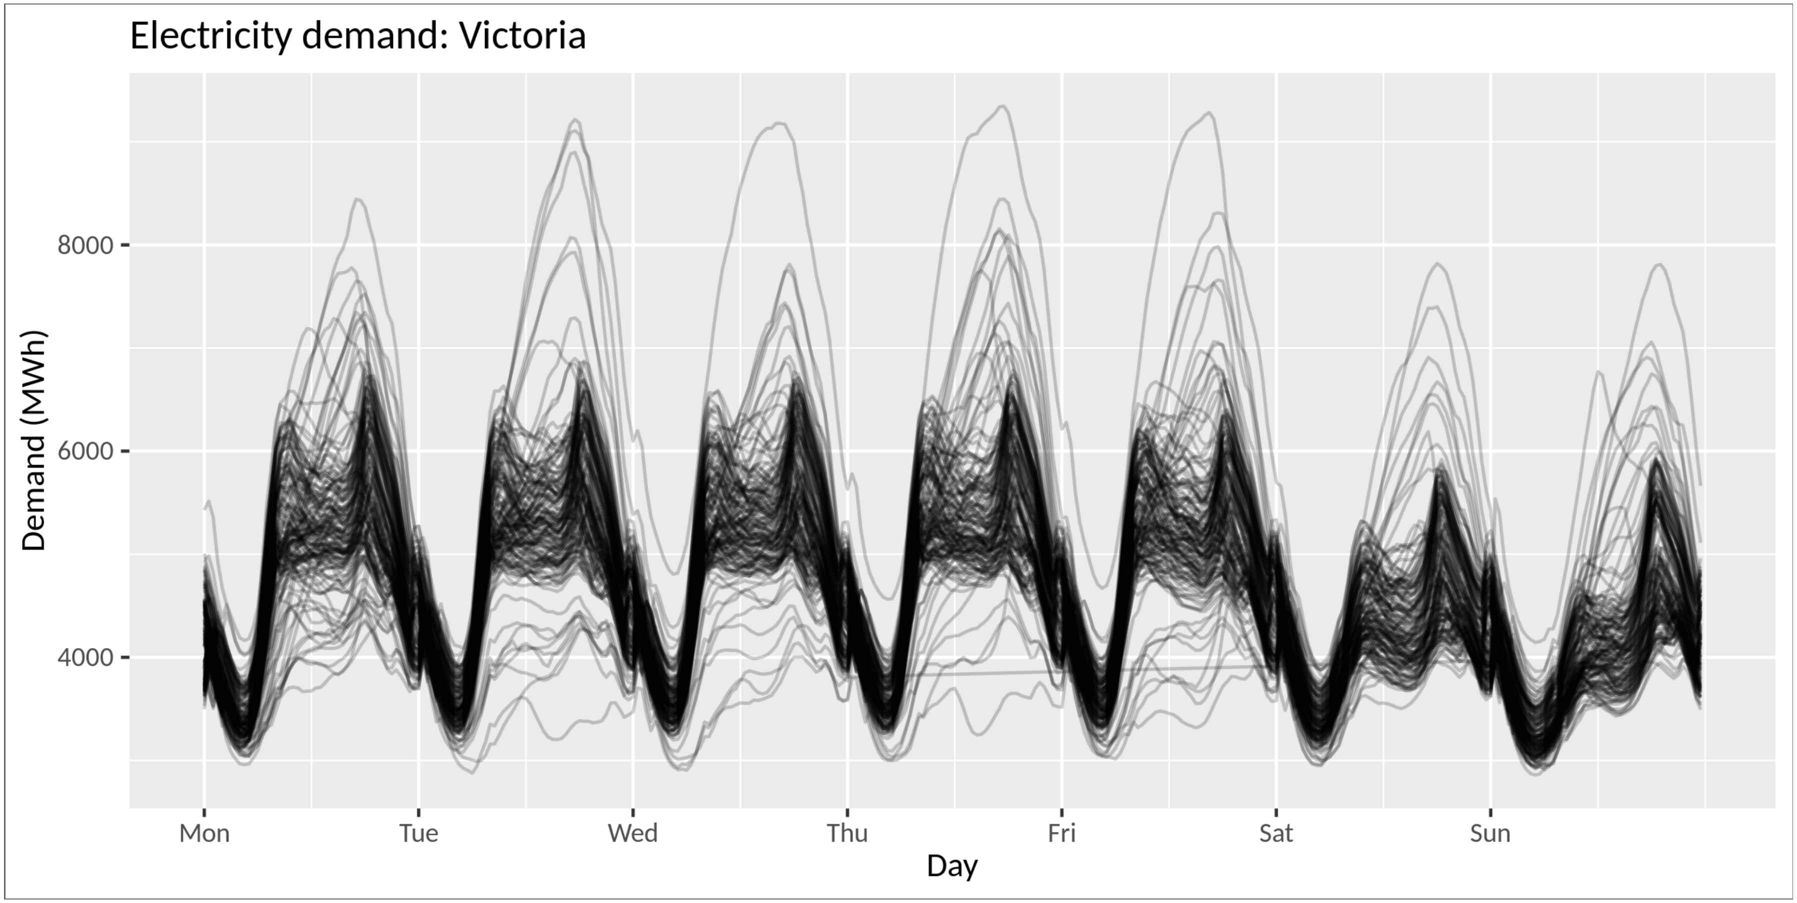

Granularity plots

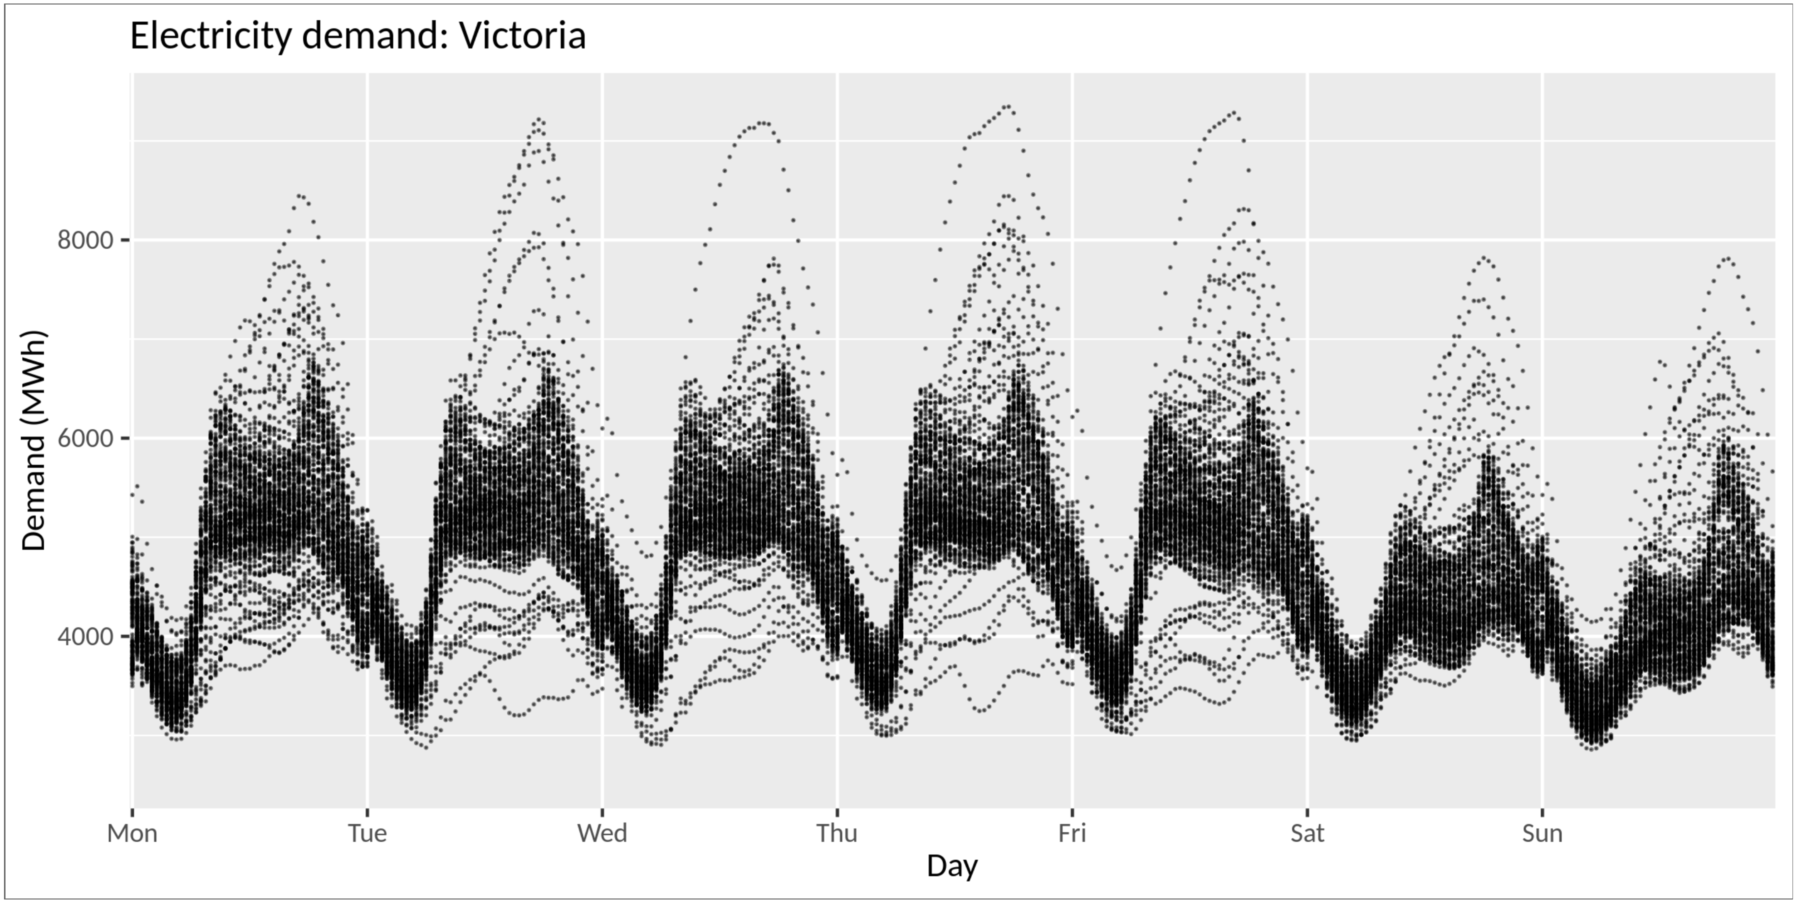

Granularity plots

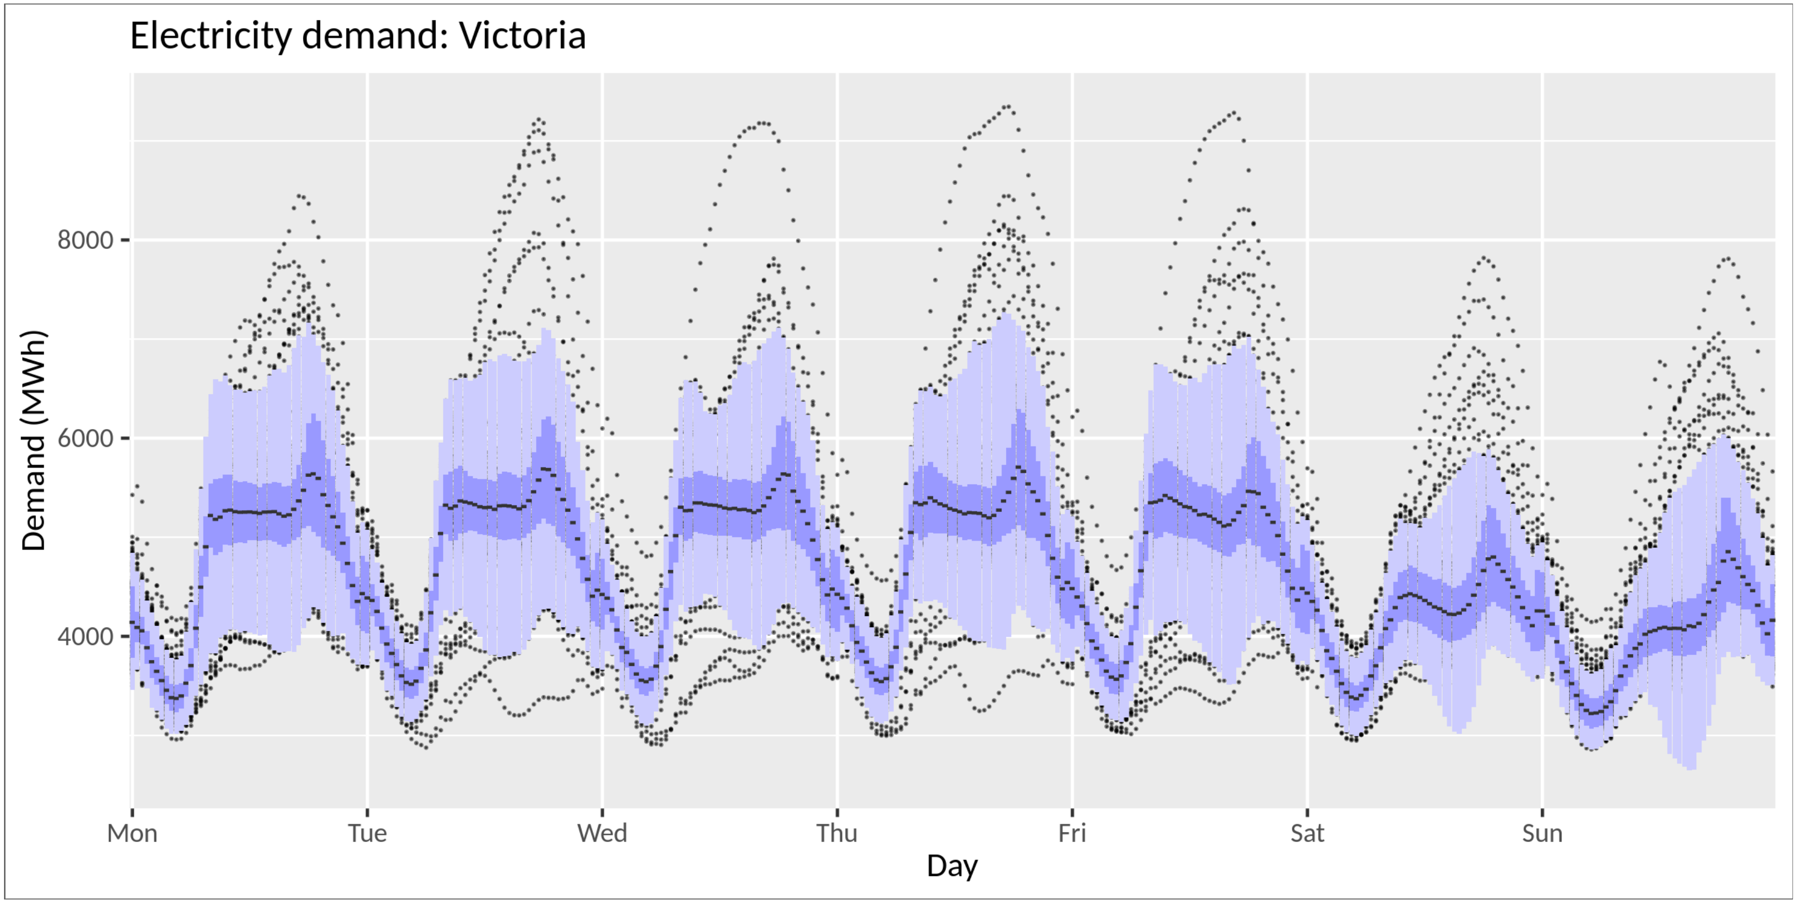

Granularity plots

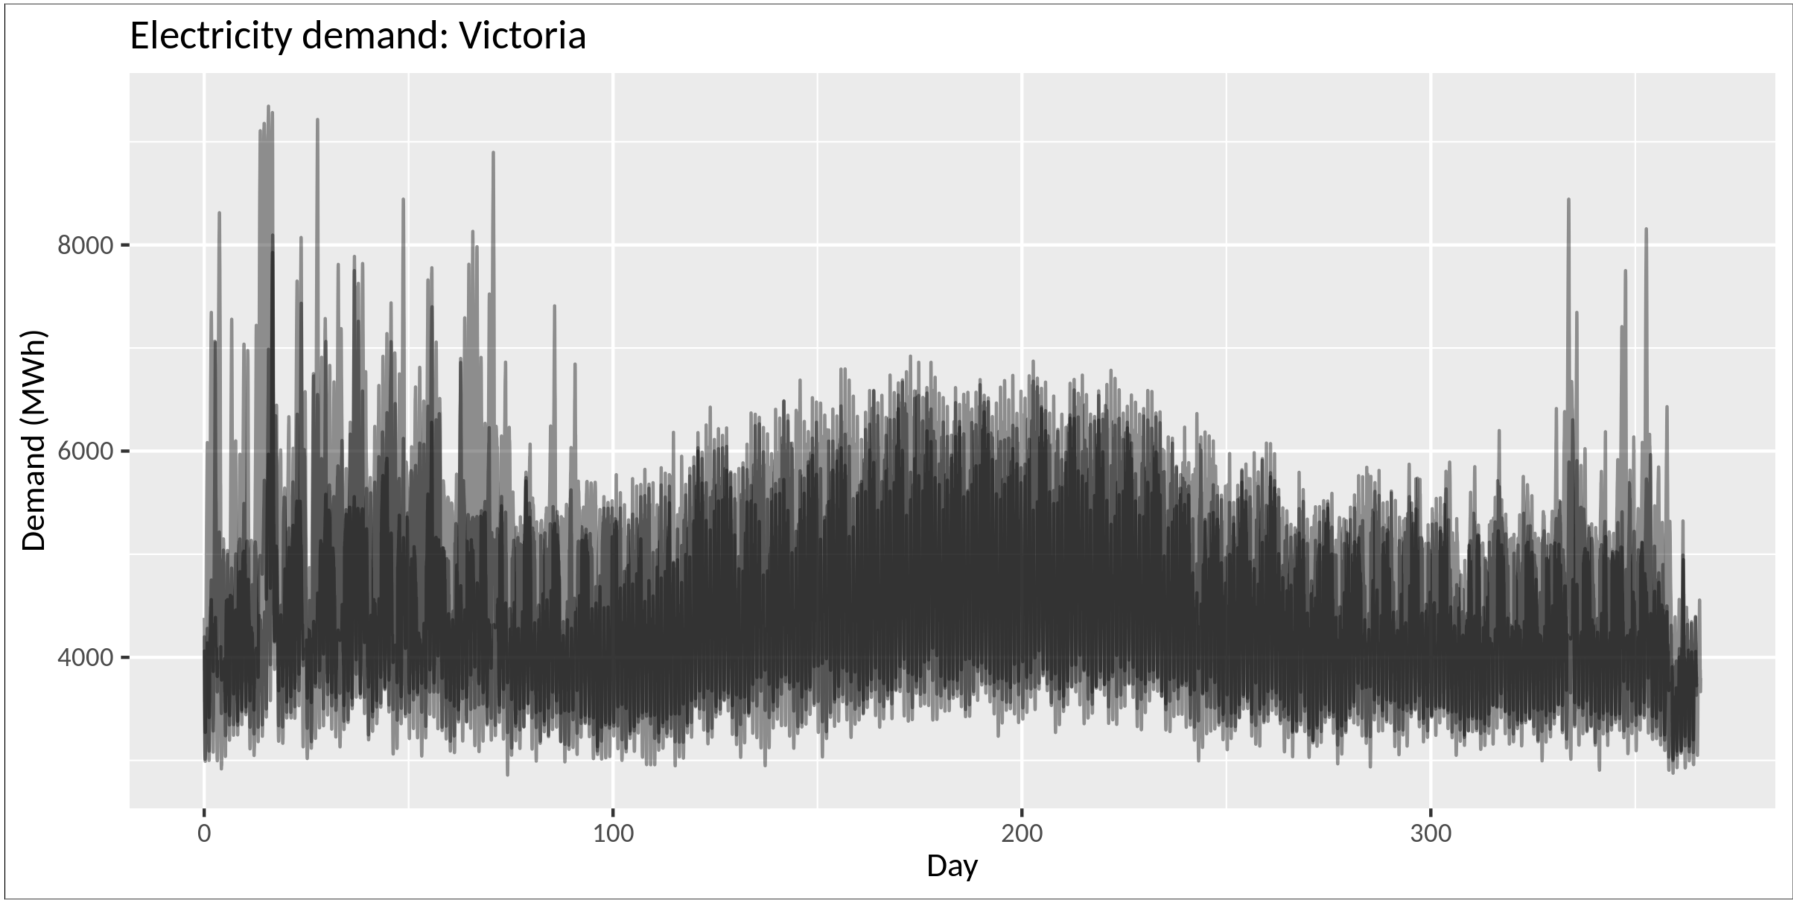

Granularity plots

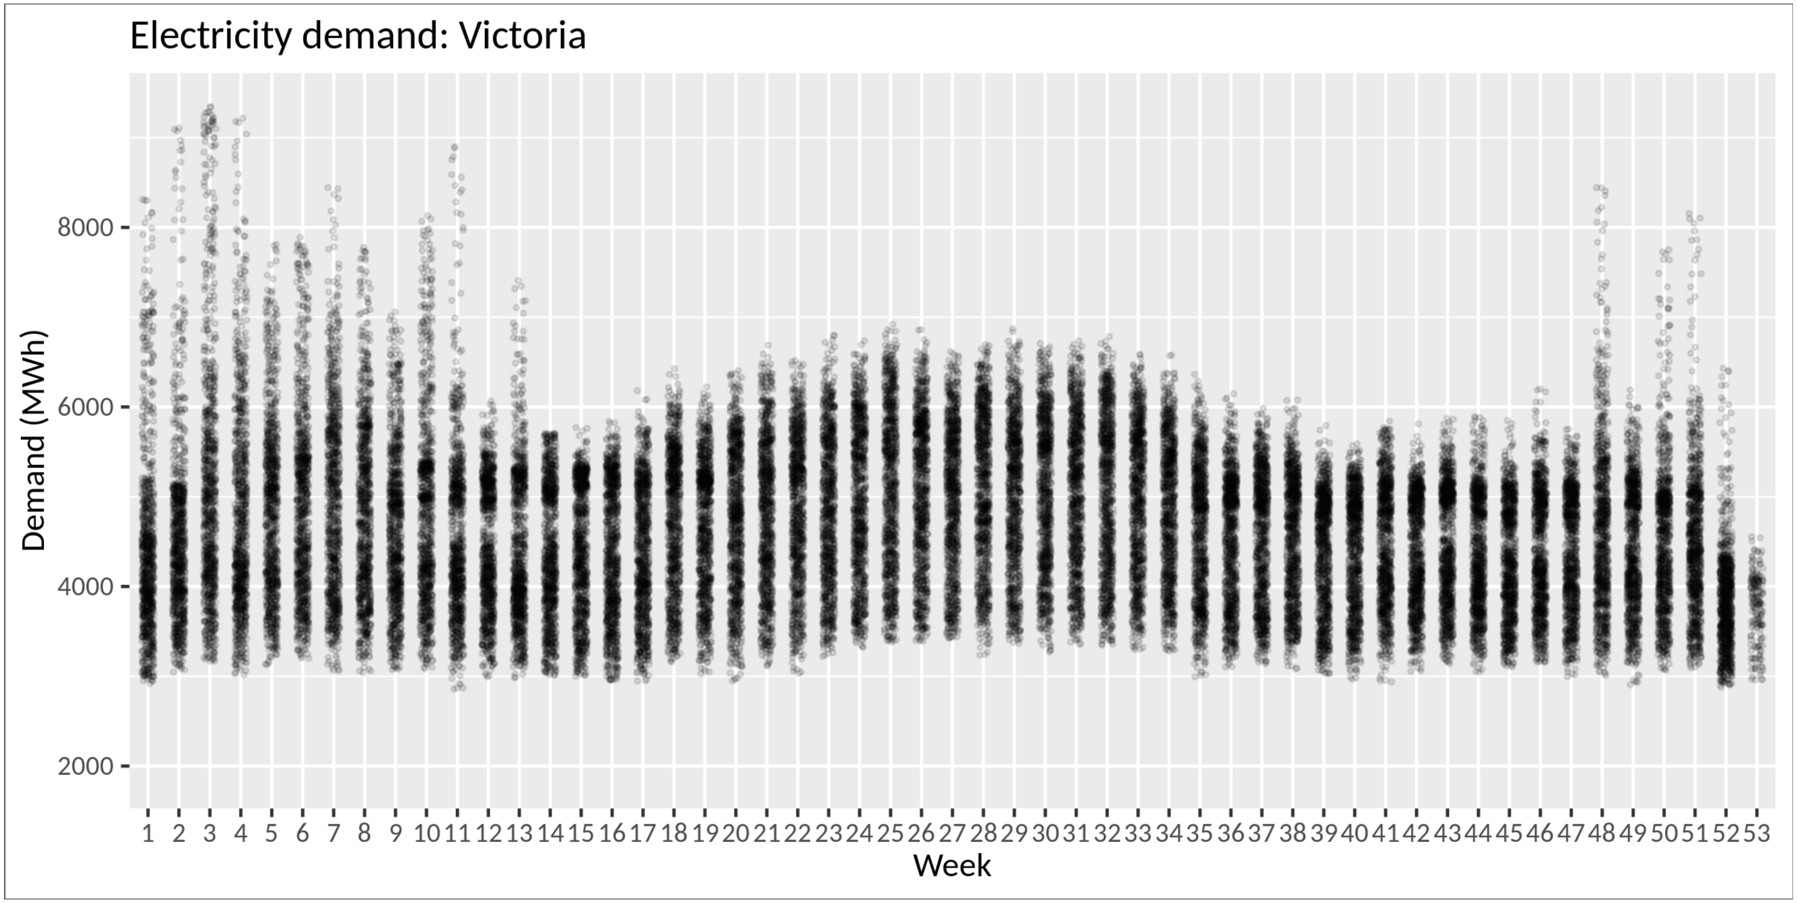

Granularity plots

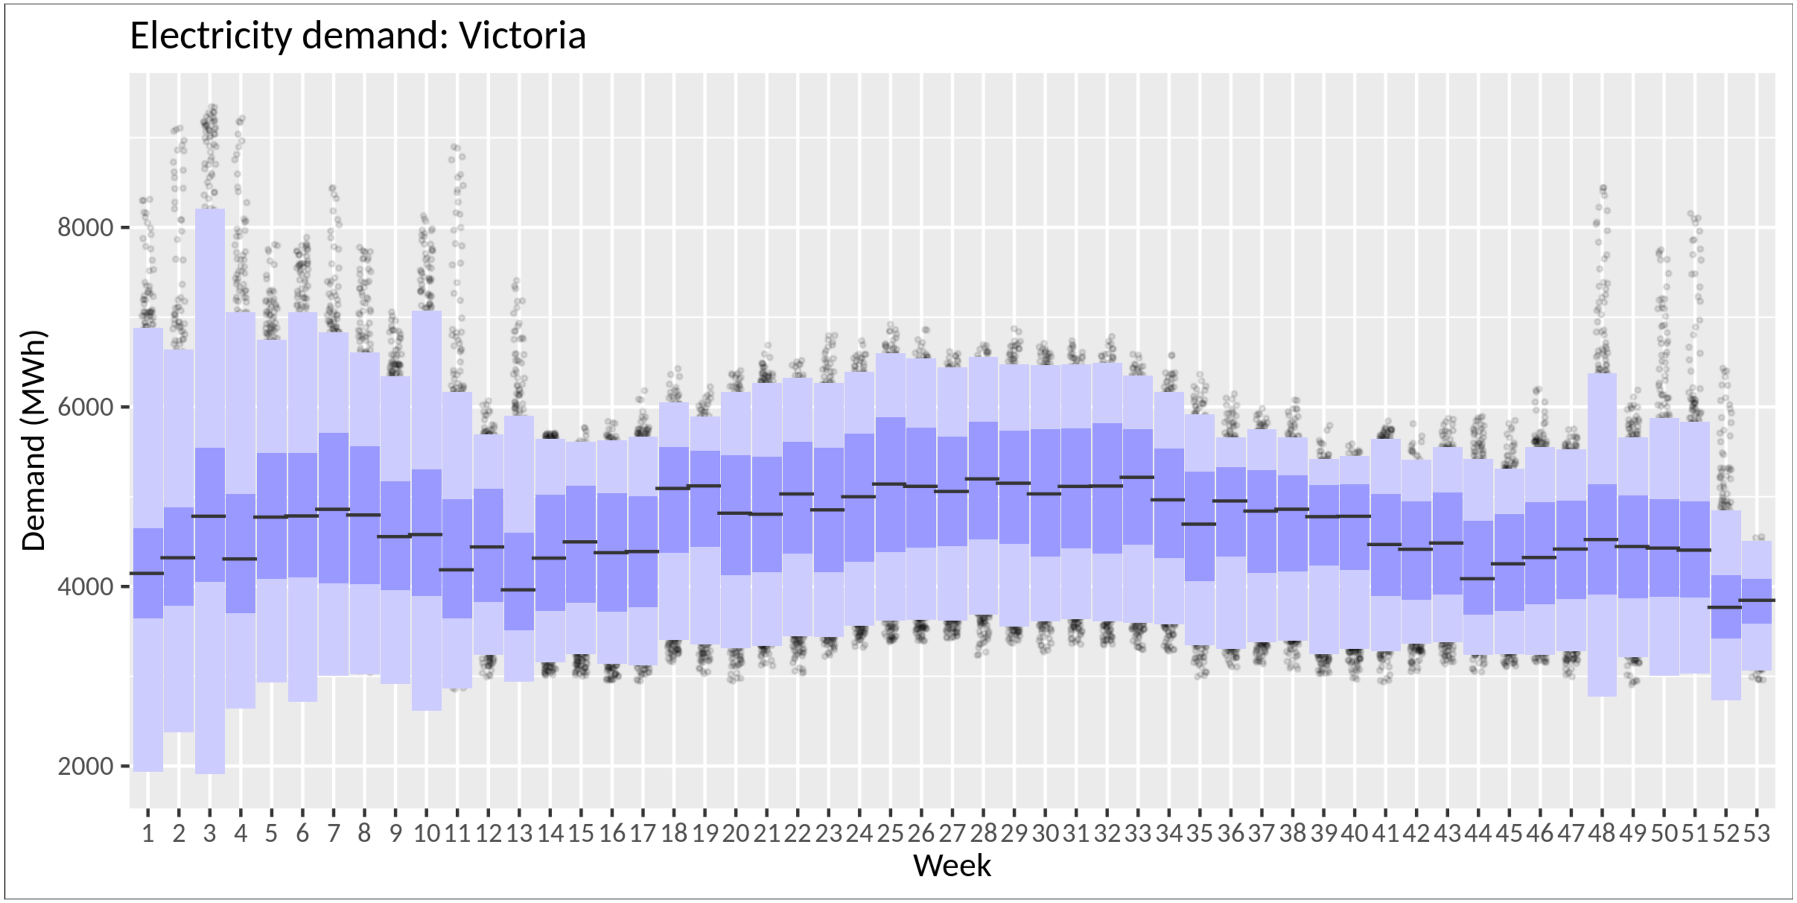

Faceted granularity plots

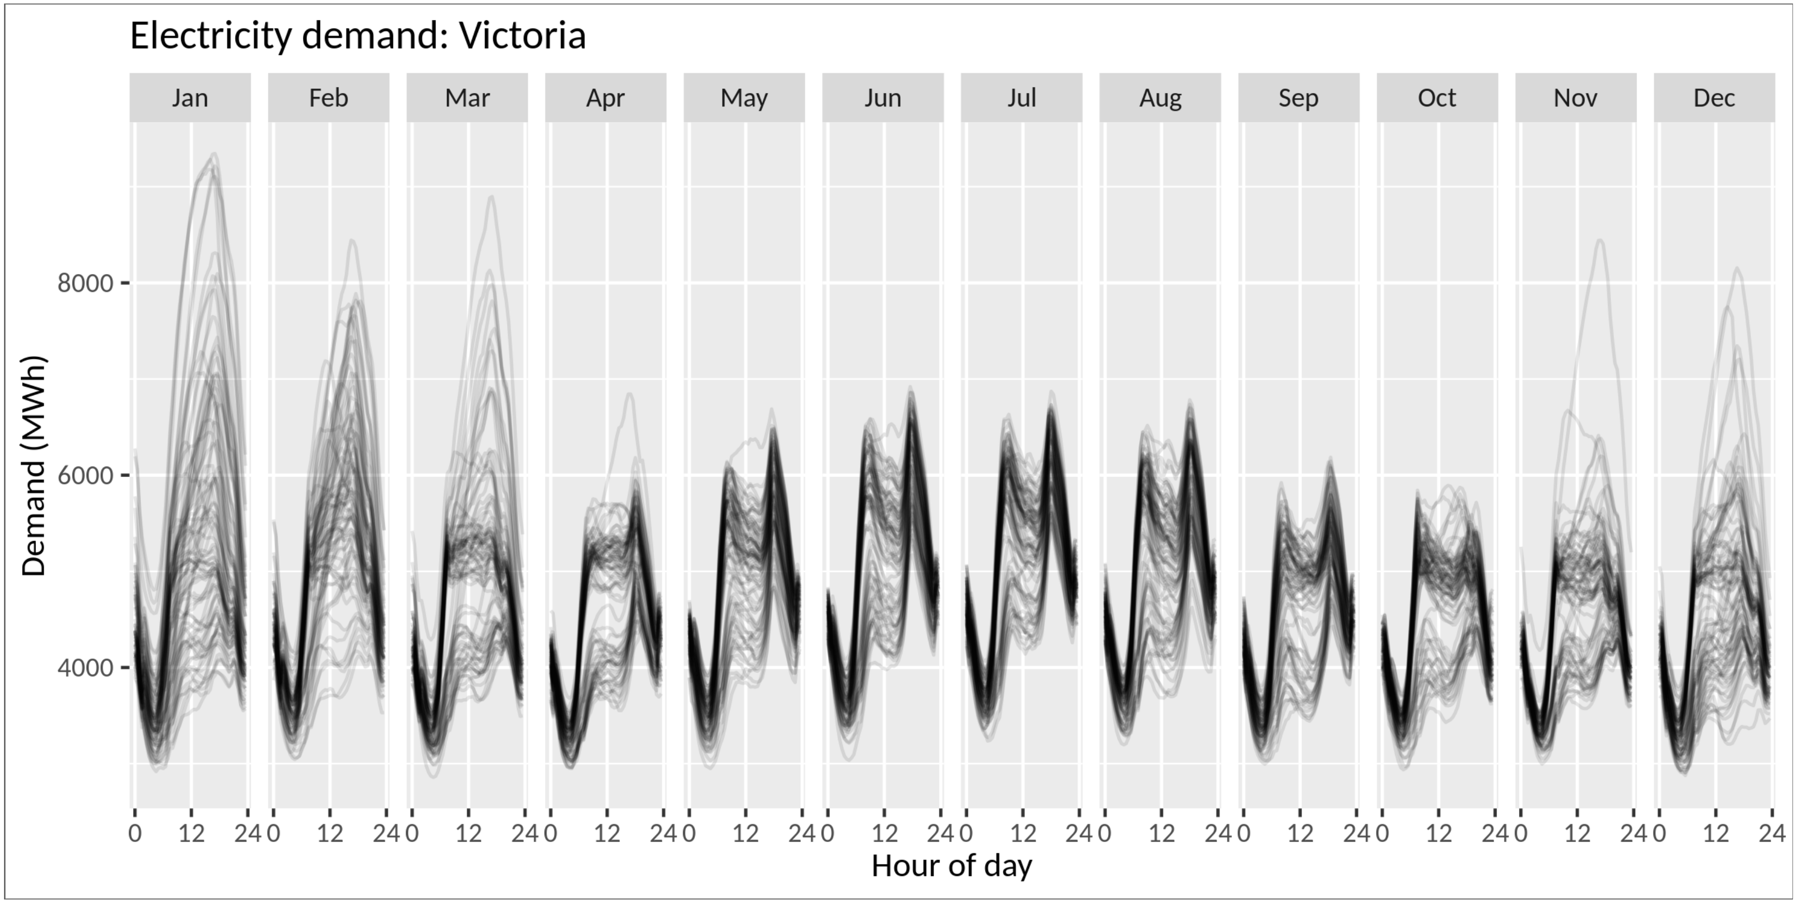

Faceted granularity plots

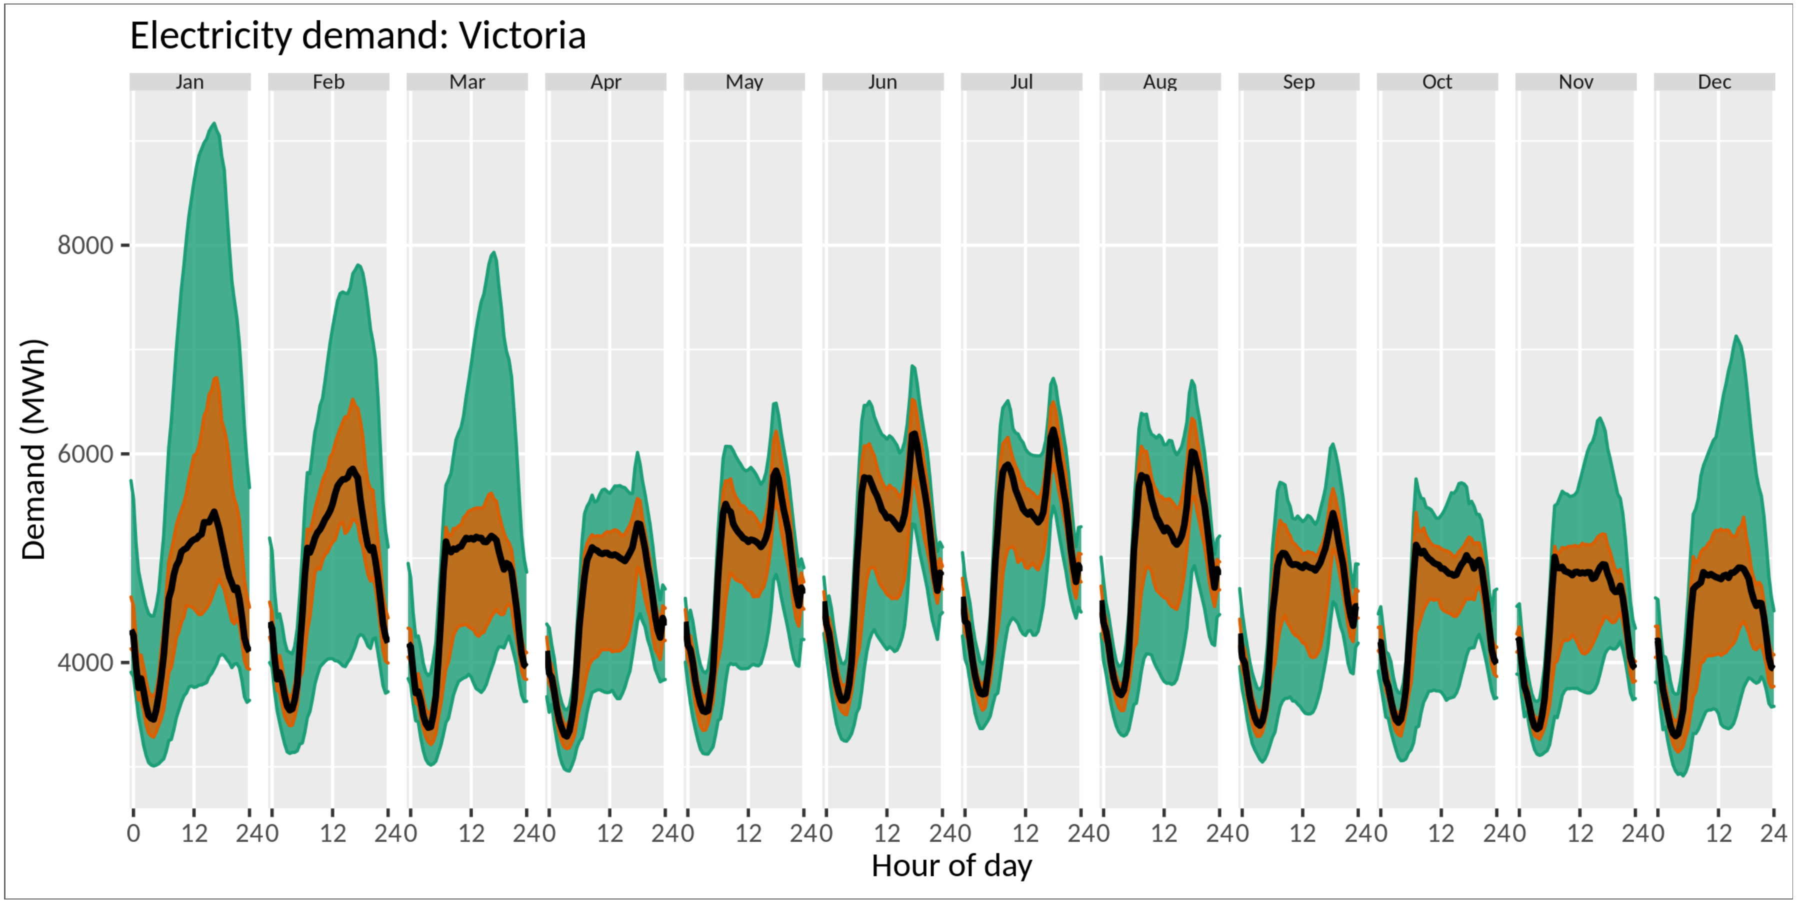

Faceted granularity plots

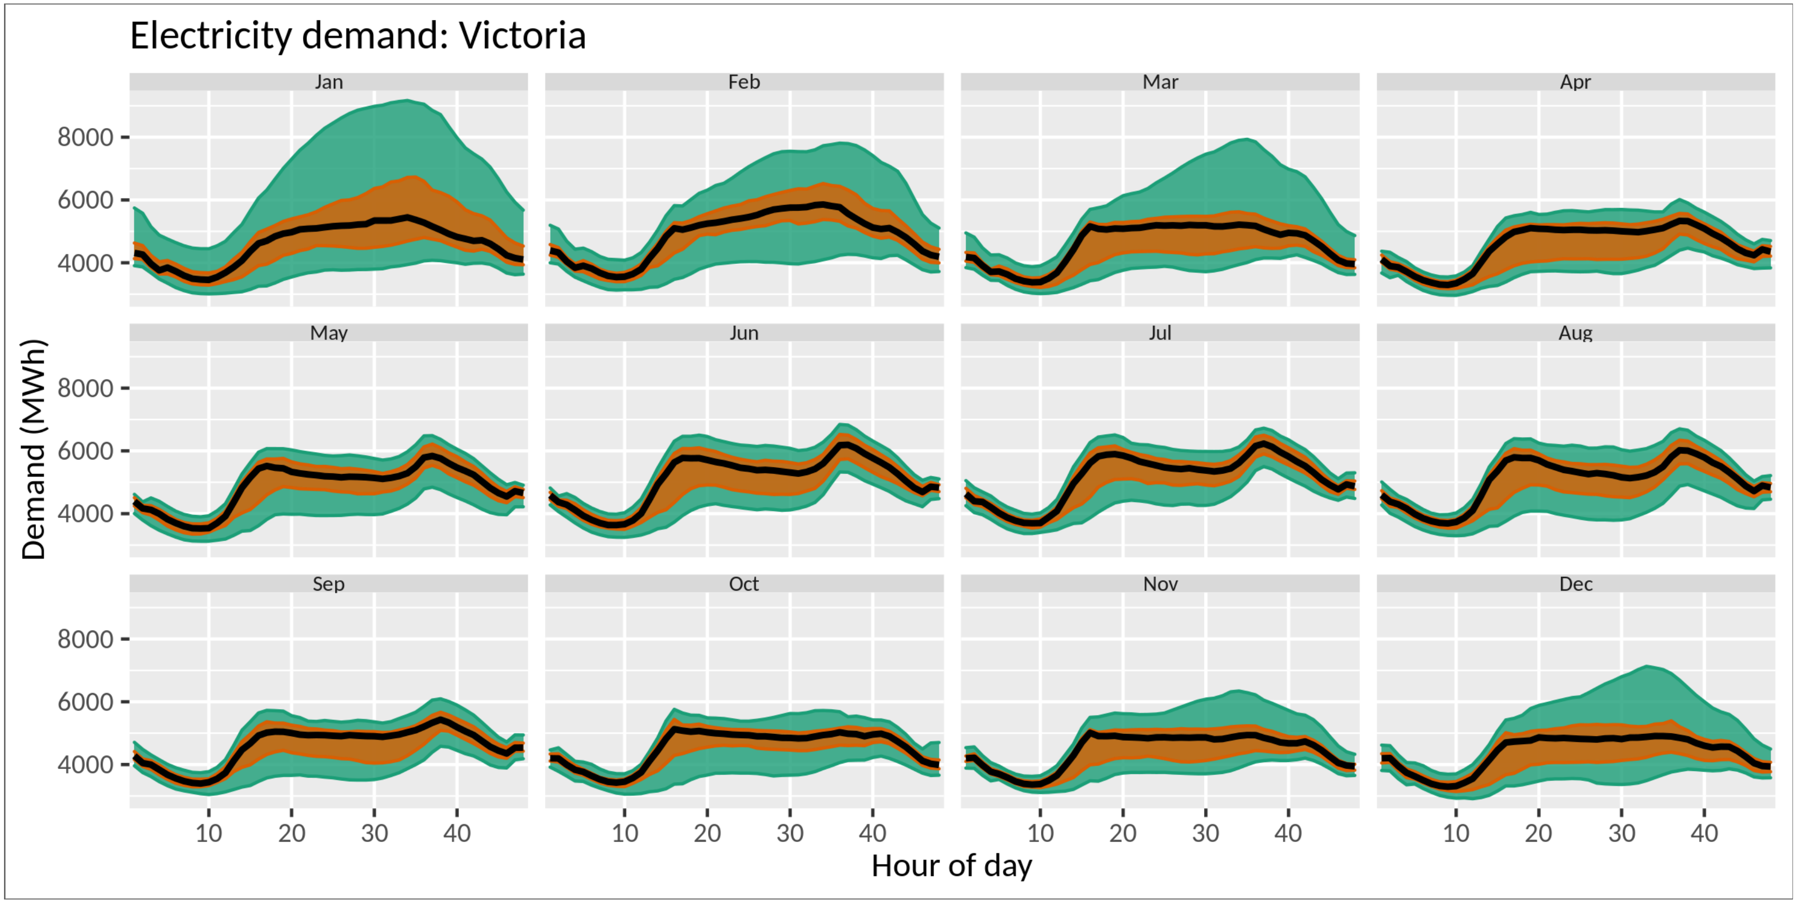

Faceted granularity plots

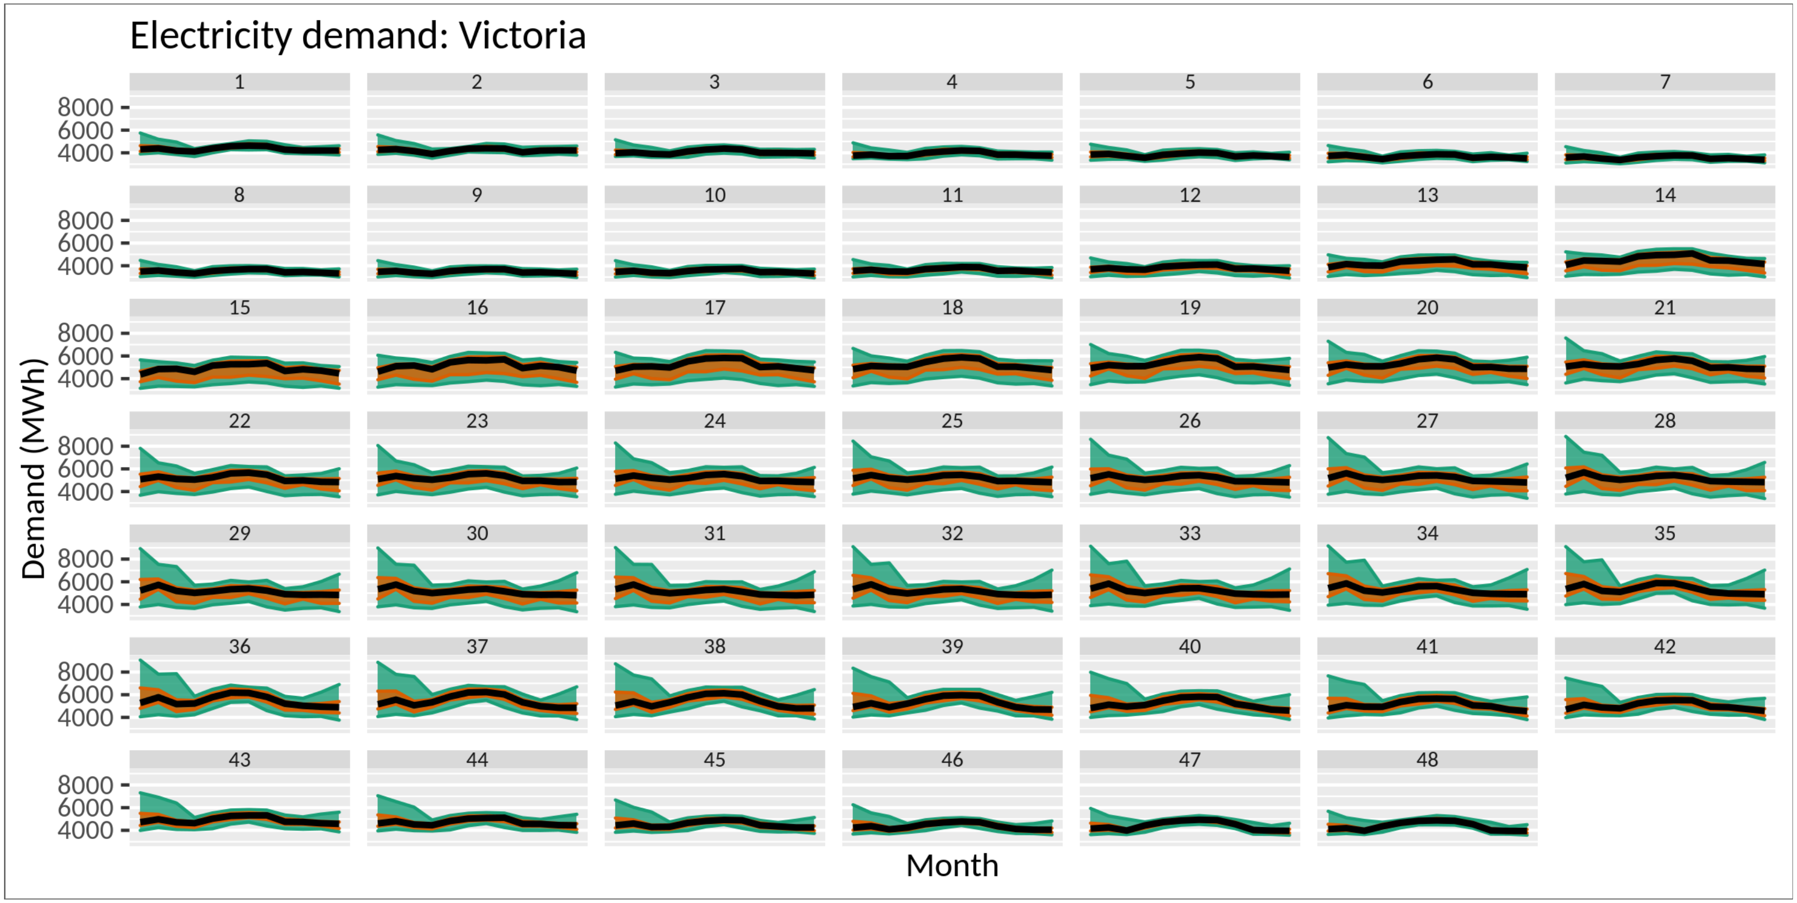

Granularities

Nested linear granularities

`hhour`, `hour`, `day`, `week`, `fortnight`, `quarter`, `semester`, `year`

Available cyclic granularities for half-hourly data

`hhour/hour`, `hhour/day`, `hhour/week`, `hhour/fortnight`, `hhour/month`, `hhour/quarter`, `hhour/semester`, `hhour/year`, `hour/day`, `hour/week`, `hour/fortnight`, `hour/month`, `hour/quarter`, `hour/semester`, `hour/year`, `day/week`, `day/fortnight`, `day/month`, `day/quarter`, `day/semester`, `day/year`, `week/fortnight`, `week/month`, `week/quarter`, `week/semester`, `week/year`, `fortnight/month`, `fortnight/quarter`, `fortnight/semester`, `fortnight/year`, `month/quarter`, `month/semester`, `month/year`, `quarter/semester`, `quarter/year`, `semester/year`

Plot options

- raw data or distributional summary on y-axis
- granularity on x-axis
- optional granularity as facet

What is an interesting plot?

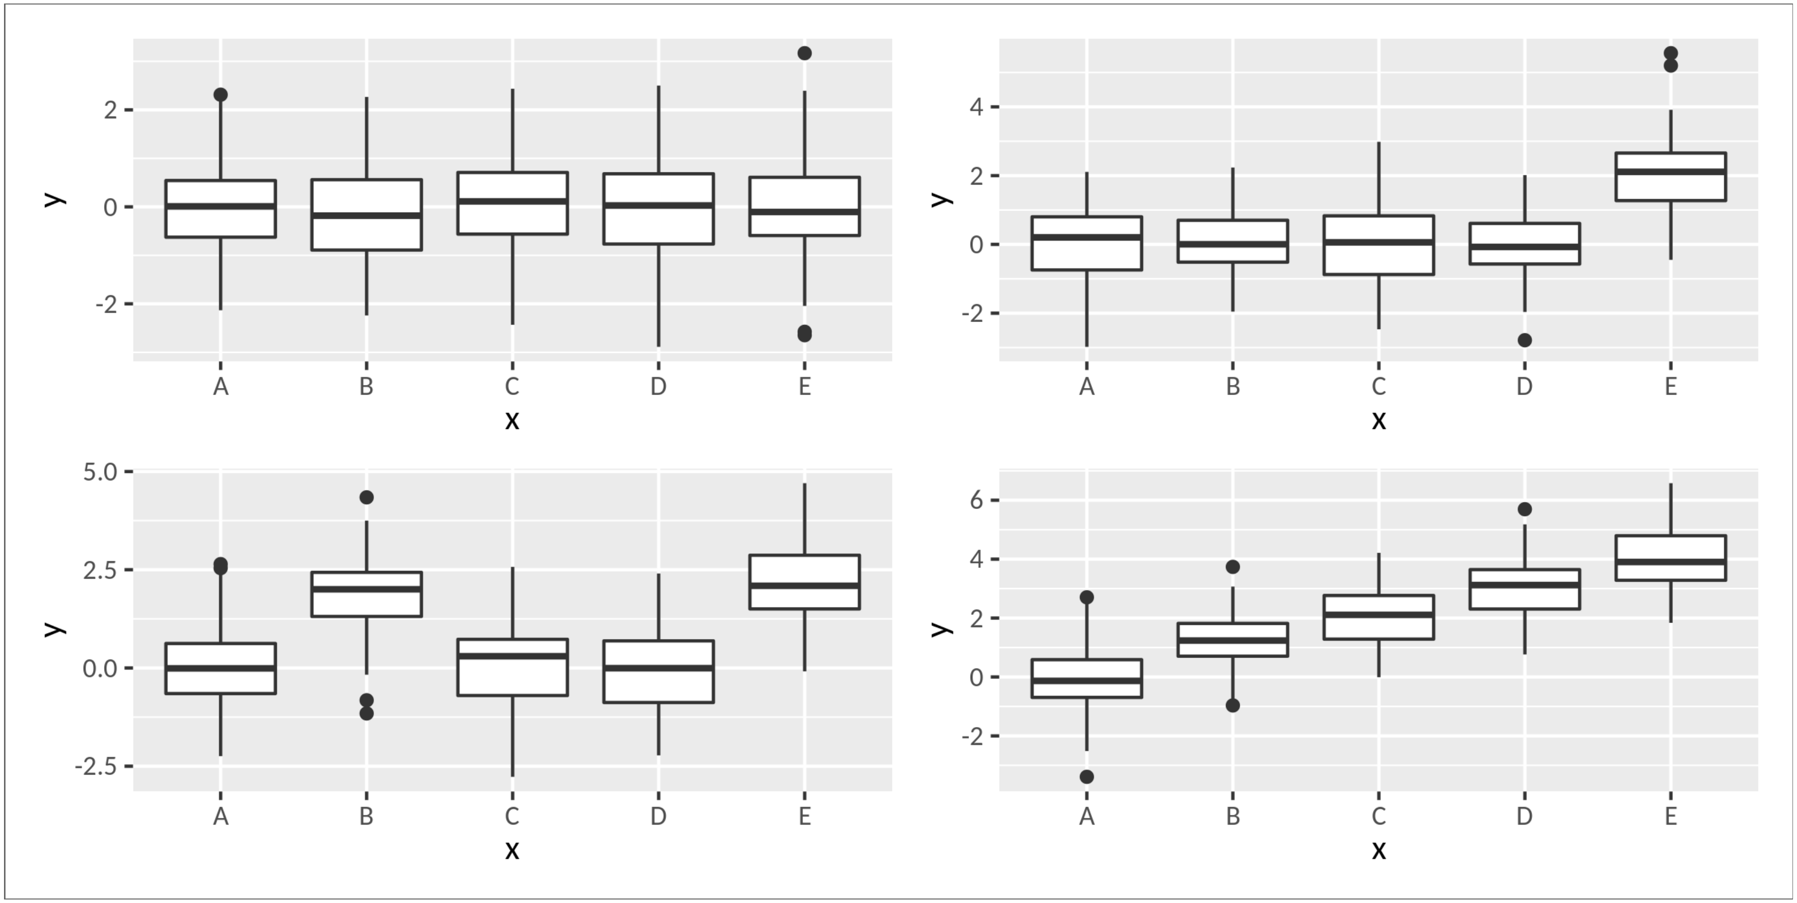

Single granularity plots

- Compute Jensen-Shannon divergences between distributions q_1 and q_2:

$$JSD(q_1, q_2) = \frac{1}{2}D(q_1, M) + \frac{1}{2}D(q_2, M),$$

  where $M = \frac{1}{2}(q_1 + q_2)$ and $D(q_1, q_2)$ is KL divergence.
- Measure effectiveness of a plot as maximum JSD for that plot (adjusted for number of levels).
- Users can be guided to view the most effective plots.

Normalization of maximum JSD

The distribution of max JSD depends on number of levels n.

Permutation approach (for small n)

- Compute max JSD after permuting the levels.
- Normalize by mean and standard deviation of permuted max JSD values

Modelling approach (for large n)

- Fit a Gumbel GLM to max JSD from simulated N(0,1) data with n as covariate.
- Standardize original data by $\Phi^{-1}()$, compute max JSD, and normalize by mean and standard deviation from model.

Single granularity plots

x	Normalized maximum JSD
hhour/week	72.8
day/year	67.0
week/year	31.8
hhour/day	24.4
day/week	21.8
month/year	15.0
day/month	-7.0
week/month	-10.5

Single granularity plots

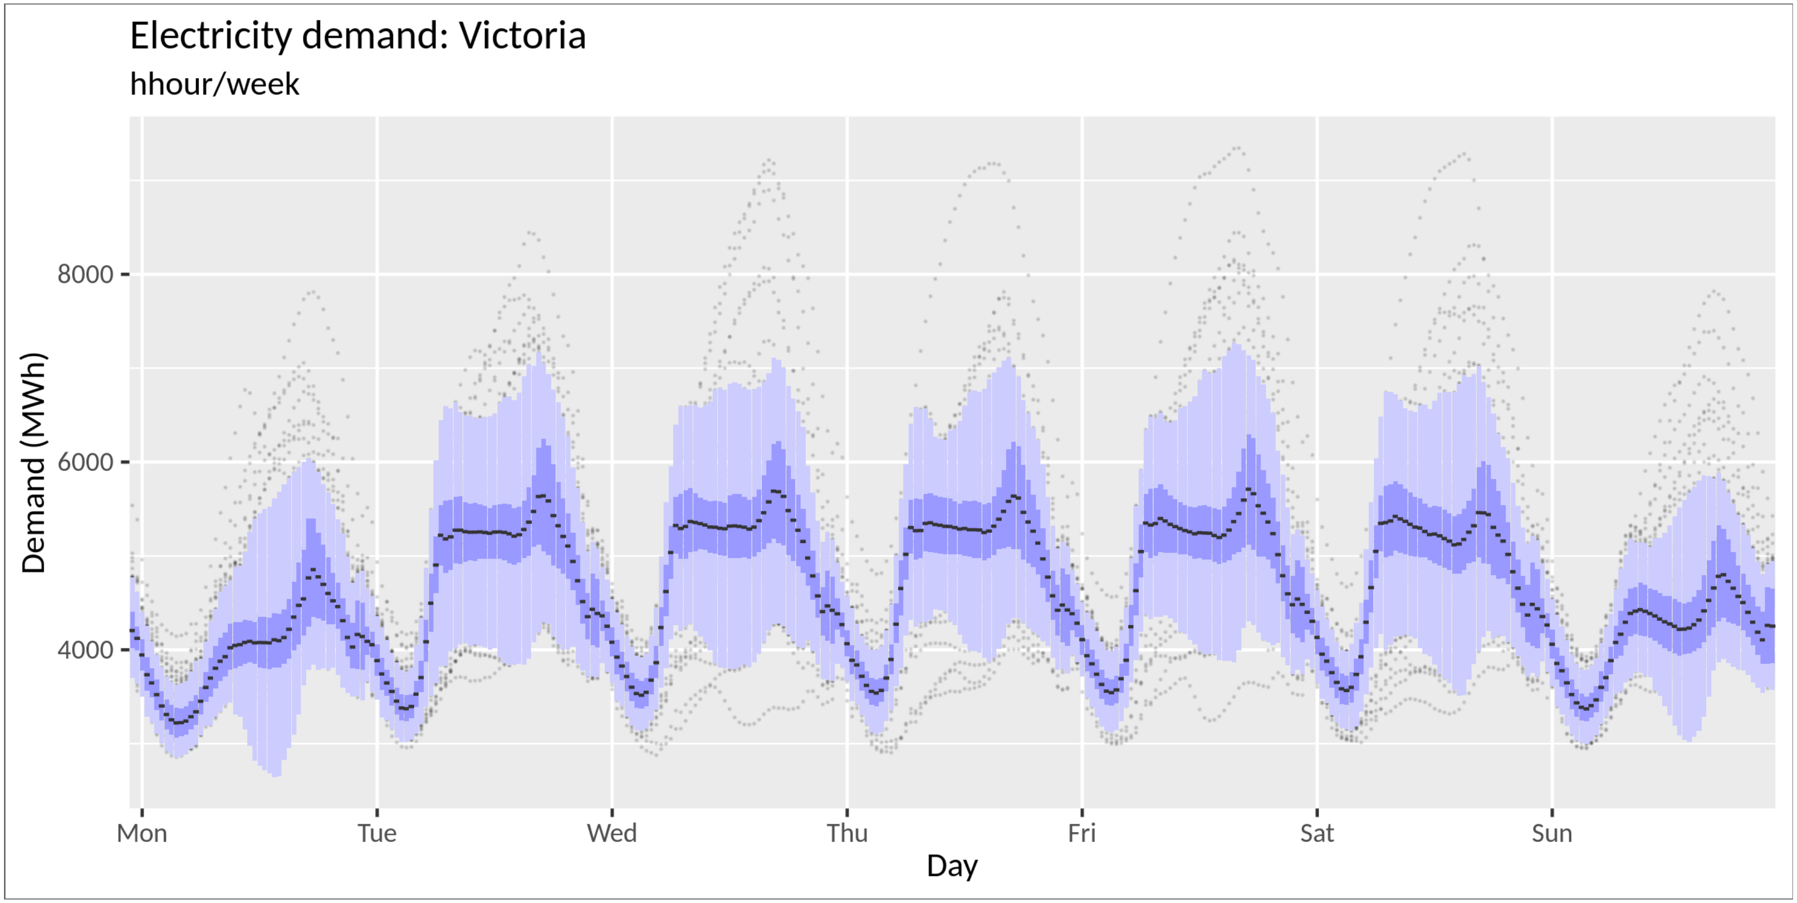

Single granularity plots

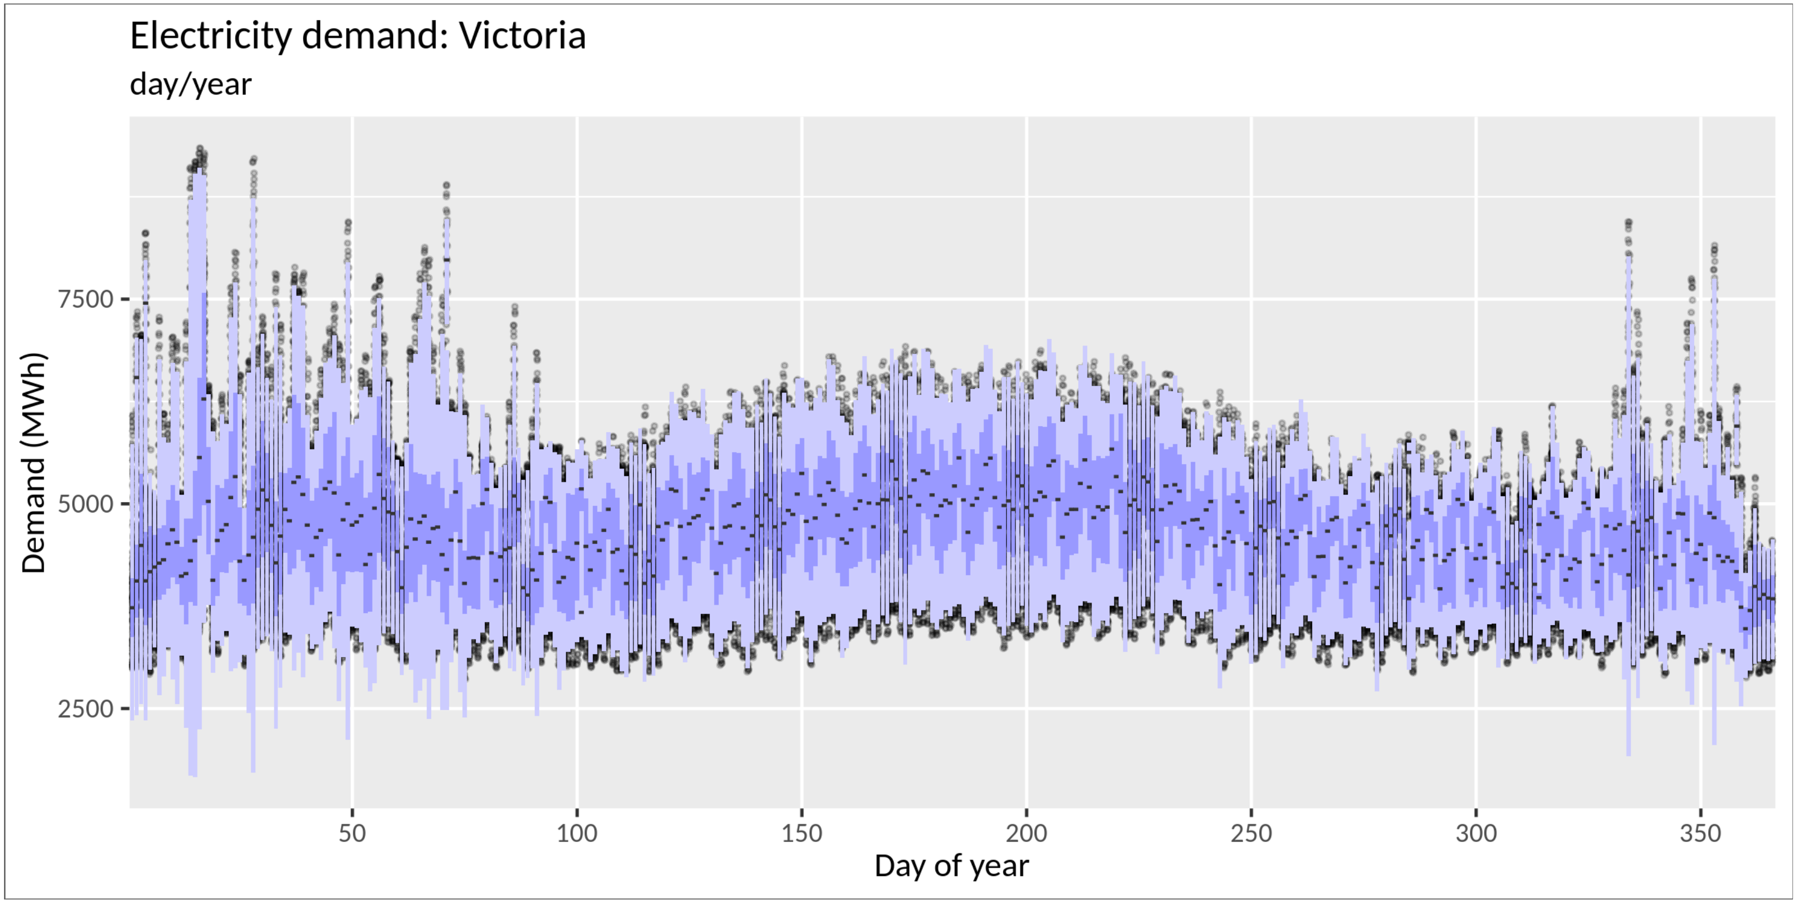

Single granularity plots

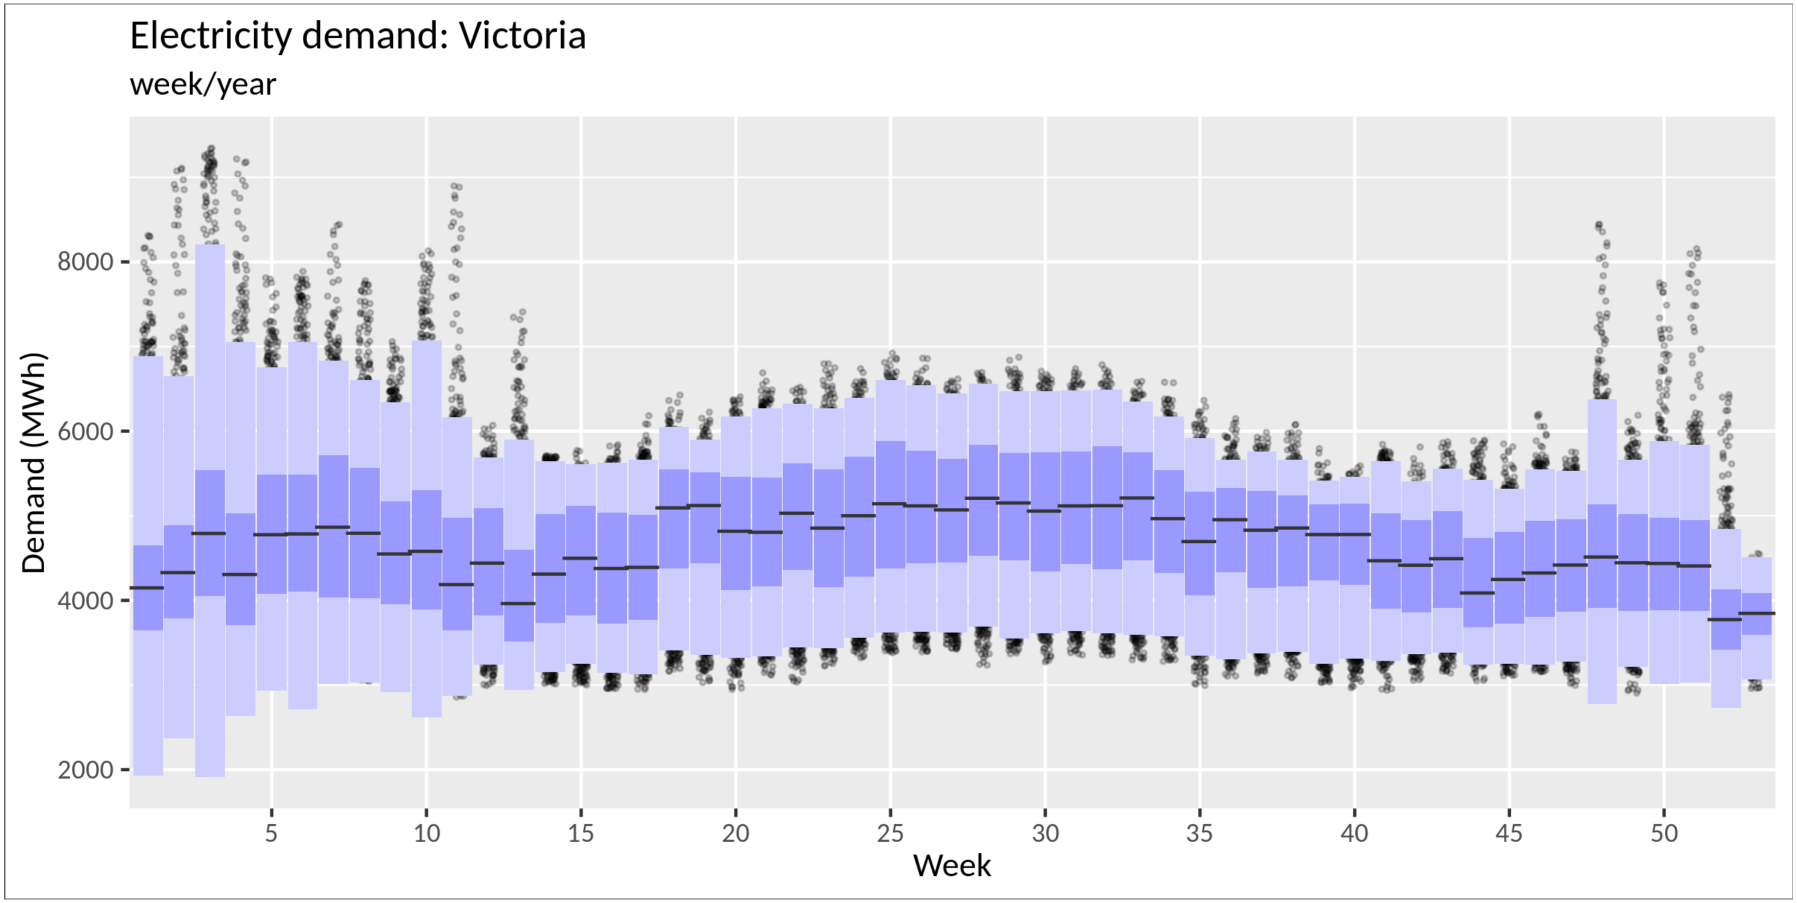

Single granularity plots

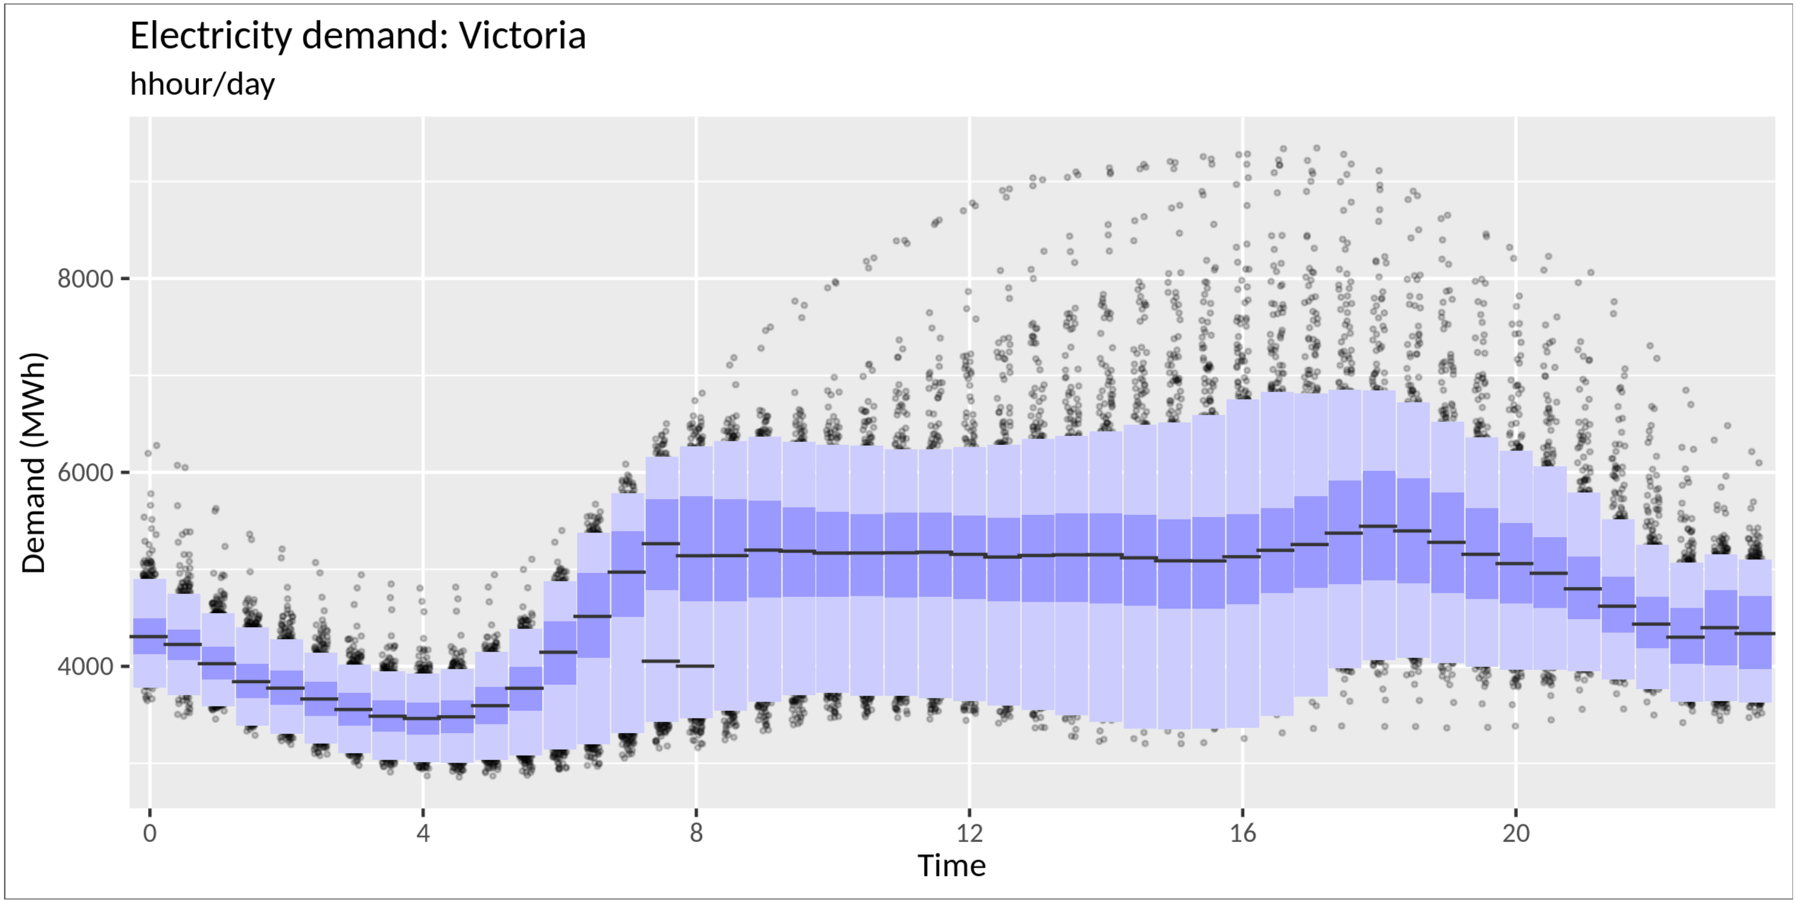

Single granularity plots

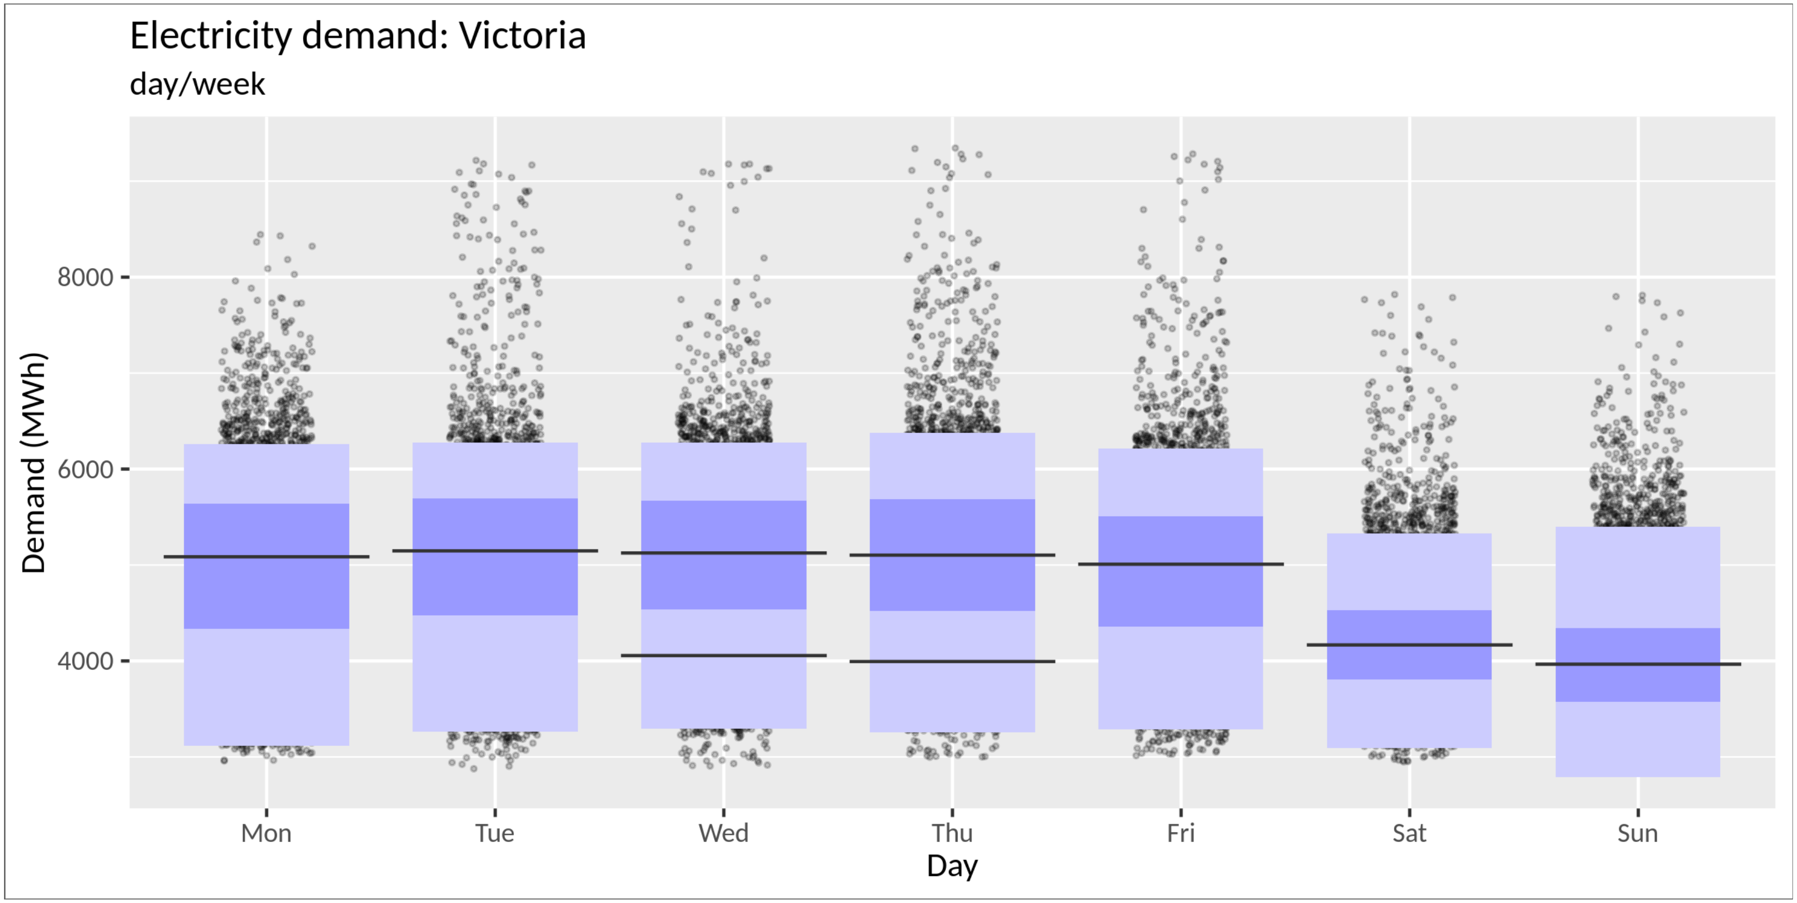

Single granularity plots

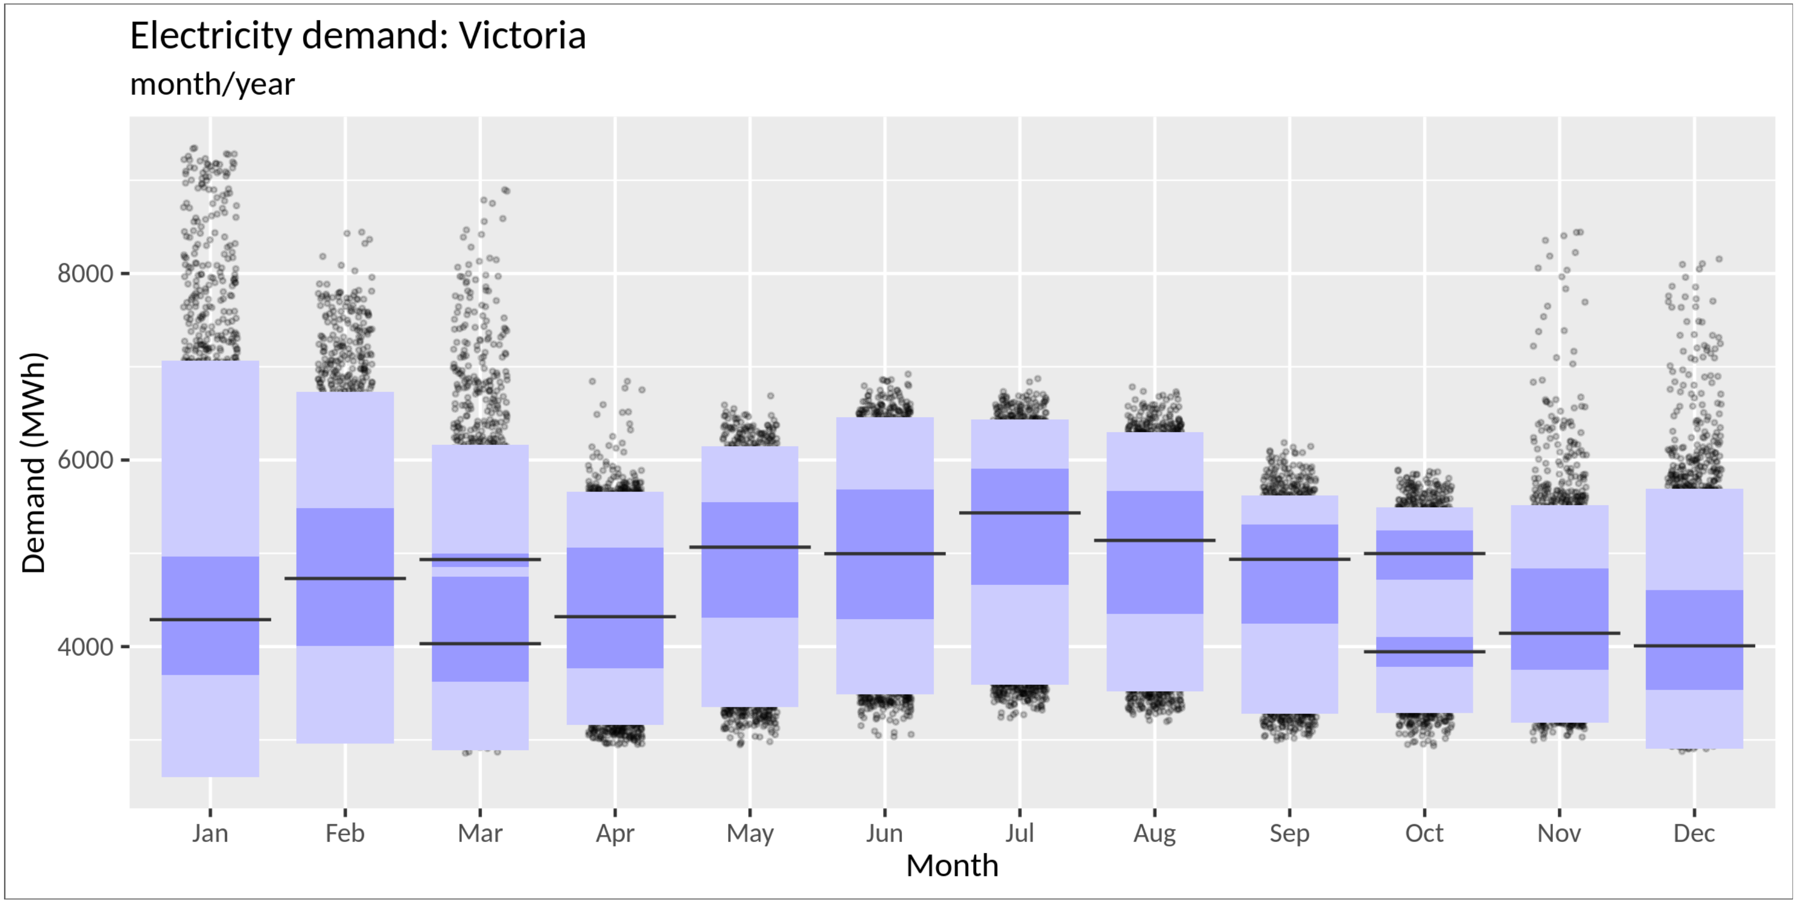

Single granularity plots

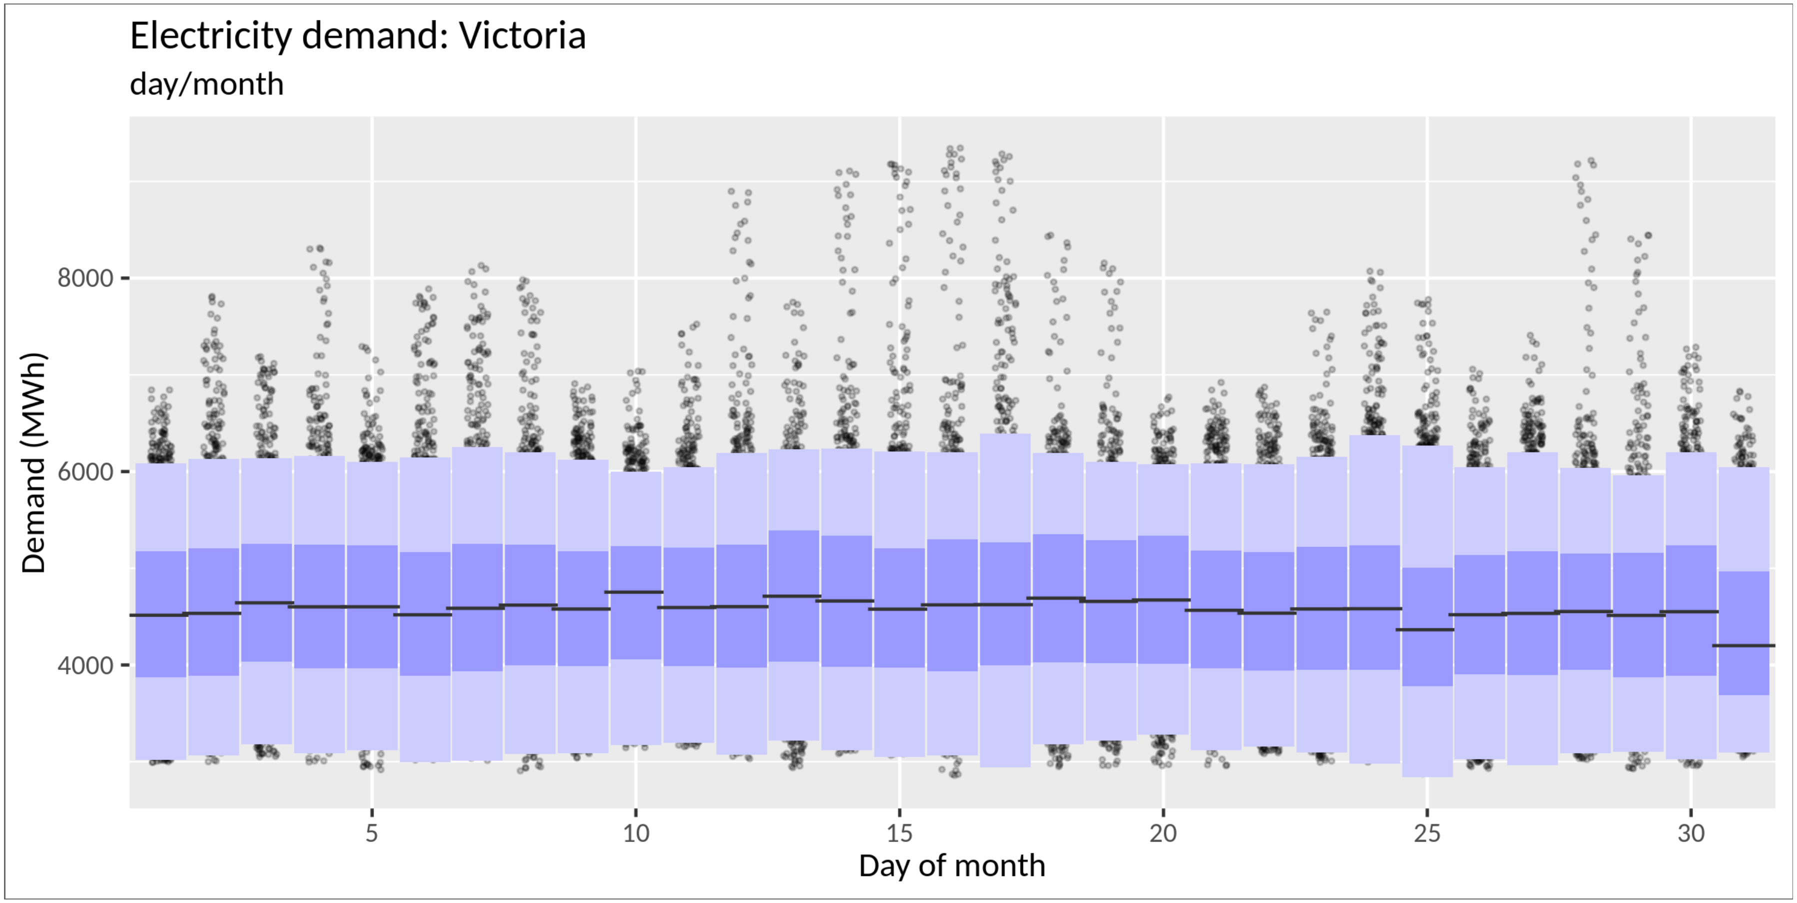

Single granularity plots

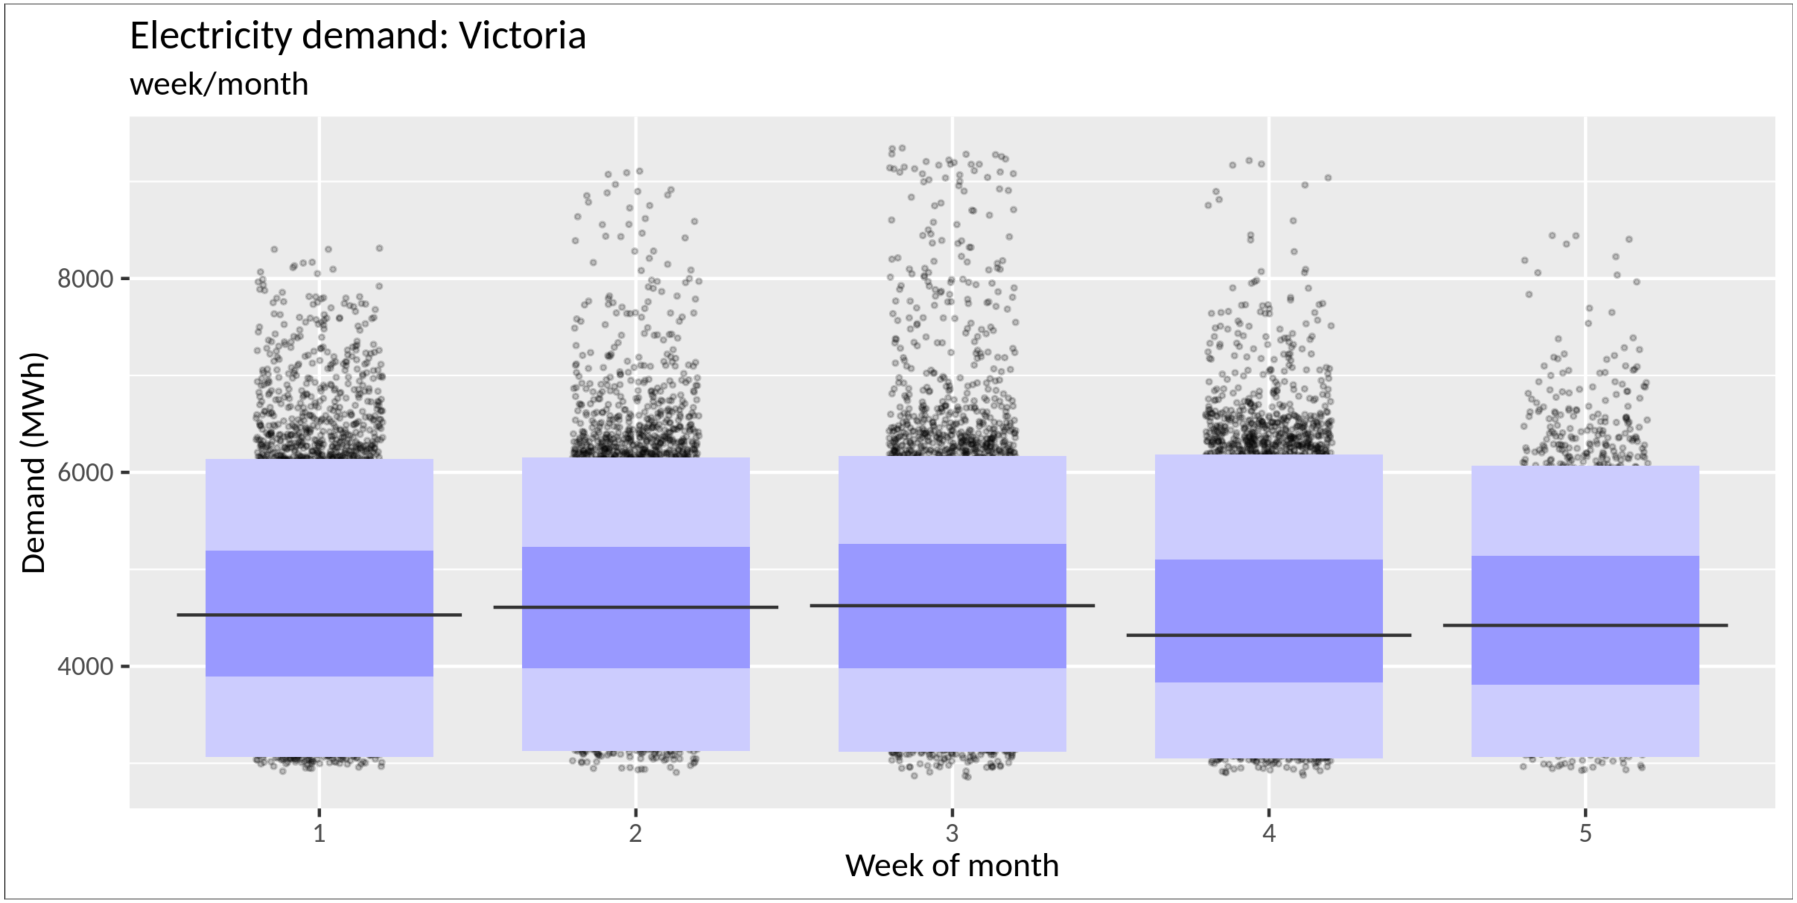

What is an interesting faceted plot?

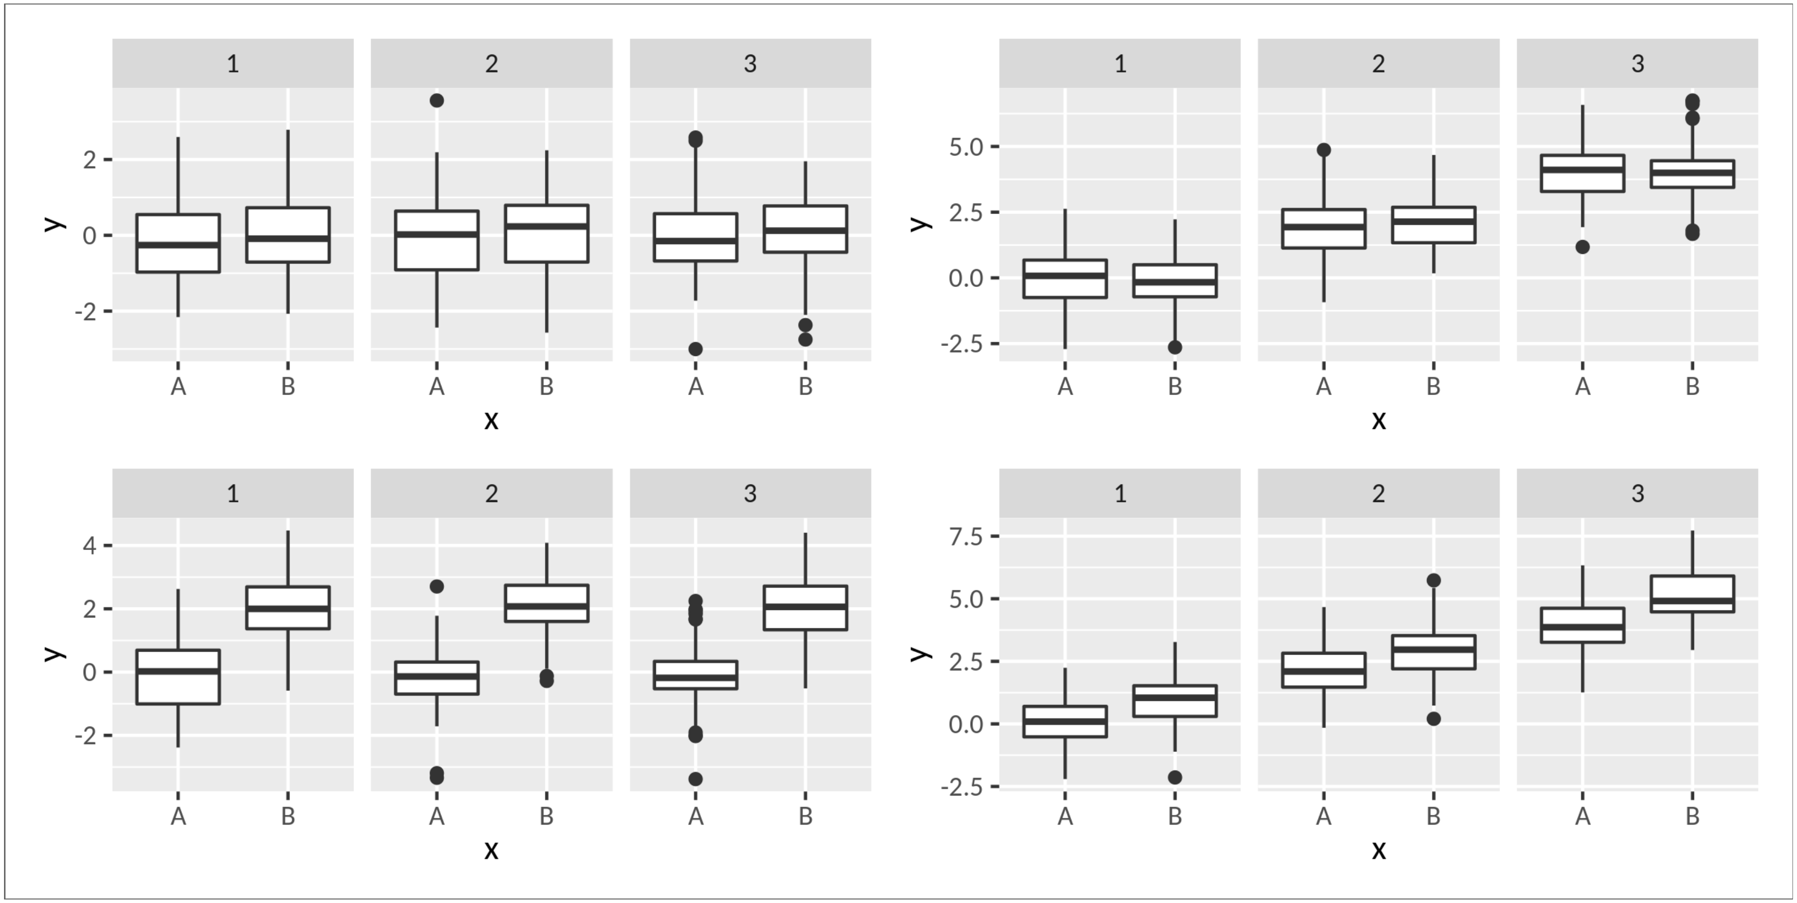

Faceted granularity plots

- Measure effectiveness of a plot as maximum JSD for that plot
  - weight within panel differences higher than between panel differences (weight 2:1)
  - normalization to adjust for number of levels and panels
- Omit combinations with empty or near-empty intersections ("clashes"). e.g., `day/year` × `month/year`
- Omit multi-step nested granularities. e.g,. `day/year`, `hhour/week`
- Omit facets with 20+ levels

Recommended faceted plots

facet	x	facet levels	x levels	Max JSD
`month/year`	`hhour/day`	12	48	123.7
`day/week`	`hhour/day`	7	48	76.6
`month/year`	`day/week`	12	7	63.3
`month/year`	`week/month`	12	5	59.7
`week/month`	`month/year`	5	12	55.9
`week/month`	`hhour/day`	5	48	51.5
`week/month`	`day/week`	5	7	47.4
`day/week`	`day/month`	7	31	44.1
`day/week`	`month/year`	7	12	37.7
`day/week`	`week/month`	7	5	23.9

Faceted granularity plots

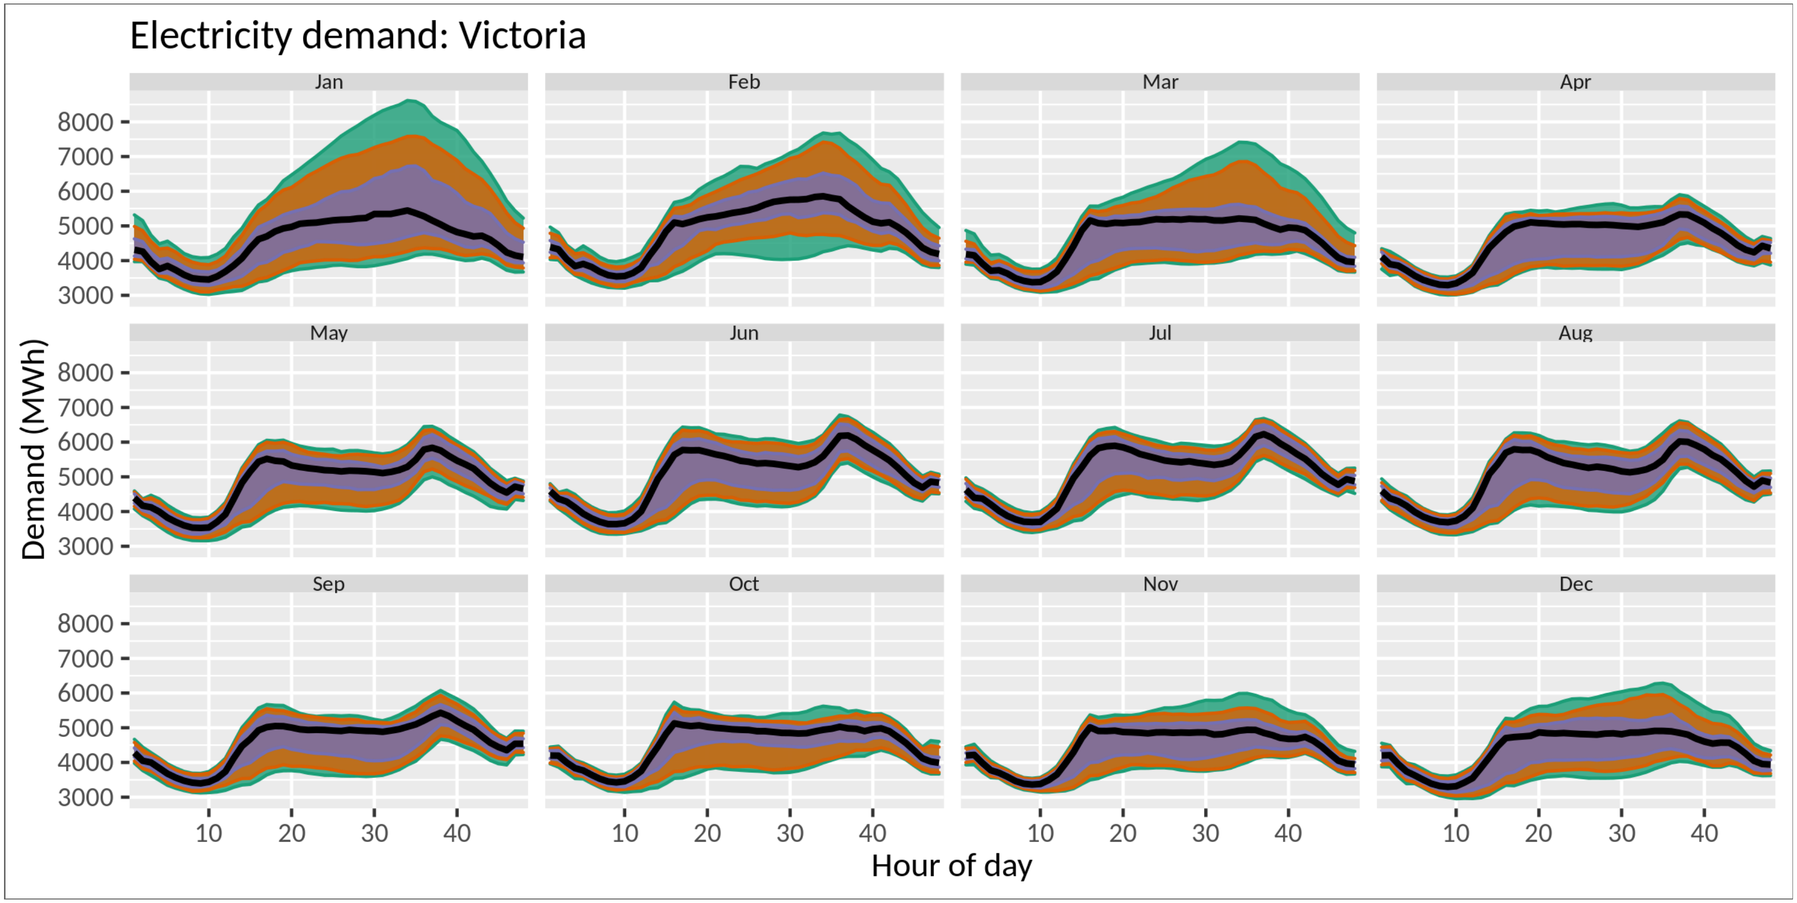

Faceted granularity plots

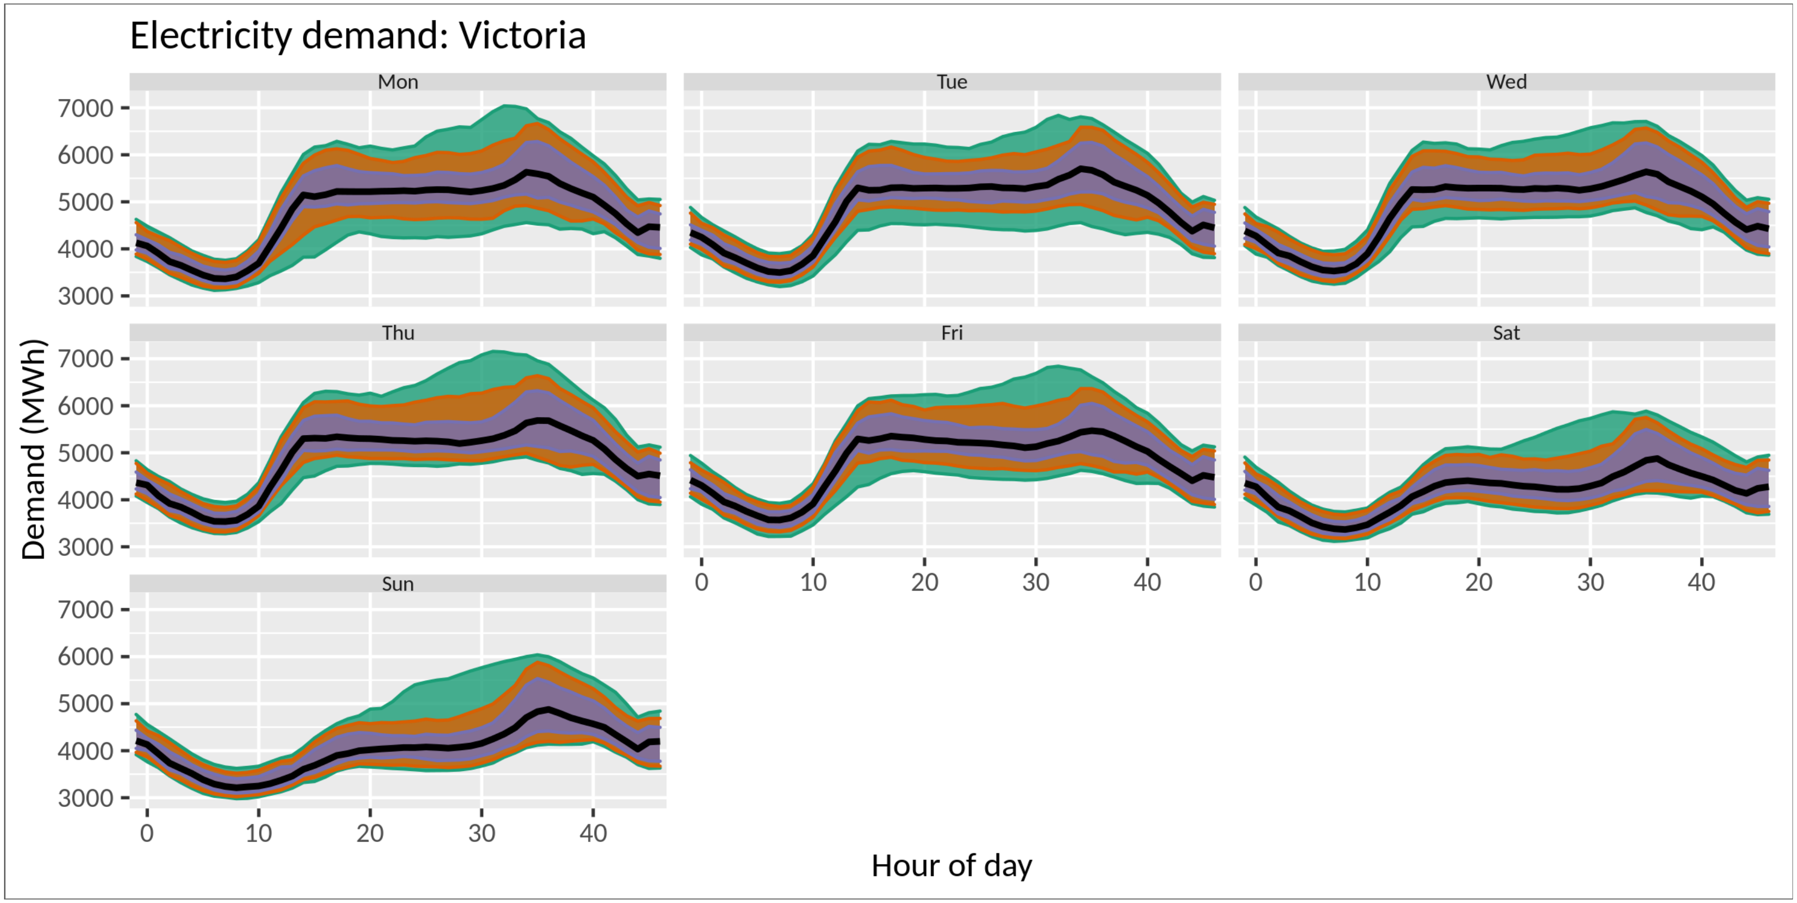

Faceted granularity plots

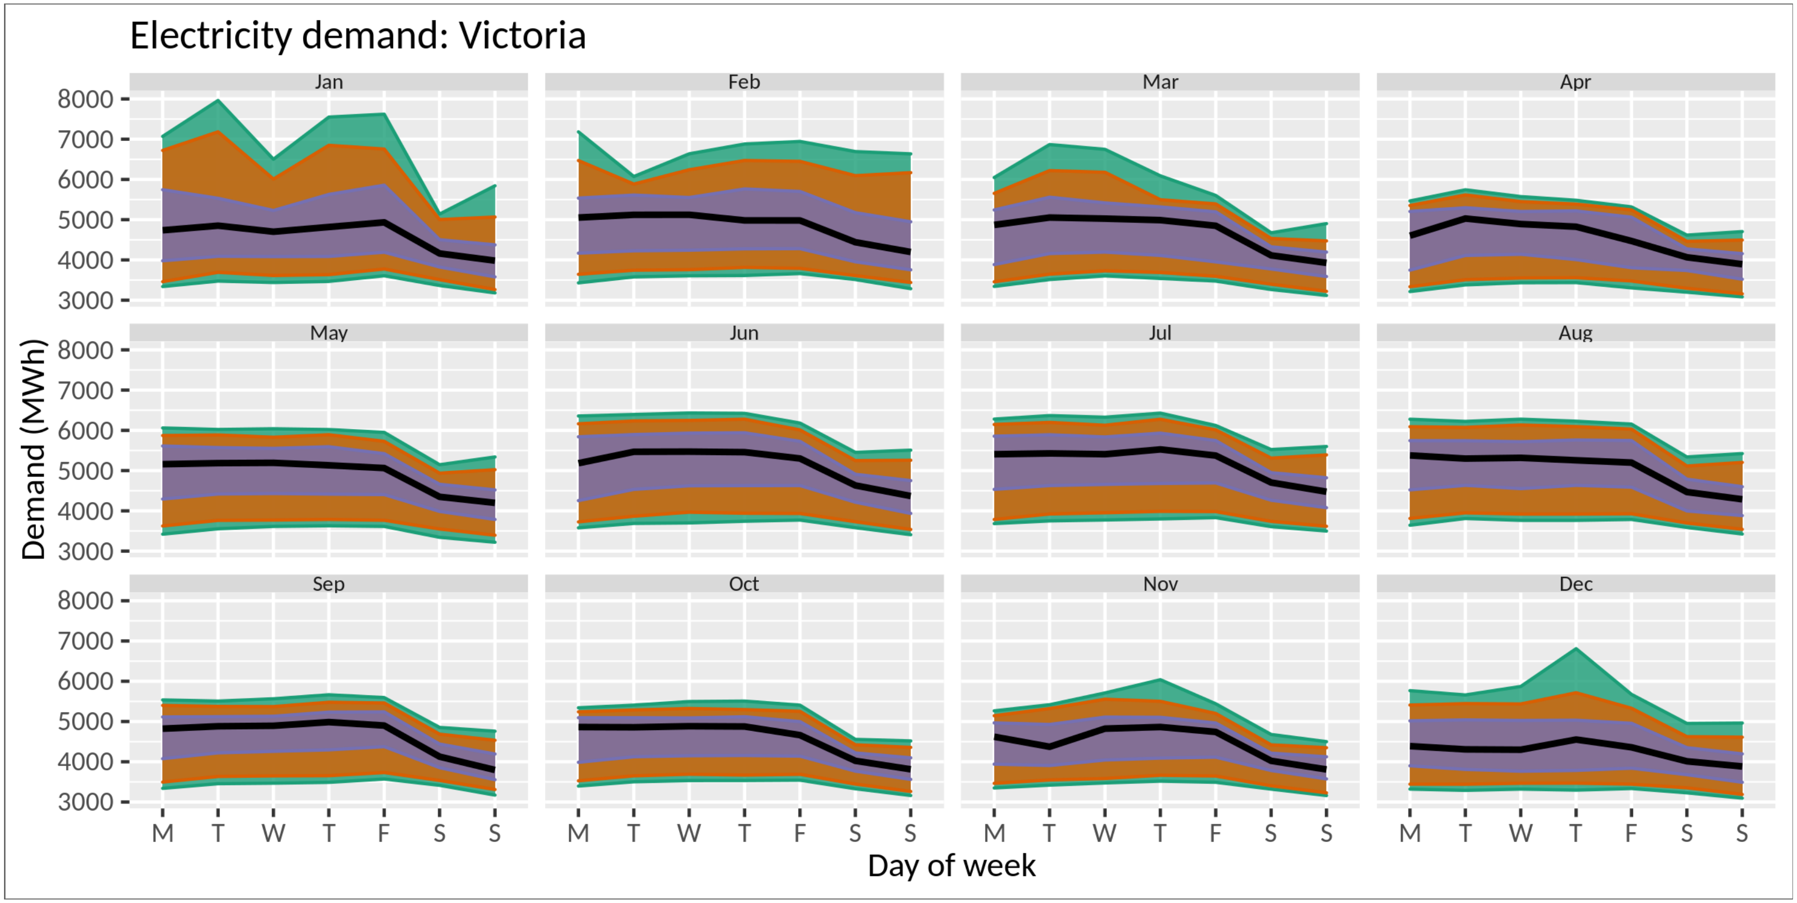

References

- Sayani Gupta, Rob J Hyndman, Dianne Cook and Antony Unwin (2022) Visualizing probability distributions across bivariate cyclic temporal granularities. *J Computational & Graphical Statistics*, **31**(1), 14-25.
- Sayani Gupta, Rob J Hyndman, Dianne Cook (2022) Detecting distributional differences between temporal granularities for exploratory time series analysis. Work in progress.

Time series decomposition for complex seasonalities

Time series decomposition for complex seasonalities

- Multiple seasonal periods, not necessarily nested
- Non-integer seasonality
- Irregular seasonal topography
- Seasonality that depends on covariates
- Complex seasonal topology

No existing decomposition method handles all of these.

Two solutions

1. MSTL: For multiple integer seasonal periods.
2. STR: For all types of complex seasonality.

MSTL

- Kasun Bandara, Rob J Hyndman, Christoph Bergmeir (2022) MSTL: A Seasonal-Trend Decomposition Algorithm for Time Series with Multiple Seasonal Patterns. *International J Operational Research*, to appear. **robjhyndman.com/publications/mstl/**
- For multiple integer seasonal periods with additive components
- Implemented in R packages `forecast` and `fable`.

MSTL

```
ec |>
odel(STL(Demand)) |>
omponents() |>
utoplot()
```

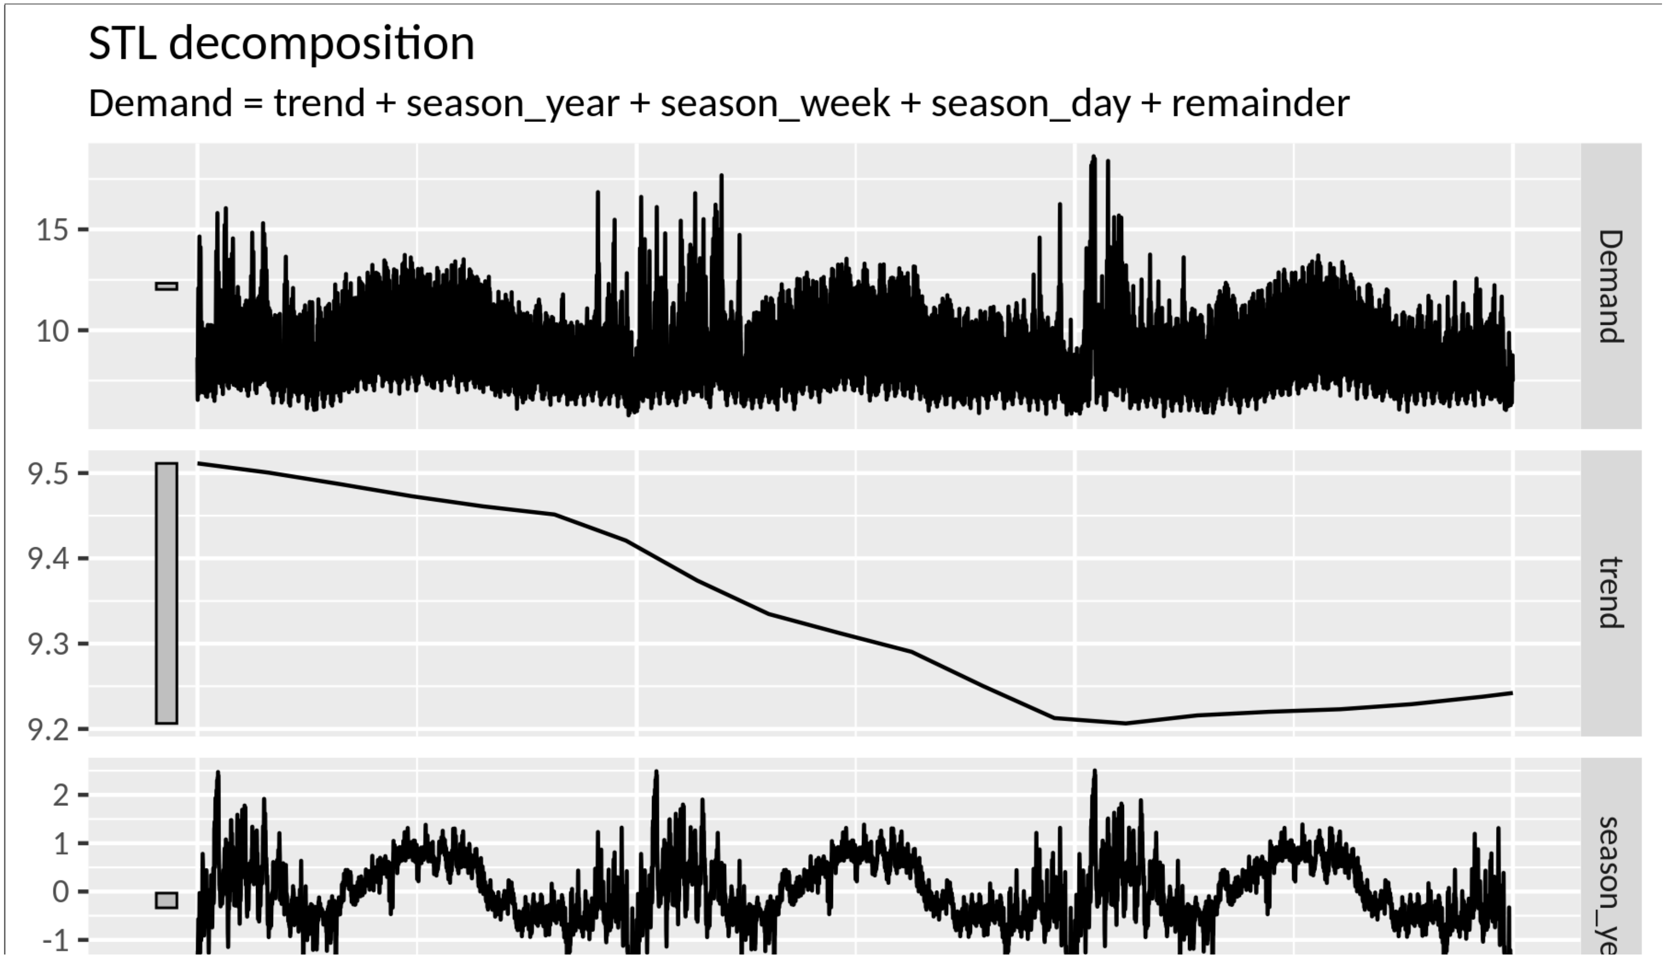

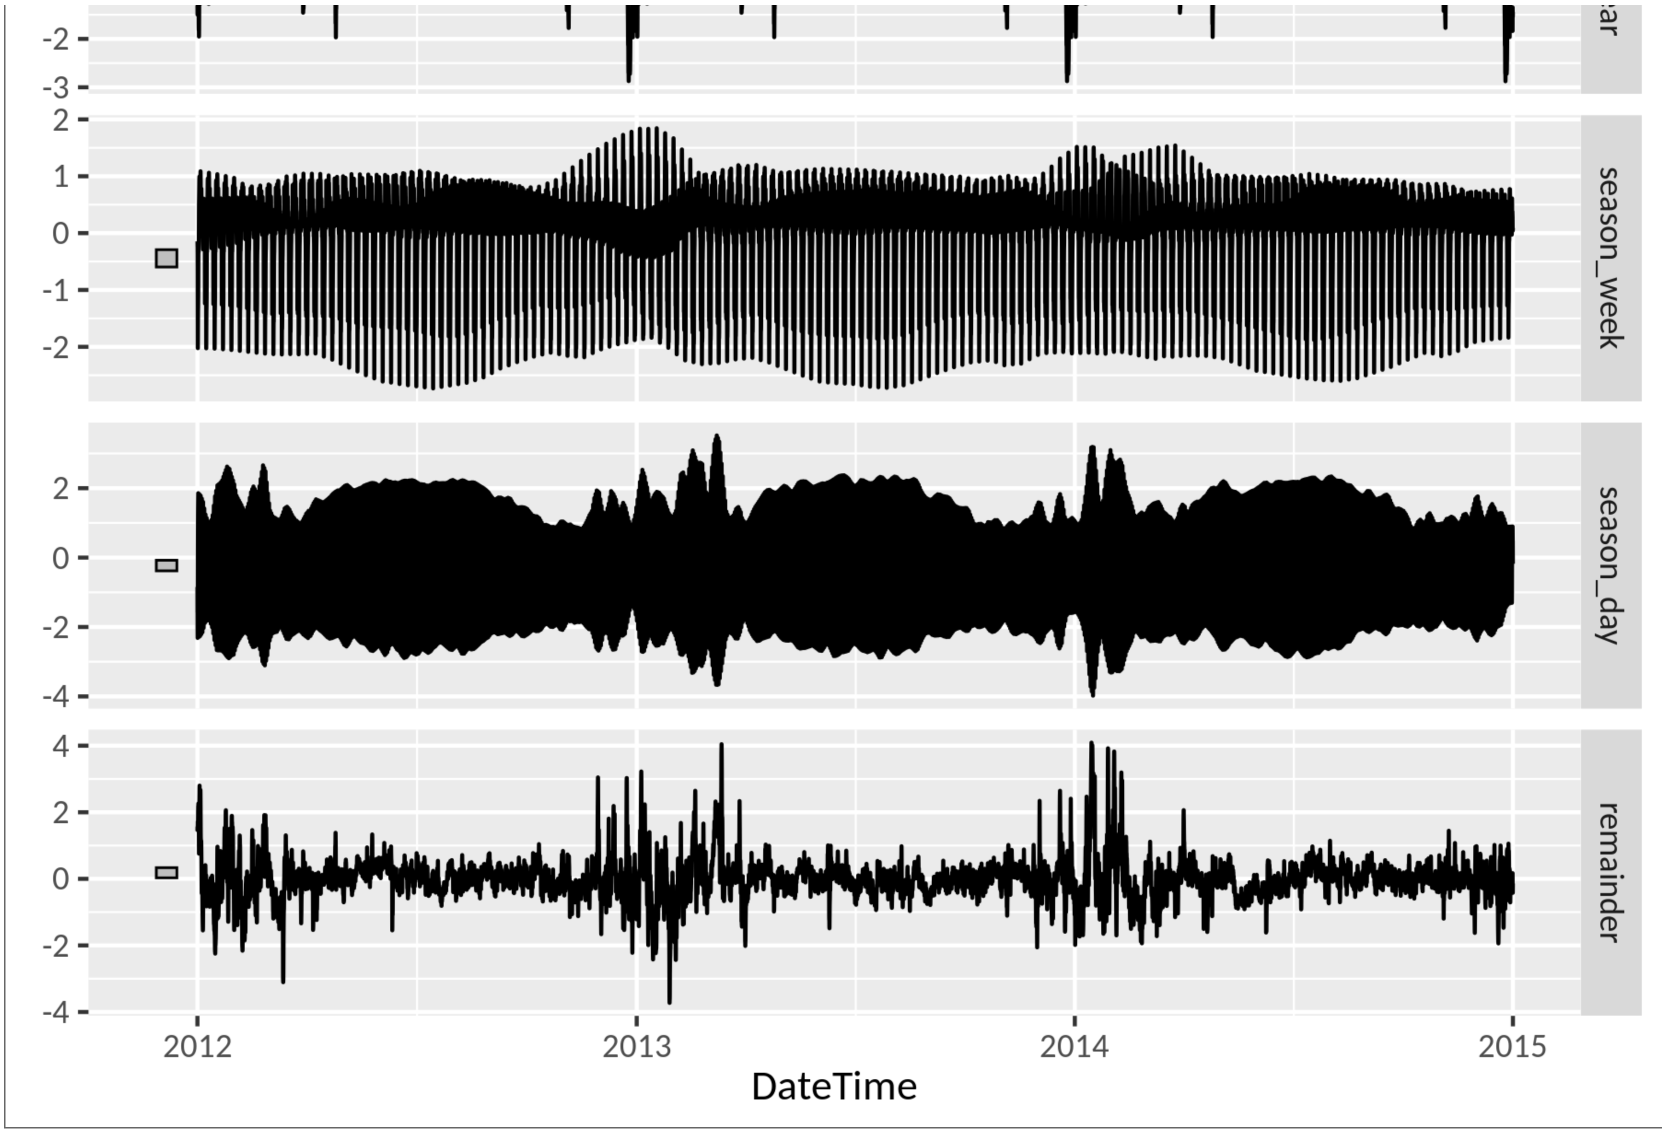

-2
-3
ar
2
1
0
-1
-2
season_week
2
0
-2
-4
season_day
4
2
0
-2
-4
remainder
2012
2013
2014
2015
DateTime

MSTL

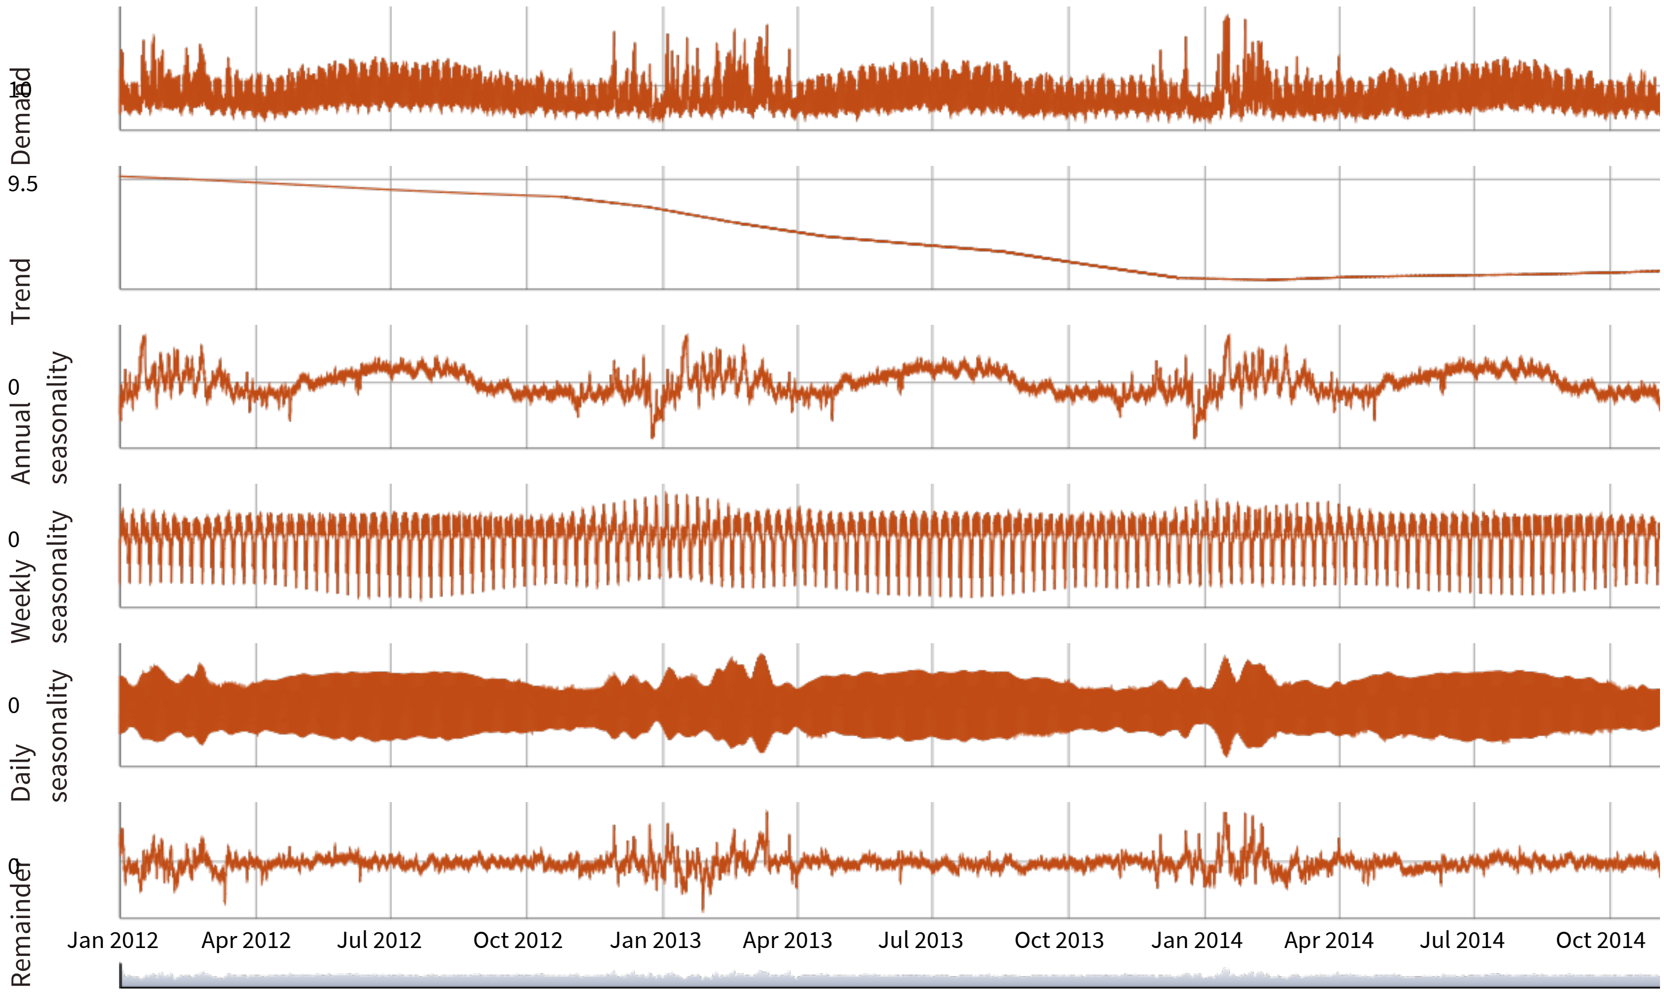
Demand
10
Trend
9.5
Annual seasonality
0
Weekly seasonality
0
Daily seasonality
0
Remainder
0
Jan 2012
Apr 2012
Jul 2012
Oct 2012
Jan 2013
Apr 2013
Jul 2013
Oct 2013
Jan 2014
Apr 2014
Jul 2014
Oct 2014

MSTL

$$y_t = T_t + \sum_{i=1}^{I} S_t^{(i)} + R_t$$

$$\begin{aligned} y_t = &\quad \text{observation at time } t \\ T_t = &\quad \text{smooth trend component} \\ S_t^{(i)} = &\quad \text{seasonal component } i \\ &\quad i = 1, \dots, I \\ R_t = &\quad \text{remainder component} \end{aligned}$$

Estimation

Components updated iteratively.

MSTL

```
: time series as vector
eriods: vector of seasonal periods in increasing order
window: seasonal window values
terate: number of  STL iterations

sonality <- matrix(0, nrow = length(y), ncol = length(periods))
eas <- y
 (j in 1:iterate) {
or (i in 1:length(periods)) {
 deseas <- deseas + seasonality[, i]
 fit <- stl(ts(deseas, frequency = periods[i]), s.window = swindow[i])
 seasonality[, i] <- fit$season
 deseas <- deseas - seasonality[, i]


nd <- fit$trend
ainder <- deseas - trend
urn(trend, seasonality, remainder)
```

MSTL

fable syntax

```
tsibble |>
  model(STL(variable ~ season(period = a, window = b) +
                       season(period = c, window = d)))
```

forecast syntax

```
vector |>
  msts(seasonal.periods = c(a, c)) |>
  mstl(s.window = c(b, d))
```

STR

- Alex Dokumentov and Rob J Hyndman (2022) STR: Seasonal-Trend decomposition using Regression. *INFORMS Journal on Data Science*, **1**(1), 50-62. **robjhyndman.com/publications/str/**
- Implemented in R package `stR`.

STR

$$y_t = T_t + \sum_{i=1}^{I} S_t^{(i)} + \sum_{p=1}^{P} \phi_{p,t} z_{t,p} + R_t$$

$T_t =$ smooth trend component

$S_t^{(i)} =$ seasonal component i (possibly complex topology)

$z_{p,t} =$ covariate with coefficient $\phi_{p,t}$ (possibly time-varying)

$R_t =$ remainder component

Estimation

Components estimated using penalized MLE

Smoothness via difference operators

Smooth trend obtained by requiring $\Delta_2 T_t \sim \text{NID}(0, \sigma_L^2)$

- $\Delta_2 = (1 - B)^2$ where $B =$ backshift operator
- σ_L controls smoothness

$$f(\boldsymbol{D}_\ell \boldsymbol{\ell}) \propto \exp\left\{-\frac{1}{2}\left\|\boldsymbol{D}_\ell \boldsymbol{\ell}/\sigma_L\right\|_{L_2}^2\right\}$$

- $\boldsymbol{\ell} = \langle T_t \rangle_{t=1}^n$
- $\boldsymbol{D}_\ell =$ 2nd difference operator matrix: $\boldsymbol{D}_\ell \boldsymbol{\ell} = \langle \Delta^2 T_t \rangle_{t=3}^n$

Smooth 2D seasonal surfaces

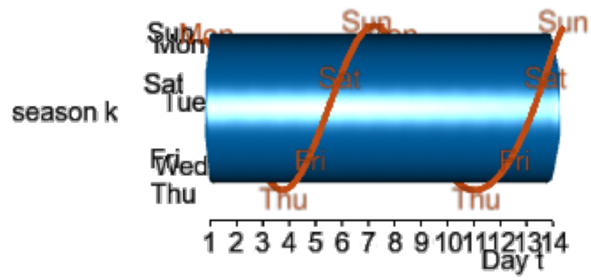

- m_i = number of "seasons" in $S_t^{(i)}$.
- $S_{k,t}^{(i)}$ = 2d season ($k = 1, \dots, m_i; t = 1, \dots, n$)
- $\sum_k S_{k,t}^{(i)} = 0$ for each t.

Smooth 2D seasonal surfaces

- $\boldsymbol{S}^{(i)} = [S^{(i)}_{k,t}]$ the ith seasonal surface matrix
- $\boldsymbol{s}_i = \text{vec}(\boldsymbol{S}_i) =$ the ith seasonal surface in vector form

Smoothness in time t direction:

$$\boldsymbol{D}_{tt,i}\boldsymbol{s}_i = \langle \Delta_t^2 \boldsymbol{S}^{(i)}_{k,t} \rangle \sim \text{NID}(\boldsymbol{0}, \sigma_i^2 \boldsymbol{\Sigma}_i)$$

$$f(\boldsymbol{s}_i) \propto \exp\left\{ -\frac{1}{2} \| \boldsymbol{D}_{tt,i}\boldsymbol{s}_i/\sigma_i \|^2_{L_2} \right\}$$

Analogous difference matrices $\boldsymbol{D}_{kk,i}$ and $\boldsymbol{D}_{kt,i}$ ensure smoothness in season and time-season directions.

Gaussian remainders

- $R_t \sim \text{NID}(0, \sigma_R^2)$.
- $\boldsymbol{y} = [y_1, \dots, y_n]' =$ vector of observations
- $\boldsymbol{Z} = [z_{t,p}] =$ covariate matrix with coefficient $\boldsymbol{\Phi} = [\phi_{p,t}]$
- $\boldsymbol{Q}_i =$ matrix that extracts $\langle S^{(i)}_{\kappa(t),t} \rangle_{t=1}^{n}$ from $\boldsymbol{s}_i$.
- Residuals: $\boldsymbol{r} = \boldsymbol{y} - \sum_i \boldsymbol{Q}_i \boldsymbol{s}_i - \boldsymbol{\ell} - \boldsymbol{Z}\boldsymbol{\Phi}$ have density

$$f(\boldsymbol{r}) \propto \exp\left\{-\frac{1}{2}\left\|\boldsymbol{r}/\sigma_R\right\|_{L_2}^2\right\},$$

MLE for STR

Minimize wrt $\boldsymbol{\Phi}$, $\boldsymbol{\ell}$ and $\boldsymbol{s}_i$:

$$-\log \mathcal{L} = \frac{1}{2\sigma_R}\left\{ \left\| \boldsymbol{y} - \sum_{i=1}^{I} \boldsymbol{Q}_i \boldsymbol{s}_i - \boldsymbol{\ell} - \boldsymbol{Z}\boldsymbol{\Phi} \right\|_{L_2}^2 + \lambda_\ell \left\| \boldsymbol{D}_\ell \boldsymbol{\ell} \right\|_{L_2}^2 \right.$$
$$\left. + \sum_{i=1}^{I} \left(\|\lambda_{tt,i} \boldsymbol{D}_{tt,i} \boldsymbol{s}_i\|_{L_2}^2 + \|\lambda_{st,i} \boldsymbol{D}_{st,i} \boldsymbol{s}_i\|_{L_2}^2 + \|\lambda_{ss,i} \boldsymbol{D}_{ss,i} \boldsymbol{s}_i\|_{L_2}^2 \right) \right\}$$

Equivalent to linear model

$$\boldsymbol{y}_+ = \boldsymbol{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}$$

- $\boldsymbol{y}_+ = [\boldsymbol{y}',\ \mathbf{0}']'$
- $\boldsymbol{\varepsilon} \sim N(\mathbf{0}, \sigma_R^2 \boldsymbol{I})$

$$\boldsymbol{X} = \begin{bmatrix} \boldsymbol{Q}_1 & \dots & \boldsymbol{Q}_I & \boldsymbol{I}_n & \boldsymbol{Z} \\ \lambda_{tt,1}\boldsymbol{D}_{tt,1} & \dots & 0 & 0 & 0 \\ \lambda_{st,1}\boldsymbol{D}_{st,1} & \dots & 0 & 0 & 0 \\ \lambda_{ss,1}\boldsymbol{D}_{ss,1} & \dots & 0 & 0 & 0 \\ 0 & \ddots & 0 & 0 & 0 \\ 0 & \dots & \lambda_{tt,I}\boldsymbol{D}_{tt,I} & 0 & 0 \\ 0 & \dots & \lambda_{st,I}\boldsymbol{D}_{st,I} & 0 & 0 \\ 0 & \dots & \lambda_{ss,I}\boldsymbol{D}_{ss,I} & 0 & 0 \\ 0 & \dots & 0 & \lambda_{\ell}\boldsymbol{D}_{tt} & 0 \end{bmatrix}$$

STR

Three seasonal components, quadratic temperature regressors

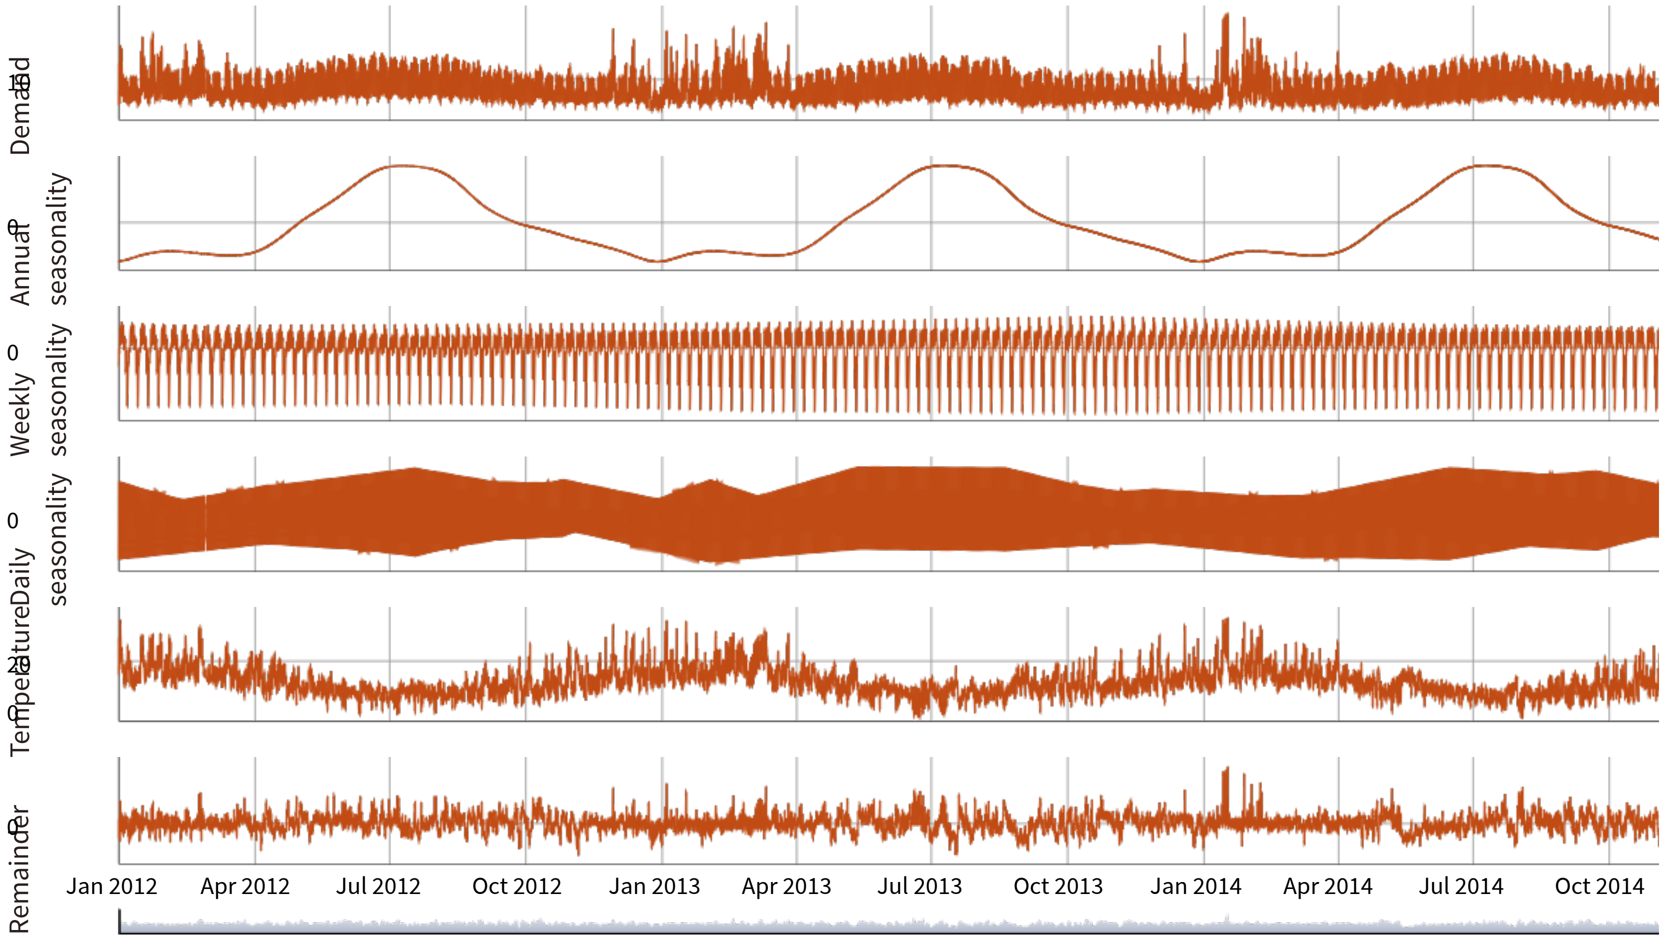

STR outliers

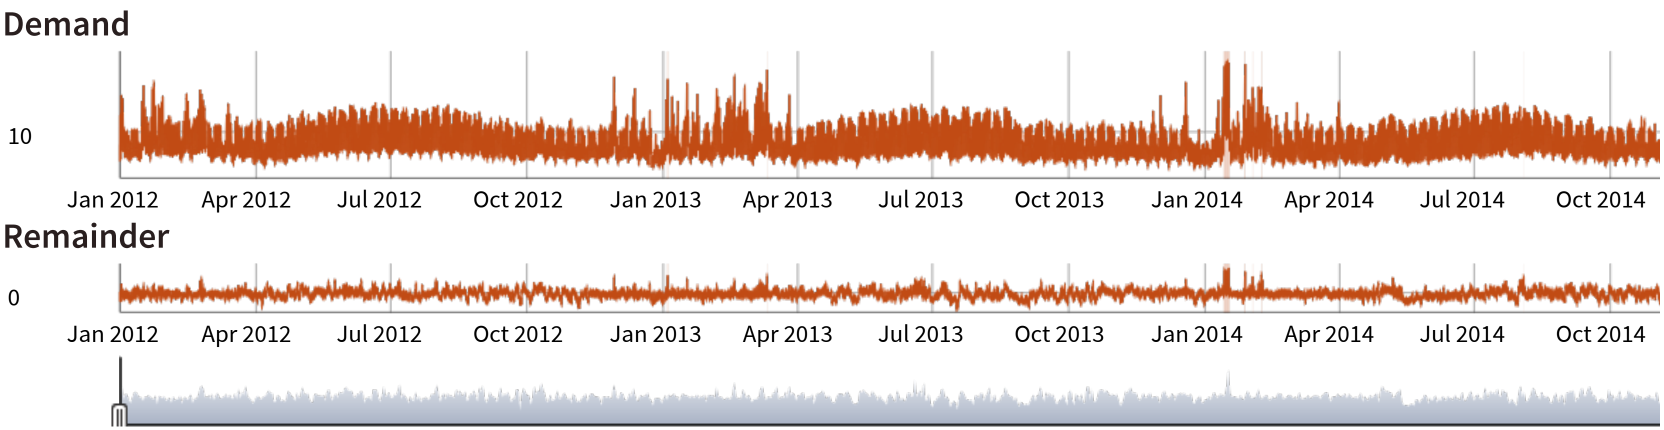

Show 5 entries Search:

	Date	Time	Temperature	Demand	Remainder
1	2014-01-17	16:00	34.53	18.51	9.40
2	2014-01-17	15:00	33.93	18.46	9.22
3	2014-01-16	15:00	34.25	18.43	8.84
4	2014-01-17	14:00	33.32	18.32	8.78
5	2014-01-14	16:00	33.89	18.06	8.62

Showing 1 to 5 of 50 entries

Previous 1 2 3 4 5 … 10 Next

R packages

gravitas

hakear

forecast

feasts

stR

Thanks to my collaborators

OTexts.com/fppit

Forecasting: Principles and Practice

Published by OTexts™ with bookdown

Benvenuti nel nostro libro online sulle previsioni.

Questo testo si propone di fornire un'introduzione completa ai metodi di previsione e di fornire informazioni sufficienti su ciascun metodo affinché i lettori possano utilizzarli in modo sensato. Non viene offerta una discussione approfondita dei dettagli teorici che stanno alla base di ogni metodo, anche se molti di questi dettagli sono disponibili nei riferimenti presenti alla fine di ogni capitolo.

Il libro si rivolge a tre tipologie di pubblico:
(1) professionisti chiamati a fare previsioni in azienda senza aver ricevuto una formazione specifica in materia;
(2) studenti universitari di economia;
(3) studenti di master impegnati in un corso di previsione. Noi stessi lo utilizziamo per gli studenti di master e per gli studenti del terzo anno della Monash University, in

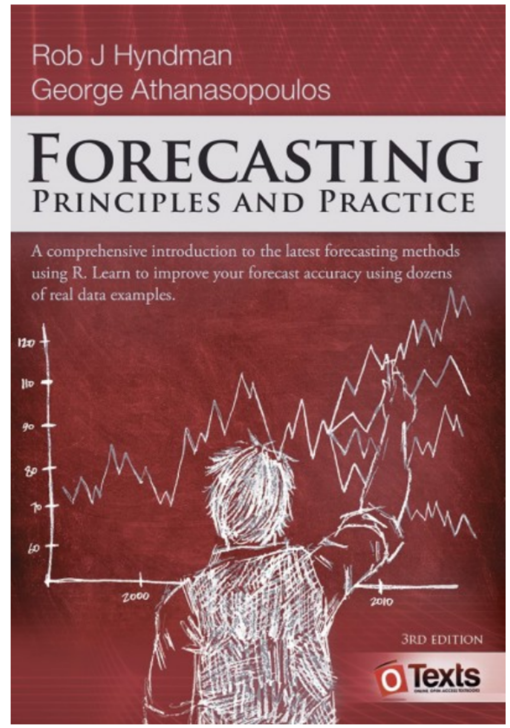

Compra una versione cartacea

For more information

Slides: robjhyndman.com/seminars/padova2022.html

Find me at:

robjhyndman.com
@robjhyndman
@robjhyndman
rob.hyndman@monash.edu