

NBER WORKING PAPER SERIES

MARKOV FORECASTING METHODS
FOR WELFARE CASELOADS

Jeffrey Grogger

Working Paper 11682
<http://www.nber.org/papers/w11682>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2005

I thank Curt Eberwein for helpful comments and the USDA and the Joint Center for Poverty Research for financial support. Opinions expressed herein are my own and do not necessarily represent the views of USDA or JCPR. Any errors are my own. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

©2005 by Jeffrey Grogger. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Markov Forecasting Methods for Welfare Caseloads
Jeffrey Grogger
NBER Working Paper No. 11682
October 2005
JEL No. I3

ABSTRACT

Forecasting welfare caseloads, particularly turning points, has become more important than ever. Since welfare reform, welfare has been funded via a block grant, which means that unforeseen changes in caseloads can have important fiscal implications for states. In this paper I develop forecasts based on the theory of Markov chains. Since today's caseload is a function of the past caseload, the caseload exhibits inertia. The method exploits that inertia, basing forecasts of the future caseload on past functions of entry and exit rates. In an application to California welfare data, the method accurately predicted the late-2003 turning point roughly one year in advance.

Jeffrey Grogger
Harris School of Public Policy
University of Chicago
1155 E. 60th Street
Chicago, IL 60637
and NBER
jgrogger@uchicago.edu

I. Introduction

Forecasting welfare caseloads has become more important than ever. One reason is the magnitude of recent fluctuations. Figure 1 exhibits the number of welfare cases in California between July 1985 and March 2005. The general patterns there are similar to those in other states. After growing slowly during the late 1980s, caseload growth accelerated starting around 1989. When the caseload peaked in March 1995 (about a year later than the national caseload peak), almost 915,000 households were receiving aid under the Aid to Families with Dependent Children (AFDC) program. In the late 1990s the caseload plummeted, reaching a low of just under 445,000 cases in November 2003. Since then it has grown a bit, standing at nearly 463,000 cases as of March 2005.

Whether the recent rebound represents a true turning point, foretelling further growth in the caseload, or a relatively minor fluctuation about a new, stable level, is a matter of keen interest for both forecasters and program administrators. Turning points are particularly difficult to forecast, and public aid caseloads represent no exception to this rule. Dynarski et al. (1991) attempted to forecast Food Stamp caseloads, which fluctuated in a manner similar to the welfare caseload during the 1980s and 1990s. Despite their extensive modeling effort, Dynarski et al. concluded that their model did not yield "highly accurate" forecasts, and that "none of the ... models would have captured the increase in participation that began in 1989" (p. xi).

From the viewpoint of program administrators, turning points impose a particular burden. Unforeseen swings in the caseload complicate the appropriations process. They may also cause logistical problems for the administration of eligibility and service delivery operations.

Perhaps most importantly, recent changes in the funding of welfare programs mean that substantial fluctuations in welfare caseloads could have broader implications for state spending and tax policy. The Personal Responsibility and Work Opportunity Restoration Act (PRWORA) replaced AFDC with Temporary Assistance for Needy Families (TANF) and changed the funding formula used to finance welfare expenditures. Whereas AFDC expenditures were shared between federal and state governments, TANF is funded with a block grant. This means that sufficient growth in TANF cases could force states to curtail eligibility, reduce benefits, reduce spending on other state programs, or raise taxes to support higher spending. Although the threat of such welfare-driven fiscal dislocations is currently low due to the high funding level of the TANF block grant, recent proposals to reduce the federal block grant could present states with difficult choices if caseloads were to rise substantially.

The purpose of this paper is to introduce a new method for forecasting welfare caseloads. I refer to the technique as Markov forecasting because it is motivated by results from the theory of Markov chains. I apply the technique to caseload data from California. It provides forecasts of the caseload a little more than one year in advance. Most importantly, it predicts turning points reasonably well.

It differs in important ways from the two traditional approaches to forecasting, known as time-series forecasting and econometric forecasting. Time-series forecasting treats future caseloads as a function of past caseloads and past forecasting errors (Box and Jenkins, 1970). Its advantage is that it can be applied relatively quickly and cheaply to any time series. Its disadvantage is that it takes no special account of the nature of the data, and thus represents something of a black box. The traditional alternative,

econometric forecasting, models the caseload as a function of variables thought to influence it, such as economic conditions and the policy environment (Granger and Newbold 1977). Its advantage is that it shows why the caseload changes. Its disadvantages are that such models can be costly to construct and that their forecasts require forecasts of the independent variables thought to affect the caseload.

Markov forecasting is based on a model of caseload evolution. In the simplest terms, today's caseload depends on yesterday's caseload plus entries and exits. Because today's caseload depends in part on yesterday's caseload, the caseload exhibits inertia. Markov forecasting exploits that inertia to base forecasts of the future caseload on current entries and exits. Put differently, current entries and exits affect current caseloads, but since today's caseload affects future caseloads, current entries and exits affect future caseloads as well. If turnover is not too great, then there is a lag between changes in entries and exits and the time that those changes are fully reflected in the caseload. Thus current trends in entries and exits help predict future caseloads.

Markov theory motivates a leading indicator that is based on entries and exits. The leading indicator is the key to Markov forecasting. A problem with the leading indicator is that it exhibits much more month-to-month volatility than the caseload. It stands to reason that a smoothed version of the leading indicator, which distinguishes the underlying signal from month-to-month noise, should provide a better predictor of the relatively smooth caseload. I propose a means for choosing the key smoothing parameters to optimize a measure of out-of-sample forecasting performance.

I develop this approach more formally in the third section of the paper, after describing my data in the next section. I present results in section IV. Finally, in the

conclusion, I discuss a number of ways in which the approach might be generalized in future work.

II. Data

My data come from the California Department of Social Services' Form CA237, which tracks the movement of cases on and off the rolls on a monthly basis. The data consist of monthly state-level welfare entries, welfare exits, and the caseload as of the end of the month. My sample period runs from July 1985 to March 2005. Since the late 1990s these data have been available on-line at the DSS web site.¹

Table 1 presents some summary statistics. The monthly caseload, which is depicted in Figure 1, averaged 650,632 cases over the sample period. On average, entries and exits are similar, averaging roughly 41,000 to 42,000 per month. This implies that on average, about 13 percent of the caseload turns over each month. The average exit rate, defined as exits divided by the previous month's caseload, is 6.7 percent.

Figure 2 depicts entries and exit rates. Entries are represented by the short-dashed line and the exit rate is represented by the long-dashed line. The scale for entries is on the left; the scale for the exit rate is on the right. The solid line in the figure is the ratio of entries to exits, which I discuss in more detail below.

As compared to the caseload, entries and exits exhibit much more month-to-month volatility. Nevertheless, the general trends in entries and exits help explain the substantial swings in the caseload during the sample period. As seen in Figure 2, entries rose gradually from the beginning of the sample period to about the beginning of 1994. The exit rate was roughly constant, aside from monthly volatility, until the late 1980s. It

¹ See <http://www.dss.cahwnet.gov/research/>.

then fell till early 1994. Thus the rise in the caseload in the early 1990s stemmed from both an increase in entries and a decrease in the exit rate.

After early 1994, entries declined and the exit rate rose. Just as the increase in the caseload during the early 1990s stemmed from increasing entries and a decreasing exit rate, the decrease in the caseload during the latter half of the 1990s resulted from a decrease in entries and an increase in the exit rate. Grogger, Haider, and Klerman (2003) estimated that decreasing entries accounted for about 66 percent of the caseload decline between 1995 and 2000, with increasing exit rates accounting for the remainder.

After the 1990s, entries began to rise again, at least until the beginning of 2004. By themselves, increasing entries would result in higher caseloads. However, the exit rate continued to rise until mid-2004. Thus between the beginning of 2000 and 2004, the caseload fell due to increases in the exit rate, despite increasing entries. Since the beginning of 2004, the pattern has reversed. By itself, the sharp decline in the exit rate in 2004 would have increased the caseload. Figure 1 indicates that it has indeed risen, but the increase would have been greater without the decline in entries that appears at the end of the sample period.

III. Methods

A. The Welfare Caseload as a Markov Chain

An important feature of Figures 1 and 2 is that turning points in entries and exits tend to precede turning points in the caseload. This is most obvious in the mid-1990s. Whereas entries peaked and exits reached a low in early 1994, the caseload rose until March 1995. Markov theory provides insights into why this is true. It also motivates a leading indicator of the caseload that is based on entry and exit data.

The basic building block of a Markov chain is a relation that describes the current caseload as a function of the past caseload and current entries and exits. The simplest possible model of caseload evolution is given by

$$\begin{aligned} C_t &= C_{t-1} + E_t - X_t \\ &= C_{t-1}(1 - x_t) + E_t \end{aligned} \tag{1}$$

where C_t represents the caseload in month t , C_{t-1} represents the caseload in the previous month, E_t represents current entries, X_t represents current-month exits, and x_t represents the exit rate, defined as $x_t = X_t/C_{t-1}$. The top line says simply that current cases equal last month's cases plus entries less exits. The second line expresses the same relationship in terms of entries and the exit rate, which is a more useful form for the analysis to follow.

Because the current caseload depends on the past caseload, the caseload exhibits inertia. This means that it takes time for the full effects of changes in entries or exits to be reflected in the caseload. Markov forecasting exploits this lag to forecast future caseloads on the basis of current entries and exits.

Equation (1) is referred to as a first-order Markov chain because the process depends on only current entries and exits and the first lag of the caseload. Higher-order models could be employed for forecasting, but they would require household-level data rather than the state-level aggregated data I use here. The reason is that households are only at risk of exiting during the second month of their spell if they do not exit during the first month. It is impossible to distinguish first-month exits from second-month exits with aggregate data. I discuss higher-order models further in Section V.

A key property of a Markov chain is its steady state (see Klerman and Haider 2004 and references therein). The steady state represents the target toward which the caseload would converge if entries and the exit rate were to remain constant indefinitely

at values E and x , respectively. For the simple model in equation (1), the steady state is given by

$$\bar{C} = \frac{E}{x}. \quad (2)$$

In the welfare context, the notion of a steady state may seem to be of little value. As shown in Figure 2, entries and exit rates are never constant for very long.

Nevertheless, the notion of the steady state motivates the leading indicator of the welfare caseload which is the key to Markov forecasting. I refer to this quantity as the implied steady state, or ISS. It is given by

$$\bar{C}_t = \frac{E_t}{x_t}. \quad (3)$$

The implied steady state has nothing to do with constant entry or exit rates; it is a function of current entries and exit rates, and thus changes over time just as entries and exits change over time. However, if entries and the exit rate were to remain constant at their current values, the ISS represents the target toward which the caseload would converge.

In fact, the ISS represents a target more broadly construed. To see this, note that the change in the caseload from month $t-1$ to month t can be written as

$$C_t - C_{t-1} = x_t(\bar{C}_t - C_{t-1}). \quad (4)$$

Since the exit rate is always positive, equation (4) says that the change in the caseload will always have the same sign as the difference between the ISS and last month's caseload. Thus if the ISS exceeds last month's caseload, then the caseload will grow. If the ISS is less than last month's caseload, then the caseload will fall.

The ISS is plotted in Figure 2 as the solid line. As entries grew and the exit rate fell in the early 1990s, the ISS rose sharply, preceding the rise in the caseload. The peak in the ISS occurred in March 1994, one year before the peak in the caseload. The ISS then fell for several years as the caseload fell. The ISS began rising again in late 2002, about a year before the caseload reached its recent low in November 2003.

B. Smoothing and Forecasting

Since the turning points in the ISS precede the turning points in the caseload, the ISS provides a leading indicator of the caseload that can be used in forecasting. One problem with the ISS is that it exhibits much greater month-to-month volatility than the caseload. Its value for forecasting would be greater if the month-to-month noise could be distinguished from the underlying signal that predicts the direction of the caseload.

One technique for distinguishing signal from noise is smoothing. In principle, one could smooth the ISS directly and use the smoothed values for forecasting. Alternatively, one could smooth the entries and exit rate, then construct a smooth ISS from the smoothed entries and exits using equation (3). This latter approach has the advantage of illustrating whether trends in the ISS stem from trends in entries, the exit rate, or both. Since this latter approach provides more information, this is the approach I adopt.

There are many different techniques one could choose for smoothing, ranging from moving averages to polynomial regression. In the current application I utilize lowess smoothing. Lowess smoothing involves running a separate smoothing regression of the dependent variable (entries or the exit rate) on the independent variable (time) for each value of the independent variable. In this case, this means that a separate smoothing

regression is run for each data point. The regression for each data point contains a limited number of observations within a neighborhood of the data point. The smoothed value of the dependent variable is the predicted value from that regression (Cleveland 1994).

A virtue of lowess smoothing is that it is local in nature. Only nearby observations are used to form the smoothed values. In polynomial regression, in contrast, all observations are used to form smoothed values for each data point. The result is that aberrations in end-point observations, for example, can affect smoothed values for all data points.

Of course, one has to define which observations are within the neighborhood of each data point. In STATA's implementation of lowess smoothing, the neighborhood is defined by a bandwidth that indicates the proportion of observations that is used in each lowess regression. Bandwidths near zero imply that very few observations are used in each regression, so the fit is highly local. A bandwidth of one implies that all observations are used in each regression, so the fit is a straight line.

An important question is how to select the bandwidth. Too small a bandwidth delivers little smoothing, defeating the purpose. Too large a bandwidth may oversmooth and miss important trends. Furthermore, the degree of smoothing will affect one's forecasts. Moreover, bandwidths are not estimated like regression coefficients, but rather chosen by the analyst. This could add an element of subjectivity to the approach. To avoid the subjectivity problem, I use the data as a guide to parameter selection, choosing the bandwidths to minimize the mean squared error of the forecasts that are based on the smoothed ISS.

More concretely, let $E_t(\beta_E)$ represent the smoothed estimate of E_t for bandwidth β_E . Let $x_t(\beta_x)$ represent the smoothed estimate of x_t for bandwidth β_x . Define $\bar{C}_t(\beta) = E_t(\beta_E) / x_t(\beta_x)$ as the smoothed estimate of the ISS given parameter vector β , where $\beta = (\beta_E, \beta_x)$. The smoothed ISS is used to construct a forecasting regression that takes the form

$$C_t = \alpha_0 + \alpha_1 \bar{C}_{t-L}(\beta) + u_t, \quad t = L+1, \dots, T \quad (4)$$

where $\alpha = (\alpha_0, \alpha_1)$ is a vector of parameters that is estimated by OLS, u_t is a zero-mean disturbance term, T is the number of observations used for estimation, and L is the lag length, that is, the number of months by which the ISS leads the caseload.

Most simply, one could carry out a grid search over various values of β and L and choose β and L to minimize the mean square error from equation (4). For each value of (β, L) , one would compute α and the mean squared error. The values of β, L and α that minimized the mean squared error would be used to construct $\bar{C}_{t-L}(\beta)$ and generate forecasts of C_t .

However, using such a within-sample fit criterion could lead to overfitting, which in turn could result in poor forecasts. To avoid overfitting, I choose β and L via cross-validation, which provides an out-of-sample fit criterion (Pagan and Ullah 1999). For each value of (β, L) , I estimate equation (4) not once, but $T-L$ times. From each regression I drop one observation, select optimal values of (β, L) , construct $\bar{C}_{t-L}(\beta)$, estimate α , and compute the squared prediction error for the omitted observation. The average of these out-of-sample squared prediction errors is used as the objective function.

For forecasting I then employ the values of (β, L) that minimized the objective function to construct $\bar{C}_{t-L}(\beta)$ and estimate α using all the observations.

The grid search was carried out for values of L ranging from 10 to 15 and for values of β_E and β_x ranging from .1 to .5 by increments of .01. The optimal parameter values fell well within the interior of the search space, as I discuss next.

IV. Results

Since my ultimate goal is to assess the forecasting performance of the approach, I employ only a subset of the sample in estimation. To assess whether the approach predicts turning points well, I end the estimation sample in November 2002, a full year before the caseload reached its minimum in November 2003. This provides 209 observations for estimation, leaving 28 observations, extending from December 2002 to March 2005, to evaluate the forecasts.

The top panel of Table 2 reports the results of the grid search. The optimal values for β_E and β_x are 0.23 and 0.20, respectively. This means that roughly one-quarter of the observations were used to construct the smoothed entry value for each data point, and one-fifth of the observations were used to construct the smoothed exit value. The optimal lag length L is 14. This means that the optimal forecasting horizon for the model is just over one year.

The smoothed entry, exit rate, and ISS series are displayed in Figure 3, where they are superimposed over the raw series. Despite the relatively low bandwidths, the series appear quite smooth. The key trends in the entry and exit series discussed above are quite prominent. The smoothed ISS seems to fit quite well, even though it was computed indirectly from the smoothed entry and exit rate series rather than estimated

directly. It achieves its maximum value in January 1994, 14 months before the caseload peak in March 1995. It achieves its minimum value in November 2002, 12 months before the recent caseload minimum in November 2003.

Figure 4 displays the smoothed ISS and the caseload. The potential forecasting power of the ISS is quite evident. For the most part, the ISS looks like the caseload to come. In principle, one could simply use a 14-month lag of the ISS to forecast the caseload, although the forecasting regression in general should do a better job.

The results from the forecasting regression are presented in panel B of Table 2. The optimal predictor of the caseload is a value equal to 92.6 percent of the smoothed ISS, lagged 14 months, plus 48,037. Both coefficients are quite large in relation to their standard errors, but since the standard errors do not account for the fact that the smoothed ISS is pre-estimated, it cannot be assumed that the usual t-statistics have standard asymptotic normal distributions. It seems likely that the standard errors are too small, since they neglect the considerable dependence that arises from pre-estimation.

Figure 4 depicts the forecasts. These were constructed by using the grid search values in Panel A of Table 2 to construct smoothed ISS values for the period December 2002 to March 2005, then using the regression parameters from Panel B to construct both point and interval forecasts 14 months in the future. In Figure 4, the caseload data used in estimation are represented as solid circles; the point forecasts are represented by the solid curve. Upper and lower 95 percent predication intervals are represented by dashed curves. Because these are based on the conventionally-computed standard error of the forecasting regression, they neglect the additional source of error that arises from the pre-estimated regressor. Thus they are probably optimistic, meaning that the true

confidence intervals should be wider. The actual values of the caseload for the period December 2002 to March 2005 are represented by open gray circles.

One feature of the forecasts is that they appear to be too low. For most of the post-estimation period, the forecasts cluster around the upper end of the 95 percent prediction interval rather than the middle. On average, the forecasts are low by about 18,200 cases, or about 4 percent of the average number of cases over the period.

At the same time, the model does a reasonably good job at predicting the turning point that appears toward the end of the sample period. The forecasts achieve their minimum value in January 2004, just two months after the caseload achieved its minimum value. Furthermore, the forecasts strongly suggest that the November 2003 low point represents a true turning point, rather than just a relatively minor fluctuation around a stable level. After the beginning of 2004, the forecasted caseload rises sharply.

Table 3 presents quarterly forecasts for the 14-month period following the end of the sample period. Values in columns (1) to (3) are those depicted in Figure 5; that is, they are constructed from estimates based on pre-December-2002 data. The point forecasts rise by 12 percent over this period, from 461,745 to 516,210. The interval forecast for the end date, May 2006, runs from 497,075 to 535,344. Relative to the March 2005 value of about 463,000 cases, this represents growth of 7.4 to 15.6 percent.

Of course, in a real forecasting setting, one would use all the data at one's disposal to generate predictions of the future caseload. This suggests constructing forecasts based on estimates obtained over the full sample period that ends in March 2005. Furthermore, comparing estimates based on different sample periods allows one to assess the stability of the forecasts.

Full-sample forecasts are shown in columns (4) through (6) of Table 3. At the beginning of the forecast period, the two sets of point forecasts differ by about 8,000 cases, or about 2 percent of the forecasted caseload. By the end of the forecast period, the forecasts are more similar, differing only by about 3,000 cases. Throughout the forecast period, the reported forecast intervals overlap to a considerable extent, indicating the true forecast intervals must overlap to an even greater extent. Thus the forecasts are fairly stable; adding nearly two-and-a-half years of new data has a relatively small effect on the forecasts, even though those data come largely from the period after the turning point had been reached. Both sets of forecasts predict that the California welfare caseload should grow fairly substantially by May 2006.

V. Conclusions

With the passage of PRWORA, states bear the risk associated with their welfare programs. Federal funding is now provided via a block grant. If the caseload were to grow dramatically, states could be forced to reconfigure their programs, reduce spending on other programs, or raise taxes. Accurate caseload forecasts cannot change the growth rate of the caseload, but they should help reduce surprises, which in turn should reduce any fiscal dislocations that arise from caseload growth.

The Markov forecasting method proposed above provides analysts with a new tool to forecast caseloads. It is fairly accurate and reasonably stable. Perhaps most importantly, it handles turning points well. Based on data that ended a year before the most recent turning point, the model forecast an increase in the California caseload beginning in January 2003, only two months after the caseload actually reached its recent minimum.

Although the method has been illustrated with welfare data, the technique is more broadly applicable. It can be used to forecast any variable that can be thought of as the accumulation of entries into and exits from a program. It should be useful in forecasting not only welfare caseloads, but also Food Stamp, foster care, and Medicaid caseloads.

There are a number of ways one might extend the analysis to produce more useful and potentially more accurate forecasts. One is to generalize the model by allowing for higher-order Markov chains. If exit rates vary by spell length, and spell lengths vary in predictable ways, then higher-order models may provide more accurate forecasts. To allow exit rates to vary by spell lengths requires household level data, however, which may not be as readily available as the aggregate data on entries and exits which can be used to estimate the first-order models described above.

Another extension would be to combine the econometric approach to forecasting with the Markov approach. The way to do this would be to model entries and exit rates econometrically as a function of variables such as benefit levels, eligibility rules, and economic conditions. One could use such a model to forecast changes in entries and exits as a function of changes in those variables, then use the forecasted entries and exits to construct Markov forecasts of the caseload.

There is a sense in which this exercise is unnecessary. Ultimately, all changes in economic or policy conditions must affect the caseload by affecting either entries or exits. Thus the effect of such changes will be implicitly captured by the approach outlined above, even if they are not explicitly modeled. At the same time, it is often helpful to explain forecasts in terms of the reasons for the observed changes, so combining the Markov approach with econometric forecasting may increase the utility of the forecasts.

References

- Box, G.E.P. and G. M. Jenkins. Time Series Analysis, Forecasting, and Control. San Francisco: Holden Day, 1970.
- Cleveland, W.S. The Elements of Graphing Data. Summit, NJ: Hobart Press, 1994.
- Dynarski, Mark, Anuradha Rangarajan, and Paul Decker. Forecasting Food Stamp Program Participation and Benefits. Princeton, NJ: Mathematica Policy Research, August 1991.
- Granger, Clive W. J. and Paul Newbold. Forecasting Economic Time Series. New York: Academic Press, 1977.
- Grogger, Jeffrey, Steven Haider, and Jacob A. Klerman. "Why Did the Welfare Rolls Fall during the 1990s? The Importance of Entry." American Economic Review 93 (2), May 2003, 288-292.
- Klerman, Jacob A. and Steven Haider. "A Stock-Flow Analysis of the Welfare Caseload." Journal of Human Resources 39 (4), Fall 2004, 865-886.
- Pagan, A. and A. Ullah, 1999, Nonparametric econometrics (Cambridge University Press, Cambridge).




Figure 1
Monthly California Welfare Cases, July 1985 to March 2005

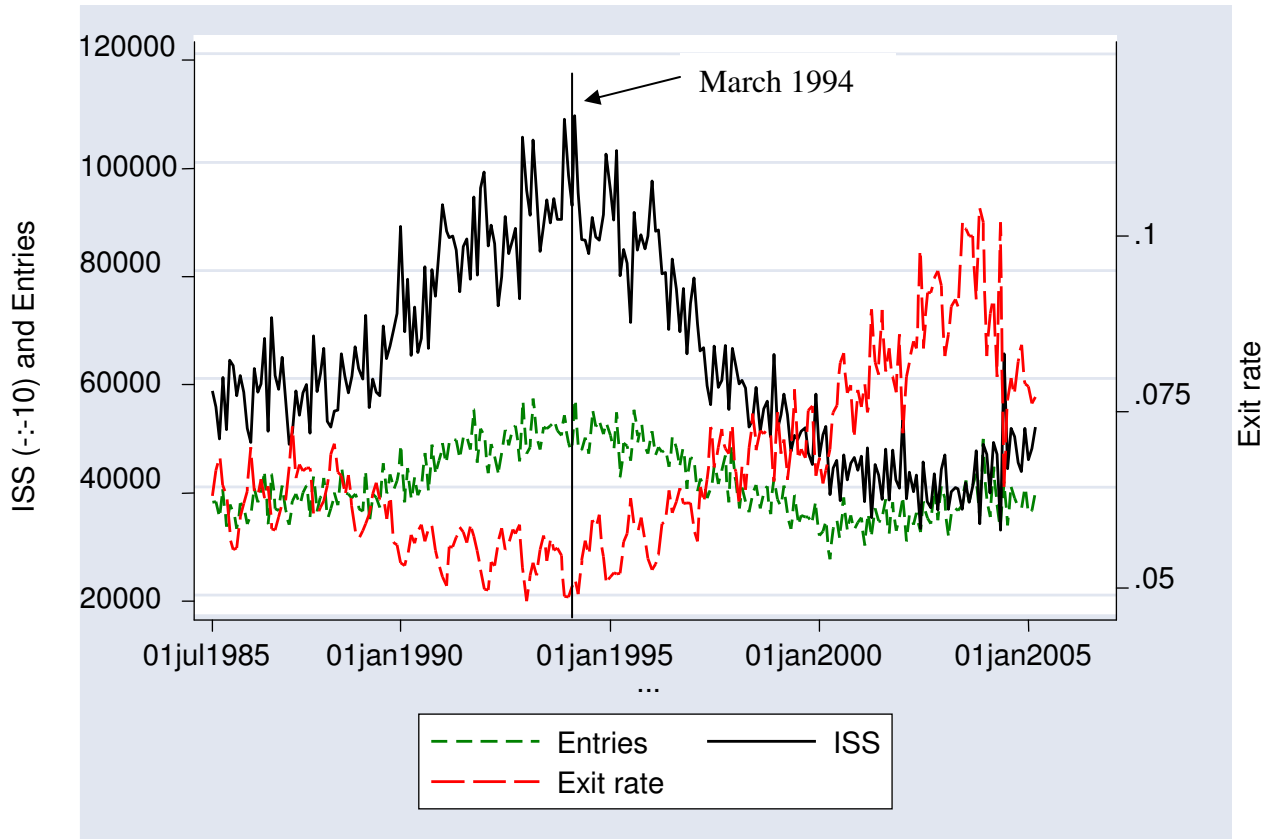


Figure 2
Raw Welfare Entries, Exit Rate, and Implied Steady State (ISS)



Figure 3
Raw and Optimally Smoothed Welfare Entries, Exit Rate, and Implied Steady State

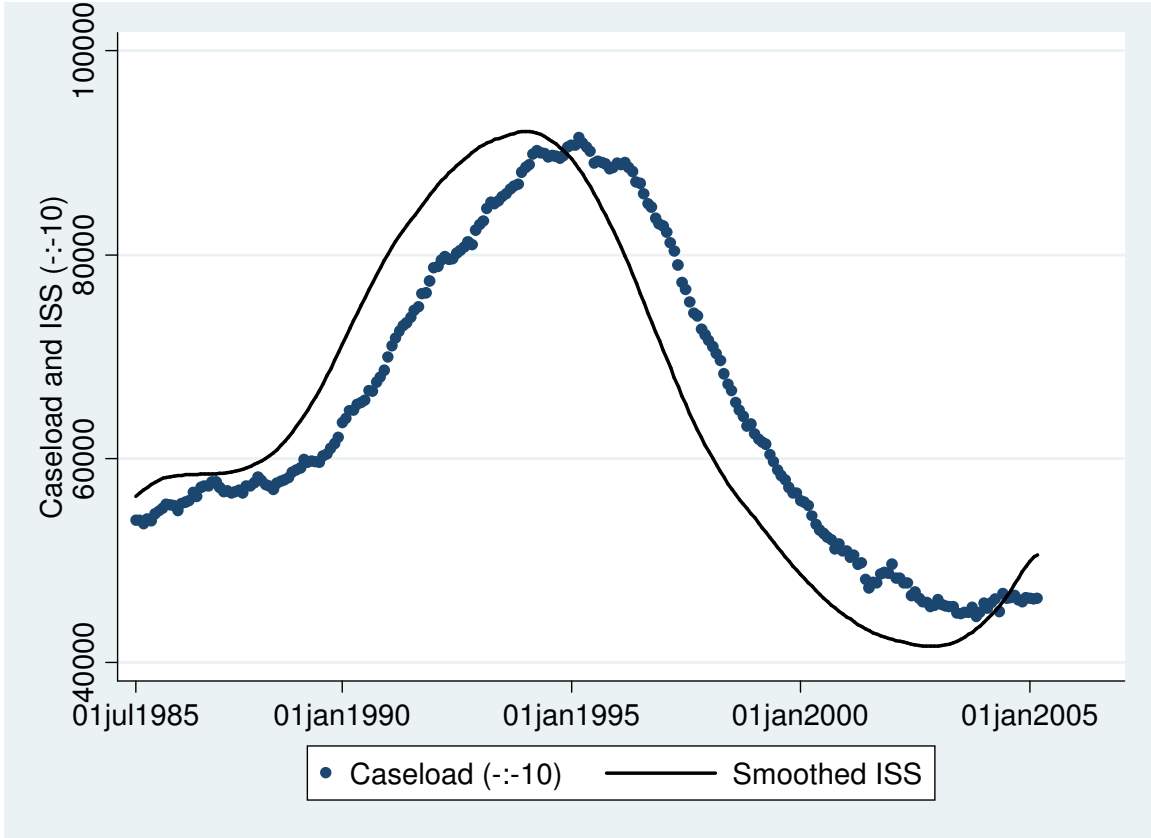


Figure 4
Welfare Cases and Optimally Smoothed Implied Steady State




Figure 5
Welfare Cases, Actual and Forecasted

Table 1
Summary Statistics

Variable	Mean	Standard deviation
Caseload	650,632	151,283
Entries	41,302	6,180
Exits	42,169	6,104
Exit rate	0.067	0.013

Sample size is 237.

Table 2
Estimation Results

<u>A. Grid search results</u>		
Bandwidth for entries (β_E);	0.23	
Bandwidth for exit rate (β_x);	0.20	
Lag length:	14	
<u>B. Regression results</u>		
Variable	Coefficient	Standard error
Constant	48037	(3194)
Smoothed ISS, lagged 14 months	0.926	(0.005)
R-square	0.992	

Sample size is 195 (September 1986 to November 2002).

Table 3
Point and Interval Caseload Forecasts, May 2005 to May 2006

End of estimation period:	November 2002			March 2005		
	Lower end, 95% prediction interval (1)	Point forecast (2)	Upper end, 95% prediction interval (3)	Lower end, 95% prediction interval (4)	Point forecast (5)	Upper end, 95% prediction interval (6)
Month						
May 2005	442,559	461,745	480,930	448,924	469,905	490,887
August 2005	454,885	474,058	493,210	460,545	481,519	502,492
November 2005	470,292	489,450	508,608	474,666	495,630	516,594
February 2006	486,699	505,842	524,985	488,566	509,521	530,476
May 2006	497,075	516,210	535,344	498,360	519,310	540,259