FORECASTING HOUSEHOLD CONSUMPTION COMPONENTS: A FORECAST COMBINATION APPROACH

Angelia L. Grant, Liyi Pan, Tim Pidhirnyj, Heather Ruberl and Luke Willard[1]

Treasury Working Paper[2]

2017-04

Date created: June 2017

Date modified: June 2017

1 All authors work in the Australian Treasury. The authors would like to thank Joshua Chan, Mardi Dungey, Linus Gustafsson, Alexandra Heath, Nigel Ray, Penny Smith and Warren Tease for helpful comments. Correspondence to: Australian Treasury, Langton Crescent, Parkes ACT 2600, Australia. Email: angelia.grant@treasury.gov.au.

2 The views expressed in this paper are those of the authors and do not necessarily reflect those of The Australian Treasury or the Australian Government.

ISBN 978-1-925504-52-1

Manager
Media Unit
The Treasury
Langton Crescent
Parkes ACT 2600
Email: medialiaison@treasury.gov.au

Forecasting Household Consumption Components: A Forecast Combination Approach

Angelia L. Grant, Liyi Pan, Tim Pidhirnyj, Heather Ruberl and Luke Willard*

June 2017

Abstract

This paper outlines a methodology for forecasting the components of household final consumption expenditure, which is necessary in order to forecast revenue collections from a number of different taxes. A forecast combination approach using autoregressive models, regressions on relative prices and the almost ideal demand system developed by Deaton and Muellbauer (1980) is found to offer a more robust forecasting framework than using one of the single models alone. In particular, the combination approach outperforms the almost ideal demand system, which is currently used by the Australian Treasury to forecast the components of consumption. The combination framework takes advantage of models that account for the persistence and longer-term trends experienced in a number of the consumption components, as well as shifts caused by evident relative price changes. A forecast combination framework is shown to be particularly useful when forecasting over a three-year forecasting period.

Keywords: Household consumption expenditure, forecast combination.

---

*All authors work in the Australian Treasury. The views expressed in the paper are those of the authors and do not necessarily reflect those of the Treasury or the Australian Government. The authors would like to thank Joshua Chan, Mardi Dungey, Linus Gustafsson, Alexandra Heath, Nigel Ray, Penny Smith and Warren Tease for helpful comments. Correspondence to: Australian Treasury, Langton Crescent, Parkes ACT 2600, Australia. Email: angelia.grant@treasury.gov.au.

1 Introduction

Forecasts for each of the expenditure components of nominal GDP are important for forecasting tax revenue collections—different compositions result in different tax revenue forecasts. A particularly important task is the forecasting of the components of household final consumption expenditure. This is because different components of consumption are subject to different taxes. For example, alcohol, tobacco and fuel are subject to excise taxes, while motor vehicles may be subject to the luxury car tax. A number of the components of household final consumption expenditure—durables, other goods, electricity and gas, and other services—are also subject to the goods and services tax.

A wide variety of models can be used to forecast the components of household consumption, with different models using different types of information. Some models are good at accounting for the persistence and longer-term trends experienced in a number of the consumption components, while other models are better at taking into account shifts caused by relative price changes. It is also the case that some models are better at forecasting over shorter time horizons, while others are better over longer time horizons.

Under these circumstances, a forecast combination approach has a number of advantages. It allows the use of information across a number of models and the use of models that perform differently across different time horizons. It is often the case that when forecasts from a variety of different models are appropriately combined, the forecast combination approach outperforms individual forecasts (see, e.g., Timmermann, 2006; Guidolin and Timmermann, 2009; Rapach et al., 2010).

This paper develops a forecast combination approach for the components of household consumption expenditure using autoregressive models, regressions on relative prices and the almost ideal demand system developed by Deaton and Muellbauer (1980).[1] The autoregressive models capture the persistence and longer-term trends in the consumption components, while the relative price regressions and the almost ideal system capture shifts in consumption components that are driven by relative price changes. At shorter forecasting horizons, models that capture short-run dynamics perform well, while at longer horizons models with trend terms and relative prices generally perform better based on root mean squared forecast errors.

[1] There are, of course, other factors that might affect household consumption, such as changes in tax policy, income uncertainty and changes in wealth.

Two forecast combinations are constructed—one based on equal weights and the other weighted based on forecasting performance according to rolling squared forecast errors.[2] The advantage of combining forecasts based on past forecast performance is that the forecast combination is robust to changes in modelling performance. That is, it accounts for the fact that certain models can improve or diminish in performance over particular time periods. However, it is also often found that equal weights perform strongly (see, e.g., Timmermann, 2006), so both approaches are considered in this paper.

The forecast combinations generally perform better than the almost ideal demand system, which is the model currently used for estimating the household final consumption components. The fuels and lubricants and the electricity and gas components are the components where the forecast combination performance is closest to that of the almost ideal demand system. The forecast combination based on past forecast performance performs better than the equal weights model for all components.

The remainder of this paper is organised as follows. Section 2 provides a brief description of each of the forecasting models used in the forecast combination framework, Section 3 discusses the forecast combination framework and Section 4 details the data used in the analysis. Section 5 reports the out-of-sample forecasting results for each of the models and for the forecast combination, and Section 6 concludes.

2 Individual Forecasting Models

This section outlines each of the models that are used in the forecast combination framework. The models are chosen *a priori* to capture the persistence and longer-term trends in the household consumption components and shifts in the consumption components that occur as a result of changes in relative prices.

The almost ideal demand system is the model currently used for estimating the components of household final consumption expenditure. This model forecasts consumption shares taking relative prices and an income term as inputs. For easy comparison between models, all are used to forecast consumption shares.

[2] For density forecasts of Australian output growth, inflation and interest rates, Gerard and Nimark (2008) use the predictive likelihood for combining forecasts from different vector autoregression models.

2.1 Autoregressive Models

The first models considered are autoregressive models. These models take into account the persistence of the shares of consumption components, with the share of consumption on a particular component modelled to be a linear combination of past shares. The implicit assumption within autoregressive models is that consumption patterns tend to depend on those in recent periods, consistent with habit-forming preferences.

The set of autoregressive models for $i = 1, \ldots, n$ consumption shares is as follows:

$$c_{i,t} = \beta_{i,0} + \beta_{i,1}t + \sum_{k=2}^{K+1} \beta_{i,k}\, c_{i,t-k+1} + \varepsilon^{a}_{i,t}, \tag{1}$$

where $c_{i,t}$ is the share of current price consumption on good i in time period t, t is a linear time trend, K is the number of autoregressive terms and $\varepsilon^{a}_{i,t}$ is assumed to be independent and identically normally distributed.

The time trend is included in the models in order to account for the fact that some of the consumption shares may be trend stationary. A version of the models is also considered without the time trend. This version is appropriate for stationary consumption shares.

The main advantage of autoregressive models is that they perform well at modelling persistence. On the other hand, the disadvantage is that they do not use any other information, such as relative price movements. Each model is estimated using ordinary least squares regressions. The forecast consumption shares are normalised to sum to 1.

2.2 Regressions on Relative Prices

The autoregressive models do not account for changes in relative prices, which can drive important shifts in the share of each consumption component. As such, the next models considered are regressions on the relative price of the consumption component.

The set of relative price regressions for $i = 1, \ldots, n$ consumption shares is as follows:

$$c_{i,t} = \delta_{i,0} + \delta_{i,1}t + \delta_{i,2}\log\left(\frac{p_{i,t}}{p_t}\right) + \varepsilon^{r}_{i,t}, \tag{2}$$

where $c_{i,t}$ is the share of current price consumption on good i in time period t, t is a linear time trend, $p_{i,t}$ is the price of consumption good i, p_t is the aggregate consumption price and $\varepsilon^r_{i,t}$ is assumed to be independent and identically normally distributed.

As for the autoregressive models, the time trend is included to capture the fact that some consumption shares may be trend stationary. A version of the models is also considered without the time trend. This version is appropriate for stationary consumption shares.

The main advantage of relative price regressions is that they take into account information about relative price shifts. A disadvantage is that they do not include dynamics in the form of past consumption shares. Each of these regressions is estimated using ordinary least squares. The forecast consumption shares are normalised to sum to 1.

2.3 Almost Ideal Demand System

The almost ideal demand system (AIDS) of Deaton and Muellbauer (1980) takes into account that the consumption of a particular good or service depends not only on its own price, but also the relative prices of other goods and services which may be either complements or substitutes. It also takes into account an income effect, with each of the shares depending on total consumption expenditure.

The AIDS demand functions for $i = 1, \ldots, n$ consumption shares are as follows:

$$c_{i,t} = \omega_{i,0} + \sum_{j=1}^{n-1} \omega_{i,j} \log\left(\frac{p_{j,t}}{p_{n,t}}\right) + \gamma_i \log\left(\frac{x_t}{p_t}\right) + \varepsilon^d_{i,t}, \tag{3}$$

where $c_{i,t}$ is the share of current price consumption on good i in time period t, $p_{j,t}$ is the price of consumption good j, $p_{n,t}$ is the price of the residual consumption good category n, x_t is the value of total consumption expenditure and p_t is the aggregate consumption price. As in the above models, $\varepsilon^d_{i,t}$ is assumed to be independent and identically normally distributed. The equations are estimated as a system and the other services component is treated as a residual to ensure the shares add to 1.

The system satisfies the homogeneity restriction and symmetry is imposed by restricting $\omega_{i,j} = \omega_{j,i}$. It is important to note that Deaton and Muellbauer (1980) test the restrictions of homogeneity and symmetry using postwar British data and find that both are decisively rejected. It is also found that the imposition of homogeneity generates positive serial correlation in the errors of those equations which reject the restrictions most strongly. For this reason, Deaton and Muellbauer (1980) conclude that the system is not “a fully satisfactory explanation of consumers' behaviour."

The main advantage of the almost ideal demand system is its strong theoretical grounding. But this strong theoretical grounding may mean that the model may be too restricted to fit the data well. In addition, the model has a large number of parameters. Given the large number of parameters and the sample period available for the data, the model is estimated without the income term (i.e. $\gamma_i \log(x_t/p_t)$).[3]

3 Forecast Combination Approach

There are two methods used to construct the forecast combinations. The first method uses equal weights. In this case, each of the $m = 1, \ldots, M$ models for each of the household consumption components $i = 1, \ldots, n$ is given a weight of $1/M$. The simple combination approach is often found to outperform other more sophisticated combination schemes.

The second method weights the forecasts using past forecast performance. More specifically, each of the $m = 1, \ldots, M$ models for each of the household consumption components $i = 1, \ldots, n$ is weighted using a J-quarter rolling weight of the inverse of the sum of the squared forecast error. The weight of model k for consumption component i at time t is:

$$w_{k,i,t} = \frac{\left(1/\sum_{j=1}^{J} e^2_{k,i,t-j}\right)}{\sum_{m=1}^{M}\left(1/\sum_{j=1}^{J} e^2_{m,i,t-j}\right)}, \tag{4}$$

where $e^2_{m,i,t}$ is the squared forecast error of model m for consumption component i at time t. The size of the rolling window J is set to 4 so that the weights depend on the performance of the past four forecast periods, with the first four periods equally weighted.[4] This strikes a balance between having relatively stable weights and weights that quickly adapt when there is a change in performance across models.

[3] The estimation without an income term undermines one of the main advantages of using a theoretical model. Given the relatively small number of observations, adding the income term would substantially increase parameter uncertainty.

[4] Using this approach, the forecast household consumption shares do not exactly add up to one. While a final normalisation of the weights could be undertaken, it makes only a small difference. The forecast shares are unchanged up to six decimal places.

The approach of combining forecasts based on past forecast performance accounts for changes in modelling performance. That is, it captures the benefits of different modelling approaches and accounts for the fact that certain models can improve or diminish in performance over particular time periods and at different forecasting horizons.

4 Data

The household consumption data are (unpublished) disaggregated series from the quarterly National Accounts, which are sourced from the Australian Bureau of Statistics. The frequency of the data is quarterly and all series are seasonally adjusted. There are nine components considered: food; alcohol; cigarettes and tobacco; durables; other goods; vehicles; fuels and lubricants; electricity and gas; and other services. The nominal shares are calculated as a share of total non-rental consumption. The data are shown in Figure 1.

The prices for each of the components are constructed from the current price and chain volume series. The full sample period is 1986Q1 to 2016Q1 and the last vintage of the full dataset is used in the forecasting exercise. The full sample period is chosen so as to accommodate the estimation of the almost ideal demand system, which has a large number of parameters.

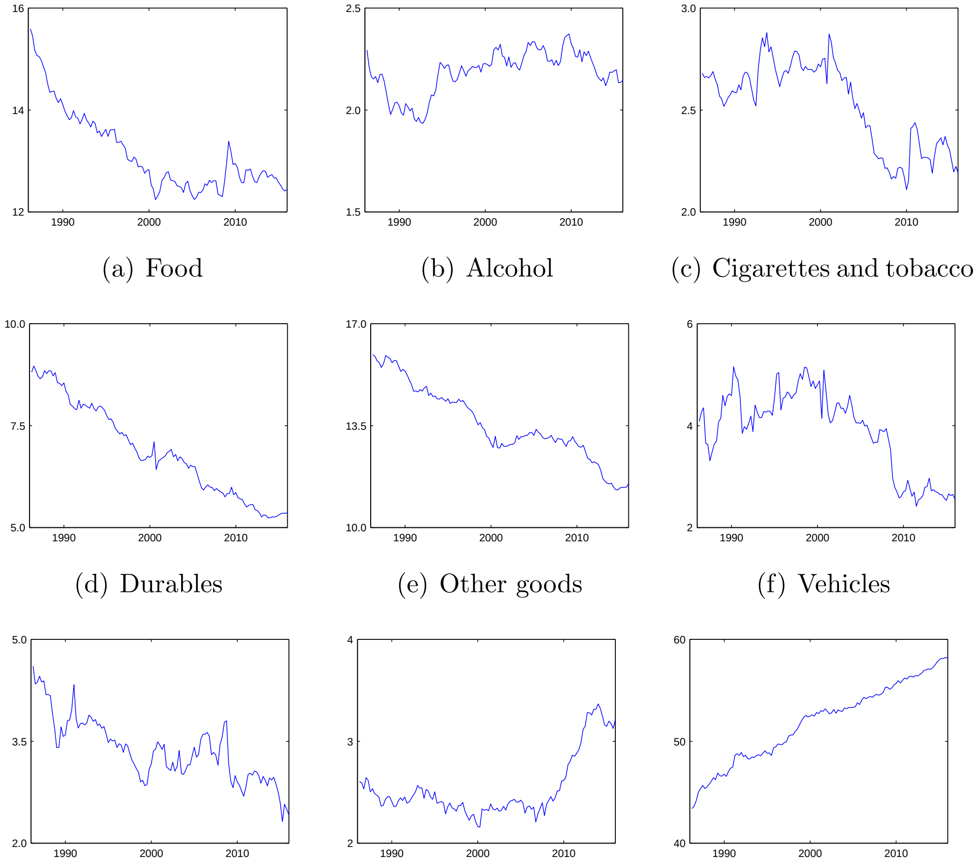

(a) Food (b) Alcohol (c) Cigarettes and tobacco

(d) Durables (e) Other goods (f) Vehicles

(g) Fuels and lubricants (h) Electricity and gas (i) Other services

Figure 1: Shares of household final consumption expenditure components.

5 Forecasting Results

This section reports the out-of-sample forecasting results for each of the models and for the forecast combinations. The out-of-sample forecasting is based on 12-step ahead forecasts, which are computed every 2 quarters. The forecast evaluation period is from 2006Q1 to 2016Q1. In assessing the forecast performance, the benchmark model is the almost ideal demand system given that it is the model currently used for forecasting the household final consumption components.[5]

[5] Formal statistical tests could be performed to assess the statistical significance of the results. However, given the relatively short evaluation period it would be difficult to obtain conclusive results.

5.1 Individual Models

Five models are estimated for each individual component of household final consumption expenditure. They are the standard AR(2) models, AR(2) models with linear time trends, relative price models, relative price models with linear time trends and the almost ideal demand system (AIDS). The lag length for the autoregressive models is chosen with a view to modelling the persistence in the data, while maintaining a parsimonious specification. The estimated AIDS does not include total consumption expenditure to ensure that the model is not over-parameterised.

Table 1 reports the root mean squared forecast error (RMSFE) relative to that of the almost ideal demand system for each of the models over different forecast horizons. At the one-quarter-ahead forecasting horizon, both sets of AR(2) models significantly outperform the almost ideal demand system for all household consumption components. These models perform particularly well for the components of food, cigarettes and tobacco, durables, other goods and other services. For example, the RMSFE for other goods under the standard AR(2) model is only 17 per cent of that of the almost ideal demand system. The forecasting gains are smaller for the components of fuels and lubricants and electricity and gas, but they continue to be better than the benchmark.

In the case of the relative price models, the model with the linear time trend generally performs much better than the model without the trend. Further, even for the components where the model with the linear time trend does not outperform the almost ideal demand system—alcohol and fuels and lubricants—the performance is not substantially worse than AIDS. The relative prices model with the linear time trend does particularly well at forecasting other services, with the RMSFE being only 17 per cent of that of the almost ideal demand system. The RMSFE for other services under the relative prices model without a time trend is 57 per cent of that of the almost ideal demand system.

As the forecasting horizon increases, the performance of the AR(2) models for some components deteriorates relative to the AIDS. For example, at the three-year-ahead horizon, the forecast for the fuels and lubricants component of the benchmark model is 53 per cent (i.e. $1 - 1/2.14$) better compared to the standard AR(2) model. In contrast, the relative price model with the linear time trend continues to perform relatively well for most of the components. Consequently, it can generally be concluded that, at shorter forecasting time horizons, models that capture short-run dynamics perform well, but that, at longer horizons, models with trend terms and relative prices tend to perform better.

Table 1: Root mean squared forecast error relative to almost ideal demand system (values less than 1 indicate better forecast performance than the benchmark).

	Household consumption components*								
	(a)	**(b)**	**(c)**	**(d)**	**(e)**	**(f)**	**(g)**	**(h)**	**(i)**
One-quarter-ahead forecast									
AR(2) model	0.21	0.40	0.23	0.23	0.17	0.38	0.82	0.76	0.21
AR(2) model with trend	0.22	0.45	0.22	0.21	0.16	0.37	0.74	0.71	0.16
Relative price model	0.81	0.85	1.30	1.25	1.43	0.97	3.71	1.11	0.57
Relative price model with trend	0.75	1.08	0.32	0.33	0.45	0.81	1.06	0.80	0.17
One-year-ahead forecast									
AR(2) model	0.37	0.87	0.72	0.29	0.34	0.81	1.77	1.56	0.33
AR(2) model with trend	0.45	1.09	0.70	0.30	0.32	0.78	1.58	1.44	0.12
Relative price model	0.66	0.94	1.16	1.16	1.42	0.96	3.36	1.10	0.44
Relative price model with trend	0.75	1.30	0.38	0.34	0.54	0.83	1.04	0.89	0.14
Two-year-ahead forecast									
AR(2) model	0.26	0.90	0.89	0.27	0.54	1.00	1.83	2.58	0.51
AR(2) model with trend	0.43	1.33	0.89	0.34	0.49	0.95	1.33	2.45	0.17
Relative price model	0.46	0.88	1.10	1.11	1.45	0.96	3.23	1.08	0.33
Relative price model with trend	0.75	1.36	0.40	0.32	0.65	0.83	0.96	1.01	0.11
Three-year-ahead forecast									
AR(2) model	0.24	0.86	1.06	0.31	0.68	1.12	2.14	3.20	0.61
AR(2) model with trend	0.47	1.56	1.17	0.40	0.57	1.06	1.04	3.22	0.17
Relative price model	0.18	0.97	1.07	1.03	1.40	0.96	3.16	1.11	0.15
Relative price model with trend	0.74	1.63	0.78	0.34	0.65	0.85	0.87	1.14	0.10

* The labelling corresponds with Figure 1: (a) food; (b) alcohol; (c) cigarettes and tobacco; (d) durables; (e) other goods; (f) vehicles; (g) fuels and lubricants; (h) electricity and gas; and (i) other services.

5.2 Forecast Combinations

This section reports the forecasting results for the two forecast combinations—the combination based on equal weights across all of the models, and the combination with weights proportional to the sum of squared forecast errors (SSFE) over the past four quarters. Table 2 reports the root mean squared forecast error relative to that of the almost ideal demand system for both of the forecast combinations over different forecast horizons.

Table 2: Root mean squared forecast error relative to almost ideal demand system (values less than 1 indicate better forecast performance than the benchmark).

	Household consumption components*								
	(a)	**(b)**	**(c)**	**(d)**	**(e)**	**(f)**	**(g)**	**(h)**	**(i)**
One-quarter-ahead forecast									
Forecast combination, equal weights	0.38	0.49	0.53	0.48	0.52	0.66	1.15	0.82	0.25
Forecast combination, SSFE weights	0.22	0.45	0.45	0.24	0.36	0.52	0.91	0.80	0.19
One-year-ahead forecast									
Forecast combination, equal weights	0.45	0.76	0.67	0.53	0.55	0.84	1.34	1.13	0.20
Forecast combination, SSFE weights	0.37	0.82	0.64	0.31	0.43	0.83	1.05	1.03	0.12
Two-year-ahead forecast									
Forecast combination, equal weights	0.44	0.76	0.72	0.55	0.59	0.93	1.27	1.56	0.17
Forecast combination, SSFE weights	0.32	0.74	0.70	0.37	0.56	0.92	0.92	1.19	0.16
Three-year-ahead forecast									
Forecast combination, equal weights	0.44	0.81	0.75	0.56	0.58	0.99	1.30	1.84	0.12
Forecast combination, SSFE weights	0.30	0.65	0.73	0.42	0.61	0.97	0.99	1.29	0.12

* The labelling corresponds with Figure 1: (a) food; (b) alcohol; (c) cigarettes and tobacco; (d) durables; (e) other goods; (f) vehicles; (g) fuels and lubricants; (h) electricity and gas; and (i) other services.

At the one-quarter-ahead forecasting horizon, both forecast combinations perform better than the almost ideal demand system, with the exception of the equal weight model for fuels and lubricants. The forecast combinations perform particularly well for the food and other services components. In the case of other services, the root mean squared forecast error for the equal weight model is only 25 per cent of that of the almost ideal demand system and 19 per cent for the model weighted by forecast performance. This is an important result given that the other services category accounts for a large share of consumption subject to the goods and services tax.

It is also the case that, at the one-quarter-ahead forecasting horizon, the forecast combination based on past forecast performance performs better than the equal weighted model. The performance of the forecast combinations does not, however, outperform the standard AR(2) models or the AR(2) models with linear time trends. This reinforces the conclusion that models that capture short-run dynamics perform well at shorter forecasting time horizons.

Figure 2 shows the model weights for the one-quarter-ahead forecast for the food component. For most periods, each of the models have a non-negligible weight in the forecast combination based on past forecast performance. In other words, all models are contributing to produce the final forecast. In addition, the weights are evolving over time. At the beginning of the forecast period, the AR(2) model and the AR(2) model with a linear time trend perform well and account for most of the weight. Over time the weight of the relative prices model increases, illustrating that the forecast performance of this model improves towards the end of the forecast period. As such, combining the forecasts from all of the models using time-varying weights means that the forecast combination can quickly adapt to changes in model performance.

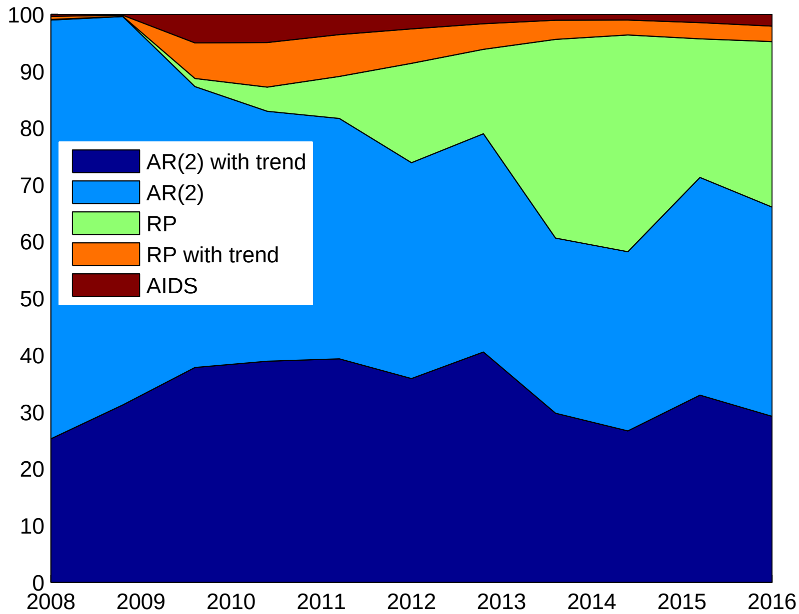

Figure 2: Weights for the one-quarter-ahead forecast for the food component.

At the one-year-, two-year- and three-year-ahead forecasting horizons, the forecast combinations also generally perform better than the almost ideal demand system, with fuels and lubricants and electricity and gas being the only components where the forecasting performance is not uniformly better than that of the almost ideal demand system. The forecast combination using past forecasting performance uniformly outperforms the combination based on equal weights at the longer forecasting horizons.

At the longer forecasting horizons, the forecast combinations perform significantly better than the autoregressive models. In contrast, at the three-year-ahead forecasting horizon, the relative prices model with linear time trends performs better than the forecast combination based on past forecasting performance for five out of the nine household consumption components. This shows that models with trend terms and relative prices tend to perform better over longer forecasting horizons, while the autoregressive models are better at forecasting over shorter time horizons.

The varied forecasting performance across the different individual models for the different components of household consumption expenditure and across different forecasting time horizons highlights the benefit of a forecast combination framework. The forecast combination based on forecasting performance takes advantage of models that account for the persistence and longer-term trends in a number of the consumption components, and the shifts caused by relative price changes. Moreover, as a model outperforms its competitors in the recent past, a higher weight is given to that successful model. In this way, the forecast combination approach quickly adapts to changes in model performance. A forecast combination framework is particularly useful when it is necessary to forecast over a three-year forecasting period.

6 Concluding Remarks

This paper outlines a methodology for forecasting the components of household final consumption expenditure. It uses a forecast combination approach with autoregressive models, regressions on relative prices and the almost ideal demand system developed by Deaton and Muellbauer (1980). The forecast combination that weights the forecasts based on forecasting performance according to rolling squared forecast errors generally performs better than the currently-used almost ideal demand system. The forecast combination takes advantage of the forecasting performance across the different individual models for the different components of consumption expenditure and across different forecasting horizons. The forecast combination is particularly useful when it is necessary to forecast over a three-year forecasting period, given significant differences in forecasting performance of models across different forecasting horizons.

References

A. Deaton and J. Muellbauer. An almost ideal demand system. *American Economic Review*, 70(3):312–326, 1980.

H. Gerard and K. Nimark. Combining multivariate density forecasts using predictive criteria. *RBA Discussion Paper 2008-02*, 2008.

M. Guidolin and A. Timmermann. Forecasts of US short-term interest rates: A flexible forecast combination approach. *Journal of Econometrics*, 150(2):297–311, 2009.

D.E. Rapach, J.K. Strauss, and G. Zhou. Out-of-sample equity premium prediction: Combination forecasts and links to the real economy. *Review of Financial Studies*, 23 (2):821–862, 2010.

A. Timmermann. Forecast combination. *Handbook of Economic Forecasting*, edited by: G. Elliott, C.W.J. Granger and A. Timmermann, North-Holland, 2006.