2017 Employment Projections Technical Report

Published September 2017

Washington State Employment Security Department
Dale Peinecke, *Commissioner*

Workforce Information and Technology Services
Cynthia Forland, Ph.D., *Chief Information Officer*

Workforce Information and Technology Services
Scott Wheeler, *Workforce Information Operations Manager*

Prepared by

Alex Roubinchtein, *Projections and Statistical Analyses Manager*
Bruce Nimmo, *Economic Analyst*

This workforce product was funded by a grant awarded by the U.S. Department of Labor's Employment and Training Administration. The product was created by the recipient and does not necessarily reflect the official position of the U.S. Department of Labor. The Department of Labor makes no guarantees, warranties, or assurances of any kind, expressed or implied, with respect to such information, including any information on linked sites and including, but not limited to, accuracy of the information or its completeness, timeliness, usefulness, adequacy, continued availability or ownership.

For more information or to get this report in an alternative format, call or email Bruce Nimmo at 360-407-4576 or bnimmo@esd.wa.gov.

Summary

Annually, the Employment Security Department (ESD) creates 2-, 5- and 10-year employment projections.

Projections results are built on state and national requirements, available data, current software tools and stakeholder input.

The projections process consists of two major steps: the creation of industry projections and the conversion of industry to occupational projections. The conversion process is based on Occupational Employment Statistics (OES) survey data.

Employment projections start with time series of covered employment processed at the county level within the unemployment insurance system. National forecasts from Global Insight are used as regressors for aggregated state employment forecasts.

Projection models for industry series are not predefined. This means assumptions are not made about which models are best for any given series. A software based optimization process selects the best combination of model outputs. The result is that model output selection may vary for each industry employment series.

We eliminate the need to manually choose the best model by allowing a software-based optimization process to select the best combination of model outputs.

Introduction

In this paper we discuss the technical processes used to produce industry and occupational projections for the Washington State Employment Security Department.

Data preparation and forecasting are done using R-software.

The projections process utilizes six models. The six models are: innovations state space exponential smoothing, naive, dynamic linear, ARIMA, hierarchy and an optimization process that combines outputs from the first four models. Only the dynamic linear and ARIMA models use regressors.

The hierarchy model is new this year. Hierarchical time series forecasting functions are found in R's *hts* package. The *hts* package specializes in forecasting time series that can be disaggregated into hierarchical structures using attributes such as geography. Forecasts are generated for each series at each level of the hierarchy. These forecasts are then combined and balanced by an optimization function within this package. The combination approach optimally combines independent base forecasts and generates a set of revised forecasts that are as close as possible to the initial univariate forecasts, but also balanced within the hierarchical structure.

Important new parameters created in this round of projections are "historical trend growth rates occurring after a major breaking point" for each series. These parameters are *historical trend growth rates.* To define these rates we used R's *BFAST* (Breaks for Additive Season and Trend) package.

Industry projections are produced at two levels: aggregated and detailed. The aggregated series are referred to as "series" and the detailed series are referred to as Industry Control Totals (ICT). For each of the series (aggregated and detailed), we produce multiple forecasts.

Selected state series projections are used as regressors for regional workforce development area (WDA) projections and state ICT projections.

In turn, state ICT projections are used as regressors for WDA ICT projections.

Base projections are benchmarked by the addition of noncovered employment (i.e., not covered by unemployment insurance). Noncovered employment comes from Current Employment Statistics (CES) data. A reconciliation optimization process balances results of different levels of aggregation between regions and the state.

Staffing patterns are created and used to transition industry projections into occupational projections. Occupational openings include openings due to growth and to turnover. Turnover rates, known as replacement rates, measure openings created when workers leave occupations. This year, specifically for Washington state, we created replacement rates using state wage files. These rates give a more realistic measure of actual openings than previous turnover rates.

In addition to projections, we produce additional products:

- Skills estimations and forecasts based on job announcements from Help Wanted Online (HWOL) skills/occupational data.
- The skill estimations are used to create matrices of related occupations based on skills. Such matrices are state specific.
- Occupations in Demand (OID) list. This list is used for determining eligibility for a retraining program (Training Benefits), as well as other education and training programs.

Industry projections

Data

- Covered employment time series
- Global Insight forecasts

Covered employment time series are based on Quarterly Census of Employment and Wages (QCEW) data. For more information see: https://esd.wa.gov/labormarketinfo/quarterly-census.

Global Insight is an international economics organization well known for their data and forecasts.

Software used

The primary software used for forecasting is R-software (R). R is an open source object oriented language with advanced statistical and optimization features. It allows programmers to operate directly on vectors and matrices. This creates significant advantages over languages and software with sequential access to data, like SAS, when producing occupational projections.

Step 1. State level aggregated industry forecast

Data preparation

Initial covered employment at the county level was aggregated into 42 industry groups (cells), presented in the file: allcodes.xlsx. Forty cells were aggregated for nonfarm employment, one for agriculture and one for private households. The cells for nonfarm employment are closely associated with employment related cells from the Global Insight model. However, to meet state employment projections requirements and Occupational Employment Statistic (OES) definition requirements, some cells were disaggregated for state projections. For example, we disaggregated transportation equipment to aerospace and other transportation equipment. The state and local government cell was disaggregated into three cells: government education, hospitals and other government. Two industries related to the information sector were also disaggregated.

We transformed some codes from the Global Insight model in order to match them with codes used in state projections. Due to these transformations, 40 state cells obtained matching relationships with Global Insight national forecasts. Two state cells, agriculture and private households, do not have related national forecasts.

A crosswalk between 4-digit NAICS codes, ICT, aggregated series codes and common combined codes can be found at: allcodes.xlsx. As can be seen in the *allcodes.xlsx* file, aggregated series do not in all cases represent an aggregation of ICT codes. The main reason is that aggregated series reflect commonly used definitions from the CES classification system, while ICT codes reflect industry definitions used in the OES system. To match CES and OES systems, we created combined codes which match aggregated series forecasts with detailed ICT forecasts.

The Global Insight model uses data with quarterly frequencies. In contrast, our historical and forecasted data use monthly frequencies. To make national forecasts usable as regressors for state forecasts, they must be interpolated from quarterly into monthly frequencies. To achieve this we used the **denton-cholette method** from the R-library **tempdisagg**. The **denton-cholette method** uses temporal disaggregation techniques to disaggregate low frequency time series to high frequency series. For an in-depth discussion of disaggregation methods, see: Journal.r-project.org

Parallel processing

When processing large numbers of series, we use R's parallel processing capability. This capability reduces processing time by distributing processes over multiple cores within a computer. The preparation for using parallel processing includes: defining the number of cores in the computer and setting the number of used cores as the number of available cores minus 1. One core must be left to run general computer functions. After the number of cores to be used are defined, core clusters need to be set up and registered with parallel processing functions. R-libraries need to be connected with registered clusters. Parallel processing has some limitations; interactive graphs are not available and failed iterations are not printed in error handling procedures. However, the speed of calculations increases significantly, by about 2.5 times when 3 of 4 available cores are activated.

The main procedure

The main industry projections procedure consists of two parts: 1) importing data for all series; and 2) processing each series. The main library used for data analyses is **dplyr**. Four additional libraries used for the processing of industry forecasts are: **forecast**, **dynlm** (dynamic linear model), **foreach** and **doParallel**. The import of data also involves the defining of data subsets, R-objects and time variables for different time intervals. Objects are held for later use in each series when indexing and cross indexing occur.

For each of the 40 state cells, which have regressors (Global Insight interpolated forecasts), we use the following four types of models:

- Exponential smoothing: innovations state space autoregressive model with an optimized selection of smoothing parameters (criteria: minimum Mean Absolute Percent Error [MAPE]).
- ARIMA: The function **auto.arima** is used to optimize selection of parameters for ARIMA, seasonal ARIMA and periods of seasonality, etc., with regressors (criteria: AIC [Akaike's information criterion]) - this is probably the most sophisticated single equation model available.
- Naive regression model with only seasonal dummies and time (linear trend) as regressors.
- Dynamic linear regression model which includes regressors (the same as for auto ARIMA), seasonal dummies and linear trend.

The exponential smoothing and naive models are autoregressive and only use historical employment time series to forecast employment. The auto ARIMA and dynamic linear regression models can include independent variables (regressors).

The state space method offers a unified approach to a wide range of models and techniques. In general, it includes equations for unobserved states and includes observation equations. Unobserved states (such as level, growth and seasonality) can be subject to change with time. Since the model can account for such changes, it is called **innovative**. The general model can be described as follows:

Let $x_t = (l_t, b_t, s_t, s_{t-1}, ..., s_{t-m+1})'$, be a state vector, where l_t - stands for level; b_t - for growth; and s_t - for seasonality. State space equations can be written in the form:

$$y_t = w(x_{t-1}) + r(x_{t-1})\epsilon_t$$
$$x_t = f(x_{t-1}) + g(x_{t-1})\epsilon_t$$

where ϵ_t is an error term with mean zero and variance δ^2. The equation $\mu_t = w(x_{t-1})$ is a one-step-ahead forecast for the states y_t - observed numbers (employment in our case). Other parameters are defined by the type of model. For instance, models with multiplicative errors use $r(x_{t-1}) = 1$ resulting in $y_t = \mu_t(1 + \epsilon_t)$. Thus, relative errors for multiplicative models are represented by $\epsilon_t = (y_t - \mu_t)/\mu_t$. As can be seen in the state space model, the term “dynamic” refers to states, rather than to observed numbers as in traditional descriptions. For more details about the state space model see: State Space Time Series Analysis.

In R’s *forecast* package, similar state space models for 30 exponential smoothing variations are subject to internal optimization. In our model specifications we chose to allow a damping parameter as a variable. This choice improved the quality of model estimations compared to the use of a default value of one.

Technical details about the models which are used in the *forecast* package can be found at: http://robjhyndman.com/papers/automatic-forecasting/.

The next two types of models are traditional regressions with dynamic, not one-step-ahead, forecasts. The dynamic linear regression model is presented in the form:

$$y_t = c + a * g_t + d * t + s_1 + ... + s_{11} + \epsilon_t$$

where observed employment numbers, y_t, are the linear function of intercept c, endogenous Global Insight forecasts, g, and 11 seasonal dummies, s. If the intercept is not used, there are 12 seasonal dummies. The parameters a and d are scalars and t is any given vector of time. All parameters are estimated by minimizing the square differences.

The naive regression model is the same with the exception of the component related to regressors, $a * g_t$.

For each time series and each model, two forecasts are produced:

- one based on a full sample; and
- one based on a 24-month hold-out sample.

For the full sample forecast, we used all available historical data from January 1990 to June 2016 for parameter estimations. We then forecast for the period from July 2016 to December 2025. Estimations for the hold-out forecast are based on historical data from January 1990 to June 2014 and then forecast from July 2014 to June 2016. As a result of this method, for each time series we have four fittings on a full sample and four hold-out sample forecasts for the following models: innovations state space exponential smoothing, naive, dynamic linear and ARIMA.

We use an optimization procedure to define weights for combining the four full forecasts. The weights are based on the performance (fitting results) of the models on both full sample and hold-out sample forecasts. This year, for the first time, we used mean absolute scale errors (MASE) as a measure of performance.[1] *MASE* is a measure of forecast accuracy proposed by Koehler & Hyndman (2006).

$$MASE = \frac{MAE}{MAE_{in-sample,naive}}$$

where

$$MAE = \frac{\sum_{i=1}^{n} |x_i - \hat{x}_i|}{n}$$

expresses the average absolute difference between each point of time n series x and $\hat{x}$ forecast of x. $MAE_{in-sample,naive}$ is the mean absolute error produced by a naive one-step-ahead forecast, calculated on the in-sample data. We use a one-step-ahead method for seasonal data, which means a 12-month step.

[1]Previously we used mean absolute percent errors (MAPE) as a measure of performance.

$MASE > 1$ implies that an actual forecast does worse than a naive forecast, in terms of mean absolute error. Thus $1 - MASE$ shows the share of variance picked up by a model.

We calculate two mean absolute scaled errors for each of the four models: for full sample fitting $MASE_{full}$ and hold-out sample forecast $MASE_{hold}$.

We define the optimum four weights $z = (z_1, z_2, z_3, z_4)$ for combining forecasts for four model $(i = 1, ..., 4)$ classes, $\sum_{i=1}^{4} for(x_t^i) * z_i$, by solving the problem, find unknown $z = (z_1, ..., z_4)$, for which:

$$MASE_{full} + MASE_{hold} \rightarrow min$$

$MASE$ is applied to the combined forecast.

This combined forecast is called an **optimum forecast**.

For two series without regressors we use the same procedure, but only with three types of models. The naive model and dynamic linear regressions become equivalent and the last is excluded from the process. Also, regressors are excluded from the auto ARIMA model.

Outcomes of the main procedure

The main procedure produces five forecasts for each time series: four models plus a combined optimum forecast. We repeat this procedure for log transformed series and thus potentially have 10 forecasts for each series.[2]

Hierarchy forecast

Hierarchy forecasts were used for the first time this year.

Hierarchical time series forecasting functions are found in R's *hts* package. The *hts* package specializes in forecasting time series that can be disaggregated into hierarchical structures using attributes such as geography. Forecasts are generated for each series at each level of the hierarchy. These forecasts are then combined and balanced by an optimization function within this package. This approach combines independent base forecasts and generates a set of revised forecasts that are as close as possible to the initial univariate forecasts, but also balanced within the hierarchical structure. Hierarchy forecasting was applied to both aggregated series and detailed Industry Control Totals (ICT):

- State series, or ICT, to state total.
- WDA series, or ICT, to state series, or ICT, to state totals.

We used two model options available in the *hts* package, *arima* and *ets*.

For technical details related to hierarchy forecasting see:
Rob J Hyndman and George Athanasopoulos.

Formal adjustments of industry forecasts

Adjustments are applied to all combinations of forecasts and historical data. An adjustment is a useful procedure for smoothing results. We used the concept of **stability controls for dynamic systems** as our smoothing method. The variance of historical employment growth over 12 months[3] was used to define confidence intervals for projected employment variances. We also arbitrarily established the lower and

[2] Estimations for some models can fail for a variety of reasons. The chance for failure increases for unstable series with small numbers involving some zeros. To avoid interruptions in loop processing, for failed series, we use **tryCatch** loops, rather than the default **do** loop. An error handling function prints I.D.'s for all failed series. Also, using the **foreach** loop, rather than the more common **for** loop, allows us to have all of the successful forecasts in output lists as well as identification of all failed series.

[3] Twelve month (or over-the-year) growth rates are used to avoid the impact of stable seasonality.

upper confidence limits at 0.96 and 1.04. These intervals represent the lower number between the historical confidence and the established limit. For each time point, if projected numbers fell within established intervals, they stayed. Otherwise, limits were applied. This process was used as the main mechanism for adjusting models.

Formally the adjustment procedure for each of the series $y_t,\ t = 1, 2, ..., 432$[4] can be described as follows.

Twelve month growth rates calculated as:

$$gr_i = y_t / y_{t-12}, \quad i = 1, ..., 420, \quad t = 13, ..., 432$$

A total of 306 growth rates represent historical data, while another 114 represent forecasted data. We calculate 95% confidence intervals for historical growth rates (*high* and *low*) and average growth rates (*mean*). In this current version of adjustments, we are using only *high* and *mean* values. To make the adjustment formulas more understandable, we introduced two new variables: *adj = high - mean* and *base = max(1,mean)*. Then adjustments to the forecasted growth rates $gr_i,\ i = 307, ..., 420$, are produced by the application of upper and lower limits as follows:

$$gradj_i = gr_i \; if \; gr_i < min(1.04, (base + adj)) \; otherwise \; gradj_i = min(1.04, (base + adj))$$

then

$$gradj_i = gradj_i \; if \; gradj_i > max(0.96, (base - adj)) \quad otherwise$$
$$gradj_i = max(0.96, (base - adj))$$

where, 0.96 and 1.04 are arbitrarily selected numbers and can be subject to change.

The order of applying upper and lower adjustments is irrelevant since values will be unaffected.

Adjusted forecasts are produced by multiplying the last year of available historical data by adjusted growth rates. Then the adjusted forecasts are combined with historical data. Adjustments are applied to each available series, up to 12, resulting in up to 24 forecast options. In this round of projections we did not apply adjustments to state level aggregated series forecasts.

Supplemental parameters used for forecast selections

Important new parameters created in this round of projections are "historical trend growth rates occurring after a major breaking point" for each series. These parameters are *historical trend growth rates.* To define these rates we used R's *BFAST* (Breaks for Additive Season and Trend) package.

The main goal of the package is to integrate the decomposition of time series with methods for detecting and characterizing change within time series. *BFAST* estimates the time and number of abrupt changes within time series. The base decomposition of time series Y_t for time t, from the beginning to the end of a period of interest, is:

$$Y_t = f(S_t, T_t, e_t), \quad where: \; S_t - seasonal, \quad T_t - trend \quad and \quad e_t - remainder$$

For instance, a graph of breaking points for construction employment from January 1990 to June 2016 is presented in *Figure 1.*

[4] Combined series include 432 months (from January 1990 to December 2025).

Figure 1:

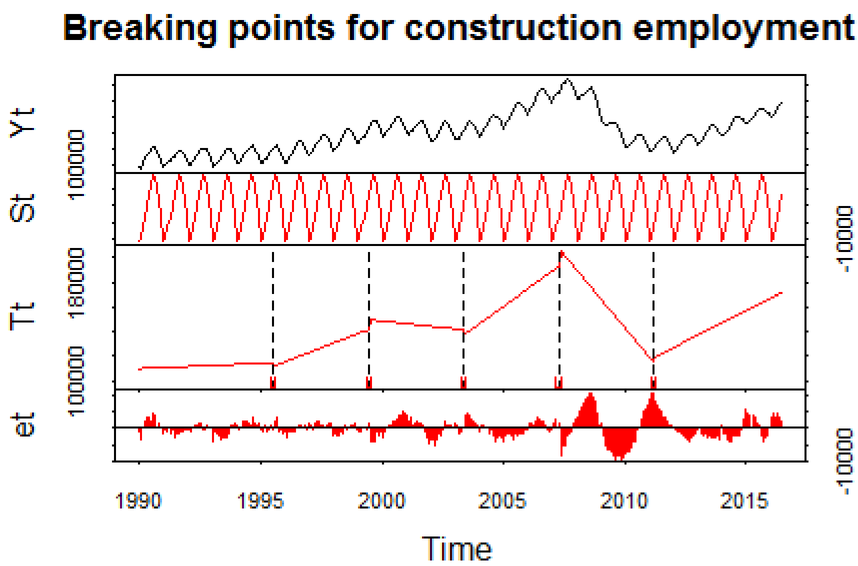

In *Figure 1*, there are five breaking points for the period under consideration: July 1995, June 1999, May 2003, April 2007 and March 2011. The confidence intervals (red marks on the T_t axis) for all breaking points are intervals from one month before and one month after breaking points.

The most significant atypical behavior of this time series is in the interval between April 2007 and March 2011. The remainders on the e_t axis are most significant. They are the largest at the last breaking point.

This construction example gives an idea of how the *BFAST* package can be used for time series evaluation. The package also has the useful function *bfastmonitor* which can be used to monitor the consistency in new data, based on observed evaluated data. Evaluated data can include all available historical data, custom specific intervals and model definitions of largest historical stable intervals. The intervals for evaluation cannot be less than 25 percent of all observed data points.

In this round of projections we used the function *bfast1* to identify one major breaking point for each series. One of the custom control features in this function is the ability to set the minimum share of time points for each of two intervals. We set our share at the level of 0.25. The graph for the same construction employment as in *Figure 1*, but with only one major breaking point is in *Figure 2*.

Figure 2:

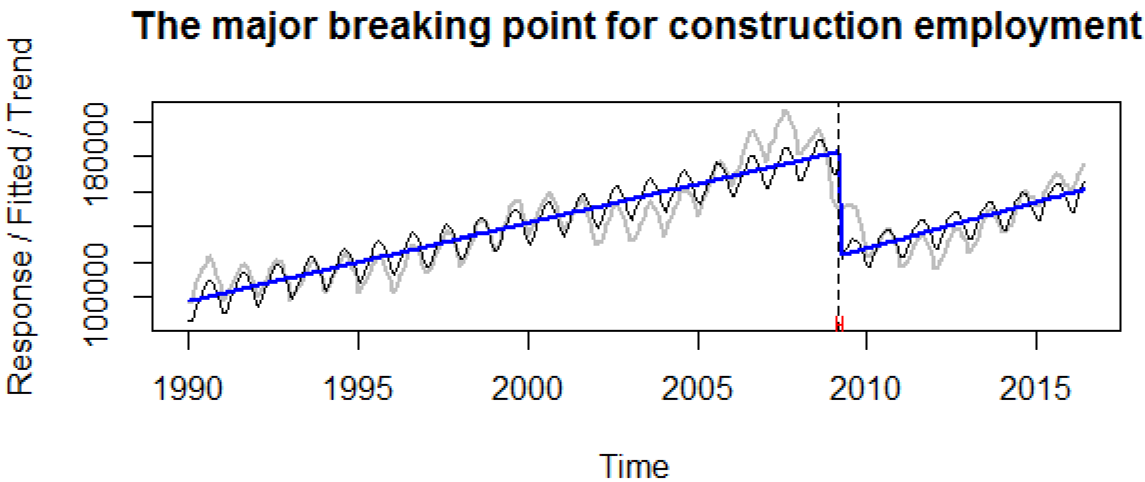

In *Figure 2*, the major breaking point for construction employment occurred in March 2009, with confidence intervals between February and April 2009. By supplementing the output from *bfast1* with the function *bfast01classify*, we can produce annualized growth rates for both intervals (before and after the major breaking point). In addition, *bfast01classify* can create significance levels for fitted models.

In our example, the growth rate on the first interval was 4.2 percent and on the second 3.7 percent. Both estimations have extremely high levels of significance. In our evaluations we mainly used growth rates for the second intervals as long as they had high significance levels.

To evaluate the "smoothness of transition" between historical and forecasted numbers, we calculated the average value for the last three years of changes between June and July and compared the results with the changes between the last month of historical data of June 2016 and the first month of forecast data, July 2016. Any big discrepancies between averaged values and the transition from last historical to a first forecasted value identifies forecasts that are not good candidates for selection.

Selection of aggregated state forecasts

At this stage of the projections process, we select just one of 12 state aggregated series forecasts (formal adjustments are not used for aggregated state series). Selected series are used as regressors in later steps. It is possible that a selected series represents a linear combination of a few forecasts. However, in this round of projections, with only one exception for *private households*, we stayed with just single series selections. This selection process is an *informal process* and is based on calculated average annual growth rates for periods used for the current round of projections. For this round of projections, these periods were: 2016Q2-2018Q2; 2015-2020 and 2020-2025. The growth rates are calculated from aggregated monthly series to proper frequencies (quarterly or annual). The following considerations were used in the informal selection of forecasts:

- historical growth rates for the entire history period and for the last interval after a major breaking point (if significance for the second interval is high);
- the latest aggregated long-term employment forecast from the Office of Financial Management (OFM) and short-term forecast from the Economic and Revenue Forecast Council (ERFC);
- previously published forecasts: our forecasts, OFM and ERFC forecasts;

- smoothness of transitions between the last month of historical data and the first month of forecasted data;
- general knowledge of underlying trends in specific industries; and
- avoidance of extreme growth and decline rates.

Our intention is to select forecasts with growth rates close to those used by OFM and ERFC. We do this unless we have convincing evidence that the OFM and ERFC forecasts are inconsistent between themselves or have significant differences with previously produced results. ERFC forecasts are used for budgetary planning purposes and require the use of adaptive controls. Consequently, forecasts should be updated often to reflect the most current data. In such cases, up-to-date data takes priority over long stable forecasting time periods. Our forecasts are mainly used for career development. Prospective students need forecasts that are stable for medium-range time periods. Frequent updates of forecasts in such cases would be disruptive. In other words, for prospective students, frequently updated forecasts lose practical value.

An example of applying general knowledge of underlying industry trends to specific industries can be seen in our pre-recessionary analyses of productivity trends. Specifically, in the construction industry, our analysis demonstrated that a combination of high employment growth coexisted with a high rate of declining productivity. The declining productivity was compensated for by large price index growth. The combination of high employment growth, low productivity and high prices could not last indefinitely and pointed to a high probability of a downward correction. Therefore, our employment projections for construction growth were more pessimistic. In fact, this type of correction happened during the great recession and created a large drop in construction employment. This drop was the largest among all major industry sectors.

A similar situation occurred in the forecasting of aerospace employment. The delay of Boeing's Dreamliner aircraft, combined with high demand, created an artificial boom in aerospace employment trends. Our projections were more in line with normal aerospace long-term declining cyclical trends. In both cases our declining trends were subject to strong criticisms. Subsequent events affirmed the practice of applying knowledge of underlying trends. The artificial aerospace conditions eventually ceased and declining cyclical trends continued.

Out of 41 forecasts with single selection, the largest number belongs to the base optimum combination - 10 cases, followed by base auto ARIMA - 7 cases and hierarchy ARIMA (*hts-arima*) - 6 cases. The log optimum and *hts-ets* models were selected four times each. In total, combining base and log transformation models among 41 series, optimum models were selected 14 times, while this year's new *hts* model - 10 times. Models from class ARIMA were selected in 17 cases, while models from class ets were selected 8 times. There were 14 selections of the optimum model, but only two selections for regression and naive models (one selection each).

Selected series forecasts are used in the following three independent, but related steps.

Step 2. Draft of state level aggregated benchmarked forecast

Actual historical covered employment numbers for the last 18 months are combined with noncovered employment from the CES program. These numbers are aggregated to two base points used for forecasts resulting in a change of frequencies. The two points are: average annual for the year 2015 and average quarterly for second quarter 2016 (2016Q2). This procedure is called *benchmarking*. Unlike benchmarking by the CES program, we do not use wedging or other adjustments to incorporate code changes. Thus, our benchmark numbers can be slightly different from CES numbers. The growth rates from selected forecasts are applied to benchmarked base numbers. The result is that we produce three required points for industry forecasts: 2018Q2, 2020 and 2025.

The results of benchmarked forecasts are rolled up to create multi-level tables that are somewhat comparable with CES tables. The table is submitted to regional economists, state agencies specified in state law and the Economic Revenue Forecast Council for their feedback.

Step 3. Detailed state level forecasts

For the most part, we repeat all procedures from Step 1's aggregated series for the detailed state level forecasts. We used selected aggregated forecasts as regressors for state detailed forecasts. We use the same formal adjustments and supplemental parameters for selection. However, in the selection process we do not use aggregated external forecasts (ERFC and OFM). Instead, in some cases, we use common combined codes rolled up from aggregated series forecasts.

Eight combined codes are the same as aggregated series and ICT codes and selected aggregated series forecasts were used in such cases. One ICT code for education, 6100, is a combination of two aggregated series: education services and government education services. We combined them to come up with an ICT forecast. The combined ICT directly matches with one combined code. Unlike the previous round of projections, selected ICT forecasts for all ICT codes, other than education, were not formally adjusted to match aggregated series forecasts at the level of combined series. Also, unlike the previous round, actual forecasts from selected models (except one case), rather than weighted forecasts, were used for projections.

Among 284 selected ICT forecasts, the largest number, 114, belongs to the optimum combination, consisting of 78 base models and 36 log transformed (log). The *ets* model was second with 85 (60 base models and 25 log). Arima was third with 56 selections (36 base models and 20 log). The hierarchy forecast was selected 19 times (10 for arima and 9 for *ets*). All hierarchy selections are for the base model since no log transformed hierarchy option was used. The naive model was selected 5 times (4 for base and one for log). The regression model was selected 4 times (3 for base and one for log).

The one ICT code for individual and family services, 6241, was very unstable with a couple of significant breaks (one of them due to code changes). None of the stand-alone models were able to provide satisfactory results for code 6241. For this problem series we selected averages among three models: ARIMA, ets and tbats. Tbats stands for "Exponential smoothing state space model with Box-Cox transformation, ARMA errors, Trend and Seasonal components." It is a fully automated autoregressive model, which includes functions from the forecast package. Tbats was not used in any other projections processes.

Step 4. Local workforce development area (WDA) forecasts

The procedures for producing and formally adjusting local level aggregated and detailed forecasts, in a mathematical sense, are the same as for the state.

We use state aggregated and detailed forecasts from previous steps as regressors for WDA aggregated and detailed forecasts.

Three possible outcomes are possible for each series:

4.1 - All options did not fail and thus we have 24 outputs for each of the series or ICT.

4.2 - All options failed and thus we do not have any forecasts.

4.3 - Some options failed and we have fewer than 24 outputs.

If all options fail,[5] we assign statewide growth rates for those series.

For outcomes 4.1 and 4.3, we use formal adjustment procedures similar to those described in **Step 1**.

Creating weighted WDA forecasts

The purpose of this step is to select for each available forecast, among all available options in each area, a weighted forecast. As in previous steps, we want a forecast that produces growth rates for periods of

[5] Some local (aggregated and detailed) series might not exist (i.e., have zero covered employment). These too can be interpreted as failed series.

interest closest to the ones at the state, i.e., regressor, level. Adjustments in this step are completely formal, conducted in R and exclude interventions.

Let's define g_s, g_m and g_l as short-term (2016Q2-2018Q2), medium-term (2015-2020) and long-term (2020-2025) average annual growth rate employment projections for each state forecast (aggregated or detailed).

Let's also define t ($t = 1, ..., 432$) as a time index for all series with monthly frequencies, and j as an index for forecast options for each available series i. For outcome 4.1, it will be 24 series. The numbers will be fewer if some forecasts failed (outcome 4.3). Let's define n_i as a subset of non failed forecasts for each series i. Then, the optimization problem for each of the available series i can be written as follows.

Find the weights w_j of aggregation for forecast options from the following conditions:

$$0 \leq w_j \leq 1, \quad j = 1, ..., n_i$$

$$\sum_{j=1}^{n_i} w_j(((\sum_{t=340}^{342} y_j^t / \sum_{t=316}^{318} y_j^t)^{0.5} - 1 - g_s)^2 +$$
$$+ ((\sum_{t=349}^{360} y_j^t / \sum_{t=301}^{312} y_j^t)^{0.2} - 1 - g_m)^2 + ((\sum_{t=421}^{432} y_j^t / \sum_{t=361}^{372} y_j^t)^{0.2} - 1 - g_l)^2 \rightarrow min$$

where $y_j = (y_j^1, ..., y_j^{432})$ vectors of employment numbers for option j.

After weights are determined, the weighted forecasts for each series i are simply calculated as:

$$y_t^f = \sum_{j=1}^{n_i} w_j * y_j^t$$

Step 5. Draft of aggregated industry forecasts

Formally adjusted aggregated industry forecasts are benchmarked in the same manner as described in **Step 2**. After benchmarking, numeric discrepancies between the state and WDAs are resolved through an informal process. Since discrepancies are normally very small, due to formal adjustments, for the most part state numbers are slightly modified to meet WDA totals. In some cases the inverse is required and WDA numbers are subject to adjustments to meet state totals.

Minimal informal interventions are possible at this stage of the projections process. Interventions are based on known discrepancies between state and local area trends. In this round of projections, an intervention was applied to electronic shopping and mail-order houses. The extreme state growth rates in this industry are mainly due to historical trends in King County and thus needed to be smoothed. In addition, the smoothing needed to occur since state level numbers are used as regressors for areas.

After adjustments are made, industry tables for each WDA are created in the same manner as in **Step 2** and submitted for internal review by regional economists.

Step 6. Final adjustments and output of industry employment projections

Generally this step is informal and involves processing feedback from state agencies and regional economists. There are generally two types of responses:

1. Responses based on expected, event-based information.
2. Responses on the level of suggestions related to major trends.

Event-based information is related to expected closures and layoffs or expected new hirings due to business expansions or relocations. A very conservative approach is followed and only events with high degrees of certainty are used.

Event-based adjustments are applied, distributed and balanced between aggregated and detailed forecasts. In addition, we incorporate informal interventions for aggregated series from **Step 5**, into detailed series.

Suggestions related to major trends are evaluated based on available data and underlying economic trends and each receive a response. We either provide reasons for rejecting suggestions or inform sources that suggestions will be incorporated into projections. Accepted suggestions are incorporated in the most conservative manner.

The main way to incorporate trend-based suggestions at the state level is by returning back to the **Selection of aggregated state forecasts** in **Step 1** and repeating all subsequent steps for affected industries. Also, it is possible to modify models for affected industries. For suggestions related to local areas, we return back to **Step 4**.

In this round of projections we did not receive any event-based suggestions. We did receive questions, rather than suggestions, related to some major trends. We responded to each question with detailed explanations of forecasting trends (including graphs with breaking points). However, no suggestion-adjustments were made in this round of projections.

After this process is complete and all aggregated and detailed projections are benchmarked, informal adjustments are required to meet the following balancing requirements. These balancing requirements must be met for each of the three projected time periods (2018Q2, 2020 and 2025):

- For each industry, totals for local areas for aggregated and detailed industries should be equal to state numbers.
- For each area, a balance between detailed and aggregated forecasts should be achieved at the aggregated combined series level.

Satisfying the above two conditions leads to a balance between state and local area forecasts at the combined series aggregation level.

While **Step 6** processes may seem complicated, for the most part they are not difficult or time consuming after all automated adjustments have been made. Discrepancies are normally not large and are handled by either a bottom-up approach where state totals are made equal to area totals, or by using a top-down approach with proportional adjustments to local area numbers so that they meet state totals.

For a few series with multiple cross-match-adjustments at the combined series level, the process is more complicated. In this round of projections we mainly used a bottom-up approach for all adjustments: from detailed areas to detailed state and aggregated areas and then to aggregated state.

The industry projections process produces two major outputs:

1. Aggregated industry projections for the state and all areas, which are rounded to the closest 100 and rolled up to create a multi-level table that is somewhat comparable with CES tables (as in **Step 2**).

2. Industry Control Totals (ICT) output for the state and all areas, which are not rounded.

The aggregated projections output is published and used for analyses in projections reports, but is not used for producing occupational projections. The non-rounded ICT output is used in subsequent steps for producing occupational forecasts. ICT industry projection numbers, rounded to integers, are also published.

Occupational projections

Data used

Occupational employment projections result from the conversion of industry employment into occupations. These conversions are based on occupation/industry ratios (i.e., staffing patterns) from the Occupational Employment Statistics (OES) survey. The survey is conducted by the Workforce Information and Technology Services (WITS) division of the Employment Security Department (ESD) in cooperation with the U.S. Bureau of Labor Statistics (BLS). WITS was formerly known as the Labor Market and Performance Analysis (LMPA) division and the Information Technology and Business Integration (ITBI) division. These two divisions have been combined to form WITS.

The full OES survey is conducted over a three-year cycle. One-third of the survey is completed each year. Occupational estimations and projections are subject to the limitations of the OES survey. The survey includes nonfarm employment and agriculture services, but excludes noncovered employment, self-employment and unpaid family members, major agriculture employment (except services), and private households.

The sample for the OES survey is designed for metropolitan statistical areas (MSAs). From the perspective of statistical accuracy for occupational projections, this level of aggregation is the most appropriate. However, for different applications like the Training Benefits Program, we use WDA aggregation levels for regional details. The direct use of OES staffing patterns for WDAs can create significant bias for a variety of reasons.

In this round of projections, the data source for the creation of staffing patterns was almost entirely raw survey data. The majority of data comes from the BLS final (i.e., not preliminary) files. These files include employment, employment distributions by wage intervals, final weights and indicators showing whether original survey responses or imputed responses are used. Response imputations come from other similar in-state areas or from other states. Imputations can have a significant influence on the OES-based staffing patterns.

The process of selecting staffing patterns for each industry and area includes calculating industry totals from raw files. Totals are calculated for weighted employment with imputation and without imputation. Totals are then compared with Industry Control Totals (ICT) for base year periods 2015 and 2016Q2. Our preference is to use data without imputations, but in some cases they do not represent significant shares of employment in ICT output. In such cases, either samples with imputations or substituted staffing patterns are used. These substitutions are mainly introduced using statewide staffing patterns. In some cases, substitutions come from other similar in-state areas. Staffing patterns can create significant bias for industries with high shares of noncovered employment, which are not part of the survey (e.g., religious organizations).

For a few industries, combined staffing patterns were used between areas. This mainly occurred for the King County and Snohomish County WDAs. This was a necessary step because King and Snohomish counties were combined in the OES survey sample. National staffing patterns are used as a last resort and for this year's projection cycle was only necessary for one industry, private households.

Some problems are unavoidable and significantly influence final occupational estimations and projections. For example, doctors can be employed by clinics or hospitals, but often are employees of independent associations or are self-employed. For this reason, staffing patterns for medical institutions are bound to be biased. Also noteworthy this year was the limited use of some results from the 2012 OES green supplemental survey for agricultural industries. The green supplemental survey allowed us to create staffing patterns for agriculture, based on weighted sample responses. This year we also used older survey responses to account for a major employer which has been missing from the latest surveys.

To manage the staffing pattern process, we added two new columns to our ICT files. The new columns indicate the *type* and *area of origination* for staffing patterns. For instance, if an original staffing pattern is used, the area of origination will be the same as the area for industry employment. If an original is not used, the area of origination might be the numeric indicator zero (0) for statewide substitution or the numeric indicator 45 for the combination of King and Snohomish counties, etc.

Occupational projections use national inputs. The inputs are self-employment and unpaid family worker ratios, replacement rates for each occupation and change factors (which we modify).

Step I. Making staffing patterns with selected change factors

The new *type* and *area of origination* columns are used to create base staffing patterns for all areas. *Type* is used to define the source file for selection, while a combination of *area of origination* and *ICT code* is the key for selecting records from source files. The source files are:

- Raw survey data with and without imputations.
- Extracts from the 2012 OES green supplemental survey for agricultural industries.
- National staffing patterns (used only for private households).
- Older raw sample responses for a large missing employer.

Combined indexes in the *area of origination* column are split apart for the selection process of original area codes. Thus, area 45 will be split into a 4 and a 5. For both areas, all available ICT codes are repeated. After selection, combined codes are restored by a simple summarization of numbers for each available industry and occupation.

Each selected vector $a_{i,j}^v$, where $i = 1, .., m$, index for occupations, $j = 1, ..., n$, index for industries and $(v \in V)$ index for areas (can be the number or combination of numbers) is normalized:

$$\sum_{i=1}^{n} a_{i,j}^v = 1, \; for\, all\, industries\, j\, and\, areas\, v.$$

The combination of vectors $a_{i,j}^v$ is often called a *matrix of staffing pattern*. It represents the normalized structures for the distribution of industry employment between occupations for the base period(s). The vectors $a_{i,j}^v$ are matched with extended ICT files by the *area of origination* and industries. After this match, the *area of origination* column is not used for further calculations and thus can be dropped from further formal descriptions. Instead, we now can use the index of actual area v for WDAs from matched files.

We define the index for base and projected periods as $t = 1, .., 5$ and for this round of projections it represents the years 2015, 2016Q2, 2018Q2, 2020 and 2025. The base staffing patterns are used for the years 2015 and 2016Q2 (t=1,2). For other periods, patterns are modified with the incorporation of limited change factors.

Change factors $c_{i,j}$ come from national data. They predict expected changes in occupational shares for each industry over 10 years. The reliability of change factors tends to be low because unlike industry employment, there are no historical time series for occupational employment.

Due to the lack of historical trends upon which to base future changes, BLS uses researchers' expectations about structural occupational changes within industries to create change factors. Within this BLS process, there is a high degree of subjective judgment. This is especially true since change factors must be developed for each occupation within an industry. Occupational outputs are very sensitive to these change factors. It is very important to evaluate the adequacy of change factors before use. Incorrect change factors can drastically increase errors in projections.

We used national change factors in combination with historical state data to create change factors for a limited number of cells. The factors were created only where state historical series were available and were consistent with the suggested change factors from national files. In such cases, we used the most conservative estimations. Change factors reflect expected changes over 10 years, and staffing patterns for projected periods must be modified accordingly:

$$c_{i,j}^2 = (c_{i,j})^{0.2}, \quad c_{i,j}^5 = (c_{i,j})^{0.5}, \quad c_{i,j}^{10} = c_{i,j}.$$

For the two base periods, change factors are not used. The value for all missing change factors can be assumed to be one, and modified staffing patterns are calculated as:

$$a_{i,j}^{v,t} = (c_{i,j})^t * a_{i,j}^v \quad t = 1, .., 5,$$

where $a_{i,j}^{v}$ are the staffing patterns for the base period. Staffing patterns for each period, industry and area need to be normalized to totals of 1.

Step II. Calculation of occupational projections

The results from the previous calculations, for each component $x_{j,v}^{t}$ of the original ICT vectors, in each time period, give as output normalized vectors for occupational distributions $a_{i,j}^{v,t}$. The base occupational employment for each period is simply calculated as:

$$e_{i,t}^{v} = \sum_{j=1}^{n} a_{i,j}^{v,t} * x_{j,v}^{t}, \quad i = 1, .., n, \quad t = 1, .., 5 \quad v = 0, ..., 12$$

due to:

$$\sum_{i=1}^{m} a_{i,j}^{v,t} = 1 \quad for\ each\ j,\ v\ and\ t\ we\ have \sum_{i=1}^{m} e_{i,t}^{v} = \sum_{j=1}^{n} x_{j,v}^{t}.$$

The totals of occupational employment for each area in each point of time equals the totals of industry employment.

The numbers for base period 2016Q2 represent distributions of industry employment between occupations according to normalized staffing patterns. Often, these too are called staffing patterns. These staffing patterns are convenient for publications, but need to be normalized for any calculations outside base periods or with modified ICT outcomes.

Step III. Calculations of self-employment and unpaid family members

Raw self-employment ratios s_i for each occupation come from national data. Based on these ratios, we calculate unadjusted estimated self-employment totals for each area for base year period 2015 as:

$$se_l = \sum_{i=1}^{m} s_i * e_{i,1}^{v}$$

We use estimated numbers of self-employed from the American Community Survey to adjust national self-employment ratios for each area. Let's define the survey numbers for each area as $self_l$. The ratio of adjustment is defined as $ratio_l = self_l/se_l$. The ratio is assumed to be the same for all periods and in this way, adjusted numbers of self-employed for each area v and occupation i are defined as:

$$ase_{i,t}^{v} = se_{i,t}^{v} * ratio_l$$

Then the total of occupational employment is defined as:

$$et_{i,t}^{v} = e_{i,t}^{v} + ase_{i,t}^{v}$$

Step IV. Adding openings due to replacement, separations and alternative state specific rates

Replacement and separation rates come from national data. Until last year, we used only one type of replacement rate, net replacements. Last year we added the national separations rate. Both of these Bureau of Labor Statistics (BLS) methods track openings created when workers leave occupations, but do not track turnover within occupations. Turnover within occupations occurs when workers stay in occupations, but change employers.

In this projections cycle, a new state specific alternative method to the BLS replacement and separations methods was created. The alternative method is based on Washington state wage records, making results state specific.

The alternative rate tracks openings created by turnover within occupations (i.e., workers stay within occupations but transfer to different companies) when workers leave one occupation for another or leave the workforce.

We estimate the numbers of annual transfers between industries, inside industries and in and out of wage files. Then we use occupation-to-industry staffing patterns (shares of occupations for each industry) to convert industry transfers to occupational transfers. Alternative replacement rates are calculated as the shares of total transfers, minus growth or decline, divided by estimated occupational employment for a base period.

From a mathematical point of view, calculations are the same for all three rates. Let's define the rates as r_i. Then the openings due to replacement or separations for each occupation for each period are defined as follows:

$$rep_{i,v} = \frac{et^v_{i,b} + et^v_{i,f}}{2} * r_i,$$

where $et^v_{i,b}$ and $et^v_{i,f}$ are employment totals for the beginning and end of the period. We calculate replacements for periods between 2016Q2 and 2018Q2, 2015 and 2020, and 2020 and 2025.

Step V. Making final outputs

Final outputs include the following results. Calculations are rounded to integers and aggregations to totals for two and three digit SOC levels:

- Total occupational employment estimations $et^v_{i,t}$ for all five periods.
- Average annual growth rates for three periods: 2016Q2-2018Q2, 2015-2020 and 2020-2025.
- Average annual number of openings due to growth $gr_{i,v}$ for each period, which are calculated by subtracting starting points from end points and then dividing the results by the number of years in the period (two or five).
- Average annual openings due to replacements $arep_{i,v}$ calculated by dividing $rep_{i,v}$ by the number of years in the period.
- Total openings due to growth and replacements are calculated as follows:

$$tot_{i,v} = max((arep_{i,v} + gr_{i,v}), 0)$$

We initially round employment estimations and then aggregate them to total two and three digit SOC codes. In this way, results are additive for each column. However, the above formula for calculating total openings introduces non-additivity into the calculations. As a result, the aggregated level of total openings might not equal the total of growth plus replacement.

Some detailed occupational employment estimations are suppressed due to confidentiality. Suppression is introduced after aggregations and normally is not reflected in aggregated results.

General use, employment projection by-products and tools

Employment projections provide a general outlook for industries and occupations in Washington state. Appendices to the main 2017 projections report describe how occupational projections are used as the basis for the Occupations in Demand (OID) list, covering Washington's 12 workforce development areas and the state as a whole. It also describes how we converted occupational projections to skills projections using specific skills extracted from Washington state job announcements. During the creation of skills projections, we produced skills-to-occupation matrices. These matrices allowed us to create a state specific tool useful for matching any given target occupation to related occupations (see *Appendix 3* in the main *2017 Employment Projections report*).

Some of the models and functions (i.e., tools) developed during the production of employment projections could be very applicable to other fields related to time series analyses. For example, R's innovations state space models are more effective for seasonal adjustment than the so called "one level" model, which currently is the most used model for seasonal adjustments. The innovations state space models allow for dynamic changes in seasonal parameters and in this way makes the implementation of additive adjustment factors very effective. In the one level model, such factors lack the ability to reflect trend changes. The use of additive parameters, in conjunction with hierarchy forecasting tools, would allow users to create seasonally adjusted series, balanced between different levels of aggregation, as a cointegrated process. Commonly used seasonal adjustment models do not allow for the direct balancing of different levels of aggregations. Such balancing is normally achieved by a top-down disaggregation or a bottom-up aggregation of seasonally adjusted series.

Useful new capabilities for time series analyses are contained in R's *BFAST* package. This package makes possible a better understanding of historical trends and the impacts of specific events, like recessions, on such trends. The *BFAST* package could be very effective for identifying pro-cyclical and counter-cyclical industries. One of the most useful features of *BFAST* is its ability to monitor the consistency of new data, based on observed evaluated data. Evaluated data could include all available historical data, custom specific intervals or the largest historical stable intervals defined by models. This package could be used, for instance, for evaluating the typical or atypical behaviors of Current Employment Statistic (CES) samples or employers' reports that are processed within the unemployment insurance system. The use of *BFAST*'s automated tools could significantly increase the speed, quality and consistency of analyses within any organization's processes.

References

Commandeur, J. F. J. and Koopman, S. J. (2007). *An Introduction to State Space Time Series Analysis*. Oxford University Press, Oxford.

Hyndman, R. J. and Khandakar, Y. (2008). Automatic time series forecasting: the forecast package in R. *Journal of Statistical Software* 27(3).

Hyndman, R. J. and Athanasopoulos, G. (2013). Forecasting: principles and practice. OTexts: Melbourne, Australia, Section 9.4, https://www.otexts.org/fpp/9/4. Accessed on 06/26/2017.

Quarterly Census of Employment and Wages. https://esd.wa.gov/labormarketinfo/quarterly-census

Sax, C. and Steiner, P. (2013). Temporal Disaggregation of Time Series, Contributed Research Article, https://journal.r-project.org/archive/2013-2/sax-steiner.pdf. Accessed on 07/06/2017.