

On the Impact of Sampling Error on Modeling Seasonal Time Series October 2009

Stuart Scott

Bureau of Labor Statistics, 2 Massachusetts Ave NE, room 1950,
Washington, DC 20212**Abstract**

The presence of sampling error in an observed time series may obscure underlying features, such as seasonality. Based on simulated series containing sampling error, estimation of the seasonal parameter in ARIMA models is examined, with and without accounting for the sampling error. In the former case, results include cases where the sampling error is incorrectly specified. Sensitivity of estimation to length of the series and relative size of the sampling error is addressed. Empirical results from a large set of employment series are presented and interpreted in light of the simulation results. The series are from the Bureau of Labor Statistics' Current Employment Statistics survey.

Key Words: signal extraction, ARIMA models, employment series

1. Introduction

The classical starting point for analyzing an economic time series is to consider it in terms of trend or trend-cycle, seasonal, and irregular or noise. In practice, most monthly or quarterly series come from surveys, in which case we may appropriately add a sampling error component. This component adds nothing to economic analysis, so we may wish to extract the signal in order to understand better the underlying behavior of the time series. The aim of this paper is to encourage and illustrate treatment of sampling error in modeling or seasonally adjusting time series. It analyzes a simple simulation experiment to indicate how the method can work and to point out potential pitfalls. Signal extraction for seasonal time series began about 30 years ago with work of Hillmer and Tiao (1982) and has a much longer history in mathematics and other areas of science. While not addressing the specific issues of this paper, Gomez and Maravall (2001) give some theoretical results on signal extraction for seasonal time series with a nice example.

My motivation comes from work on variances for X-11 seasonally adjusted series using a model-based method due to Bell & Kramer (1999). The method calls for estimating a signal model in the presence of sampling error (SE). An ongoing test of the method involves 137 series from the Current Employment Statistics program at the Bureau of Labor Statistics (BLS). Figure 1a. contains a histogram of the seasonal MA parameter θ_{12} from ARIMA modeling of the observed series. There is a broad range of values, mostly concentrated at the upper end, toward relatively stable seasonality. Figure 1b. shows the distribution for θ_{12} in the ARIMA model for the signal, based on signal extraction. The distribution has some concentration near 1 and an even greater concentration near 0 and below. In fact, the values stretch out all the way to -1. Most negative values of θ_{12} do not correspond to decomposable models; restriction of θ_{12} to $[0,1]$ encompasses a broad range of seasonality from highly varying to deterministic. Clearly, signal extraction has not worked well for these series. What has gone wrong?

Two plausible explanations are (1) instability in signal extraction and (2) misspecification of sampling error. We shall see that series length affects stability. In applying signal extraction, external information on sampling error has been used and this information may be incorrect. After all, the variance of a variance estimate can be quite large.

The next section presents a simulation experiment and explains what can be expected from basic model theory. After that, simulation results are given. Lessons for the employment application are drawn. A final section gives tentative findings.

2. A Simple Simulation Experiment

Our model for an economic time series is

$$Y = T + S + I \quad (1a)$$

$$y = Y + \varepsilon, \quad (1b)$$

where y is the observed series, Y the signal or population consisting of trend (trend-cycle), seasonal, and irregular or noise, and ε is sampling error. Hillmer and Tiao (1982) show how ARIMA models for series of the form (1a) can be decomposed. Harvey (1989) accomplishes decomposition with structural models. Bell and Hillmer (1990) pioneer work with models (1b) having a sampling error component. Tiller (1992) at BLS uses structural models to estimate and remove a sampling error component to obtain the official estimates of the unemployment rate for 50 U.S. states and the District of Columbia.

ARIMA-based seasonal adjustment accounting for sampling error can be carried out as follows using publicly available software:

- (1) adopt a model for the signal and a complete model specification for the sampling error,
- (2) apply signal extraction to obtain signal model parameters and a decomposition $y = \hat{Y} + \hat{\varepsilon}$ using REGCMPNT,
- (3) decompose the signal model and \hat{Y} using TRAMO-SEATS or X13-ARIMA-SEATS,

and

- (4) compute the seasonally adjusted value $\hat{A} = y - \hat{\varepsilon} - \hat{S}$.

Both REGCMPNT, developed by Bell (2004) and X13-ARIMA-SEATS are available from the U.S. Bureau of the Census (<http://www.census.gov/srd/www/x12a/>); the Bank of Spain offers TRAMO-SEATS (<http://www.bde.es/servicio/software/econome.htm>). This paper focuses on the second step, signal extraction.

For the simulation experiment, the signal is specified as an airline model,

$$(1 - B)(1 - B^{12})Y_t = (1 - \theta_1 B)(1 - \theta_{12} B^{12})a_t$$

with

$$\theta_1 = \theta_{12} = 0.6, \quad \sigma_a^2 = 121.$$

The sampling error is white noise, with variance 100,

$$\varepsilon_t \sim WN(0, 100).$$

Simulation steps:

- (1) decompose the signal model into models for T , S , and I ,
- (2) simulate 1600 series of length 216 (18 years) for each of the four components and form y and Y ,
- (3) carry out and analyze signal extraction for
 - (a) 10-year and 18-year series,
 - (b) alternative values for σ_ε^2 , 100, 144, and 49.

We shall see the impact of over- and underestimation of the sampling error, both for the 18-year series and for 10-year series. The latter are the central 10 years from the full series.

What does theory tell us about behavior of the observed series? The differenced observed series is

$$w_t = (1 - B)(1 - B^{12})y_t = u_t + v_t$$

where

$$u_t = (1 - .6B)(1 - .6B^{12})a_t, \quad v_t = (1 - B)(1 - B^{12})\varepsilon_t,$$

that is, w_t is the sum of MA(13) processes, the latter differenced white noise. Granger and Morris (1976) present several results for sums of independent ARMA processes, including the Proposition below. In particular, the sum of MA processes is again an MA process.

Proposition (Granger & Morris, 1976).

Suppose the components are independent.

$$ARMA(p_1, q_1) + ARMA(p_2, q_2) = ARMA(p^*, q^*),$$

where

$$p^* \leq p_1 + p_2, \quad q^* \leq \max(p_1 + q_2, p_2 + q_1).$$

Can we approximate our sum with a simple (001)(001) model? Matching moments at lags 0, 1, and 12 yields the approximation

$$w_t^* = (1 - \tau B)(1 - \tau B^{12})b_t$$

with

$$\tau = .744, \quad \sigma_b^2 = 258.5.$$

The autocovariances in Table 1 show that w^* is a close approximation. Both the MA parameters and the disturbance variance increase in value to account for the additional variability and correlation. Note that $\tau = .744$ is a “compromise” between .6 for the signal and 1 for the differenced white noise. This result suggests that, given θ_{12} for an observed series containing white noise sampling error, we may expect θ_{12} for the underlying signal to be smaller.

Table 1. Autocovariances of Differenced Observed Series w and Approximation w^*

	Lag	0	1	2	...	10	11	12	13
w		623.8	-298.7	0		0	143.6	-298.7	143.6
w^*		623.8	-298.7	0		0	143.1	-298.7	143.1

To compare more closely the signal and observed series models, we decompose the differenced series. The component models implied by an airline model are

$$\begin{aligned} (1 - B^2)T_t &= \theta^{(T)}(B)\eta_t \\ U(B)S_t &= \theta^{(S)}(B)\zeta_t, \\ I_t &\sim WN \end{aligned}$$

where

$$\begin{aligned} \theta^{(T)}(B) &= 1 - \theta_1^{(T)}B - \theta_2^{(T)}B^2 \\ U(B) &= 1 + B + \dots + B^{11} \\ \theta^{(S)}(B) &= 1 - \theta_1^{(S)}B - \dots - \theta_{12}^{(S)}B^{12} \end{aligned}$$

These formulas readily yield stationary models for the differenced components, allowing us to decompose the variance of the differenced series, displayed in Table 2. Going from the signal to the observed series, the sizable sampling error component adds much to the total variance. Most of the added variability goes to the irregular component. The variance in the seasonal component is higher, but not by much. In fact, its share of total variance has dropped in half, from 6.3% to 3.0%. This suggests that in the favorable case of white noise sampling error an overall ARIMA model may capture the seasonality in the series reasonably well.

Table 2. Variance after Differencing Simulation Models

	Trend	Seasonal	Irregular	Total
Signal	12	14	198	224
Observed	13	19	592	624

3. Simulation Results

Table 3 contains results from modeling the simulated signal series $Y^{(r)}$ and the observed series $y^{(r)}$, $r = 1, \dots, 1600$. Mean and median estimates for the MA and disturbance variance parameters are close to the true values .6 and 121 for the signal and .744 and 258 for the observed series. Figure 2 contains histograms for θ_1 and θ_{12} for 18- and 10-year series. Results for the 18-year series are quite concentrated with 90% of the distribution lying roughly within ± 0.10 for both signal and observed cases. For the 10-year observed series, the most notable features are the stronger central peak for θ_1 and some clustering of values near 1 for θ_{12} . All histograms in the paper have a common horizontal axis, so it's hard to assess symmetry in Figure 2. Note, however, that skew is negligible in all cases; the largest magnitude in Table 3 is 0.5 for θ_{12} for signal series using 10-year series.

Before examining the signal extraction results, let's review the signal extraction task. Airline models are invertible for $\theta_{12} \in (-1, 1)$. On the other hand, Hillmer & Tiao (1982) only establish that these models are decomposable for $\theta_{12} \in [0, 1)$. This range is quite adequate for representing a range of seasonal behavior from rapidly changing to essentially deterministic. Our data are the differences

$$w_t = (1 - \theta_1 B)(1 - \theta_{12} B^{12})a_t + (1 - B)(1 - B^{12})\varepsilon_t$$

and we assume full knowledge of the SE term. Maximum likelihood methods are used to estimate parameters θ_1 , θ_{12} , and σ_a^2 for the first term. Variance for this term can be increased by increasing any or all of the parameters. When the bulk of the variance is assigned to the second term, it seems intuitive that estimation can be unstable; this indeed is seen below when the SE contribution is overestimated.

Table 4 and Figures 3-5 contain the signal extraction results. In the most favorable case, signal extraction is carried out for 18-year series, supplied with the correct SE variance $\sigma_\varepsilon^2 = 100$. Table 4a. shows that both mean and median values are close to the true values for both series lengths. 90% of the distribution lies within (.44, .75) for θ_1 and within (.42, .78) for θ_{12} . A few small estimates, including two near -1 cause the θ_1 distribution to have negative skew and contribute to the substantial kurtosis. Still, signal extraction performs well in this case, with fairly tight distributions.

Table 3. ARIMA Modeling Estimates for Signal and Observed Series
for 18- and 10-year Lengths (no signal extraction)

a. Signal series

		θ_1		θ_{12}		Disturbance variance	
	Length	18	10	18	10	18	10
Mean		.60	.61	.60	.62	120	118
Quantile							
Median		.60	.60	.60	.61	119	117
5%		.51	.47	.50	.44	100	90
95%		.70	.75	.71	.80	141	148
I'quartile range		.08	.11	.09	.14	17	23
Skew		0.0	0.1	0.1	0.5	0.2	0.3
Kurtosis		0.0	0.4	0.3	0.9	0.1	0.1

b. Observed series

		θ_1		θ_{12}		Disturbance variance	
	Length	18	10	18	10	18	10
Mean		.73	.74	.73	.75	260	252
Quantile							
Median		.73	.73	.73	.73	259	251
5%		.65	.62	.63	.57	219	190
95%		.82	.87	.84	.998	305	319
I'quartile range		.07	.09	.08	.15	35	52
Skew		0.1	0.3	0.4	0.4	0.2	0.2
Kurtosis		0.3	1.0	0.8	-0.2	0.2	-0.2

Table 4. Signal Extraction Estimates as a Function of Sampling Error Input and Series Length

a. $\sigma_\varepsilon^2 = 100$, true value

		θ_1		θ_{12}		σ_a^2	
	Length	18	10	18	10	18	10
Mean		.60	.58	.60	.61	118	110
Quantile							
Median		.60	.60	.60	.60	118	109
5%		.44	.32	.42	.29	75	51
95%		.75	.81	.78	.993	165	175
I-quartile range		.12	.17	.13	.24	36	51
Skew		-4.1	-3.2	-0.1	-1.2	0.2	0.1
Kurtosis		56.6	22.7	0.8	7.2	0.4	0.1

b. $\sigma_\varepsilon^2 = 144$, overestimate

		θ_1		θ_{12}		σ_a^2	
	Length	18	10	18	10	18	10
Mean		.42	.27	.44	.44	59	48
Quantile							
Median		.46	.42	.47	.49	58	46
5%		.10	-.93	.13	-.51	19	4
95%		.67	.75	.71	.995	101	110
I-quartile range		.21	.38	.20	.38	33	52
Skew		-2.7	-1.4	-2.5	-1.4	0.3	0.6
Kurtosis		12.4	1.1	12.8	2.8	0.3	-0.2

c. $\sigma_\varepsilon^2 = 49$, underestimate

		θ_1		θ_{12}		σ_a^2	
	Length	18	10	18	10	18	10
Mean		.69	.69	.69	.70	193	185
Quantile							
Median		.69	.69	.68	.69	192	183
5%		.58	.54	.56	.49	152	127
95%		.79	.85	.82	.995	238	250
I-quartile range		.08	.11	.10	.18	35	50
Skew		-0.0	0.1	0.3	0.3	0.3	0.2
Kurtosis		0.3	1.1	0.7	-0.1	0.2	-0.1

For the 10-year series, negative skew is quite evident for both distributions, with roughly 10% of the estimates lying below 0.40 and a few having strongly negative values. For the θ_{12} distribution, we see some clustering above 0.85. 10% of the estimates are above 0.90 and 5% above 0.95. Thus, results have deteriorated with the shorter series. While the means and medians are close to 0.6, the shape of the distributions shows some unsatisfactory characteristics.

Table 4b. contains results when supplied with an overestimate of the SE variance, $\sigma_\varepsilon^2 = 144$. In terms of variance, this is close to a 50% overestimate; in terms of standard deviations, 10 vs. 12, it is only 20%. For the 18-year series, we see median estimates of 0.46 and 0.47 for θ_1 and θ_{12} and 58 for σ_a^2 . All three values are smaller, to compensate for the large amount of variation attributed to the sampling error. The histogram in Figure 4a. shows greater greater spread than in the previous case; interquartile ranges are about 75% higher. Both distributions have negative skew; negative values occur for 3.7% of the estimates of θ_1 and 2.5% for θ_{12} . Still, 80% of the distribution lies in (.21, .63) for θ_1 and a similar interval for θ_{12} . Unsurprisingly, in the presence of misinformation, signal extraction yields biased estimates, but, most of the time, its estimates provide usable models which are consistent with overall variability in the observed series.

For the 10-year series, signal extraction does not perform well. The large left tail in the histogram for θ_1 in Figure 4b. accounts for 20% of the distribution. Most of these values are negative. They stretch to -1 and result in a mean of 0.27, compared to the median value 0.42. Both tails are heavy for θ_{12} : 10% of the estimates are negative and 13% lie above 0.95. The mean 0.44 and the median 0.49 are reasonable, but, with extreme estimates occurring in both directions, the interquartile range is nearly double that for the 18-year series. We can conclude that estimation is unstable, due to (1) being constrained to yield a relatively small variance and (2) having only 107 data points for estimation of 3 parameters.

Underestimation of the SE variance yields much more stable results, even with 10-year series. Predictably, all three parameters tend to move higher, toward the values from modeling without an SE component. As seen in Table 4c., means and medians for both series lengths are near 0.70 for the MA parameters and 190 for σ_a^2 . Figure 5 shows tight distributions for the MA parameters with the 18-year series. For the 10-year series, the θ_1 distribution still has a strong peak, while that for θ_{12} has more spread, including once again some clustering of values near 1. We see that the signal extraction task is easier when more variability is attributed to the signal.

To round out our picture of signal extraction performance, let us examine correlations between parameter estimates. First, we check correlations in modeling the signal and observed series, which appear in Table 5a. Since the components have been generated independently, it is not surprising that the correlations between $\hat{\theta}_1$ and $\hat{\theta}_{12}$, associated with the trend and seasonal respectively, are negligible for both signal and observed series and both series lengths. Correlations between $\hat{\theta}_{12}$ and the disturbance variance are

Table 5. Sample Correlations between Parameter Estimates

a. Series models

		$\hat{\theta}_1, \hat{\theta}_{12}$		$\hat{\theta}_1, \hat{\sigma}_a^2$		$\hat{\theta}_{12}, \hat{\sigma}_a^2$	
	Length	18	10	18	10	18	10
Signal		.04	-.01	-.04	-.04	-.09	-.23
Observed		.04	.03	-.06	-.07	-.17	-.36

b. Signal extraction

		$\hat{\theta}_1, \hat{\theta}_{12}$		$\hat{\theta}_1, \hat{\sigma}_a^2$		$\hat{\theta}_{12}, \hat{\sigma}_a^2$	
	Length	18	10	18	10	18	10
$\sigma_\varepsilon^2 = 100$.23	.21	.42	.45	.32	.20
$\sigma_\varepsilon^2 = 144$.19	-.00	.59	.62	.46	.28
$\sigma_\varepsilon^2 = 49$.08	.06	.13	.10	.02	-.17

negative, especially for the 10-year series. Perhaps, this is related to the clustering of θ_{12} estimates near 1 (cf. Figure 5b.); low disturbance variance estimates may occur when $\hat{\theta}_{12}$ is near 1.

Turning to signal extraction, when $\sigma_\varepsilon^2 = 100$ is used, correlations are all positive. Apparently, when the series appears less variable overall, all three parameter estimates tend to be smaller, and vice versa. When $\sigma_\varepsilon^2 = 144$ is used, the unstable case, even larger values occur for the pairs $\hat{\theta}_1, \hat{\sigma}_a^2$ and $\hat{\theta}_{12}, \hat{\sigma}_a^2$. The lack of correlation for $\hat{\theta}_1, \hat{\theta}_{12}$ for the 10-year series is likely due to the heavy tails at both ends of the distribution for θ_{12} . This suggests that together all three parameters tend toward smaller values when a relatively small variance is allotted to the signal. In the underestimation case, $\sigma_\varepsilon^2 = 49$, the correlations are intermediate between the other signal extraction cases and the observed series case. It is interesting to note that the correlations for the two series lengths agree in sign and are close in magnitude in most cases; they don't disappear with the longer series.

4. Application to employment series

Our simulation provides evidence that signal extraction becomes unstable when sampling error is a relatively large proportion of overall variance for the observed series. One measure is the ratio of disturbance variances. For the simulation, this ratio is

$$\sigma_\varepsilon^2 / \sigma_b^2 = 100 / 258.5 \approx 0.4.$$

The ratio increases to 0.56 when the overestimate $\sigma_\varepsilon^2 = 144$ is used. Our 137 BLS employment series have the following distribution for this ratio:

	%
(0, 0.4]	23
(0.4, 1]	55
(1, ∞)	22

In other words, compared to the employment series, our simulation has a moderate amount of sampling error. Even so, Figure 4b. shows similarities in shape to Figure 1b. The results depicted in Figure 1b. are consistent with the hypothesis that many of the SE variances are overestimates.

Variances of these series are large in part because they in some cases correspond to fairly detailed industries. Furthermore, they are change on the log scale, and change values are relatively more variable than levels. This choice of series form has been made because (1) monthly change is the most significant statistic for economic analysis and (2) this form admits a simple and realistic form for the SE model. Even using the median of a year's worth of estimates, the SE variance values seem often to be overestimates.

For work on variances for seasonal adjustment mentioned in the Introduction, the results of this paper suggest restricting movement in the model parameters from those estimated for the observed series. Preliminary work in this direction has begun.

5. Findings and Future Work

The results in this paper illustrate that signal extraction can effectively estimate a model for the signal, given sufficient data and good sampling error (SE) information. Some pitfalls are identified in less favorable situations. These findings are limited and tentative, since they are based on a simple simulation experiment. It is likely that some theory exists to explain the results more definitively.

Detailed findings

- For effective signal extraction, more than 10 years of data are needed. Estimation is reasonably stable with 18-year series, even when SE information is incorrect. On the other hand, with 10-year series, undesirable properties emerge.
- Overestimation of SE variance can cause instability. As illustrated in Figure 4b., the combination of limited data and overestimation of SE variance very often results in unlikely or even unusable signal models. The impact differs for the regular MA parameter θ_1 and the seasonal MA parameter θ_{12} . Estimates of θ_1 move strongly in a downward direction from the true value for a large number of cases. Many estimates of θ_{12} behave similarly; also, many move upward close to 1.
- Misspecification of SE variance causes bias in estimating signal parameters. When SE variance is overestimated, the signal disturbance variance tends to be underestimated, as would be expected. However, there is positive correlation between the estimates of this parameter and the MA parameters. In particular, the seasonal MA estimates are biased downward, even when ample data are available. Underestimation of the SE variance causes signal model parameters to move in the opposite direction. In this case, seasonal MA estimates are biased upward.

- Sampling error can cause moving seasonality to appear deterministic.
Most of the histograms for θ_{12} based on 10-year series show some clustering of values near 1, well above the true value. Table 4 shows that 5% of the values are above 0.99, even when the SE information is correct.
- An ARIMA model may adequately capture the seasonality in the observed series when the sampling error is white noise.
This stems from the result in Granger and Morris in Section 3. Even when it is white noise, the sampling error contributes to the seasonality of the series. Using the results in Table 1, we see that an airline model can essentially capture the characteristics of the observed series.

Future work

This study represents only an early look at properties of signal extraction for seasonal time series. The technique becomes more useful and significant in the presence of correlated sampling error. Tiller (1992) has shown that apparent short-term trend effects can come strictly from correlated sampling error. His labor force models incorporate a complicated sampling error component, often modeled as an AR(15) process, based on the rotating panel survey design. For some of the employment series described in Section 4, sampling error has been modeled as white noise, but more often as an MA(1) process, with relatively small values of the MA parameter. Ties to theoretical results should bolster the results on sensitivity.

Acknowledgments

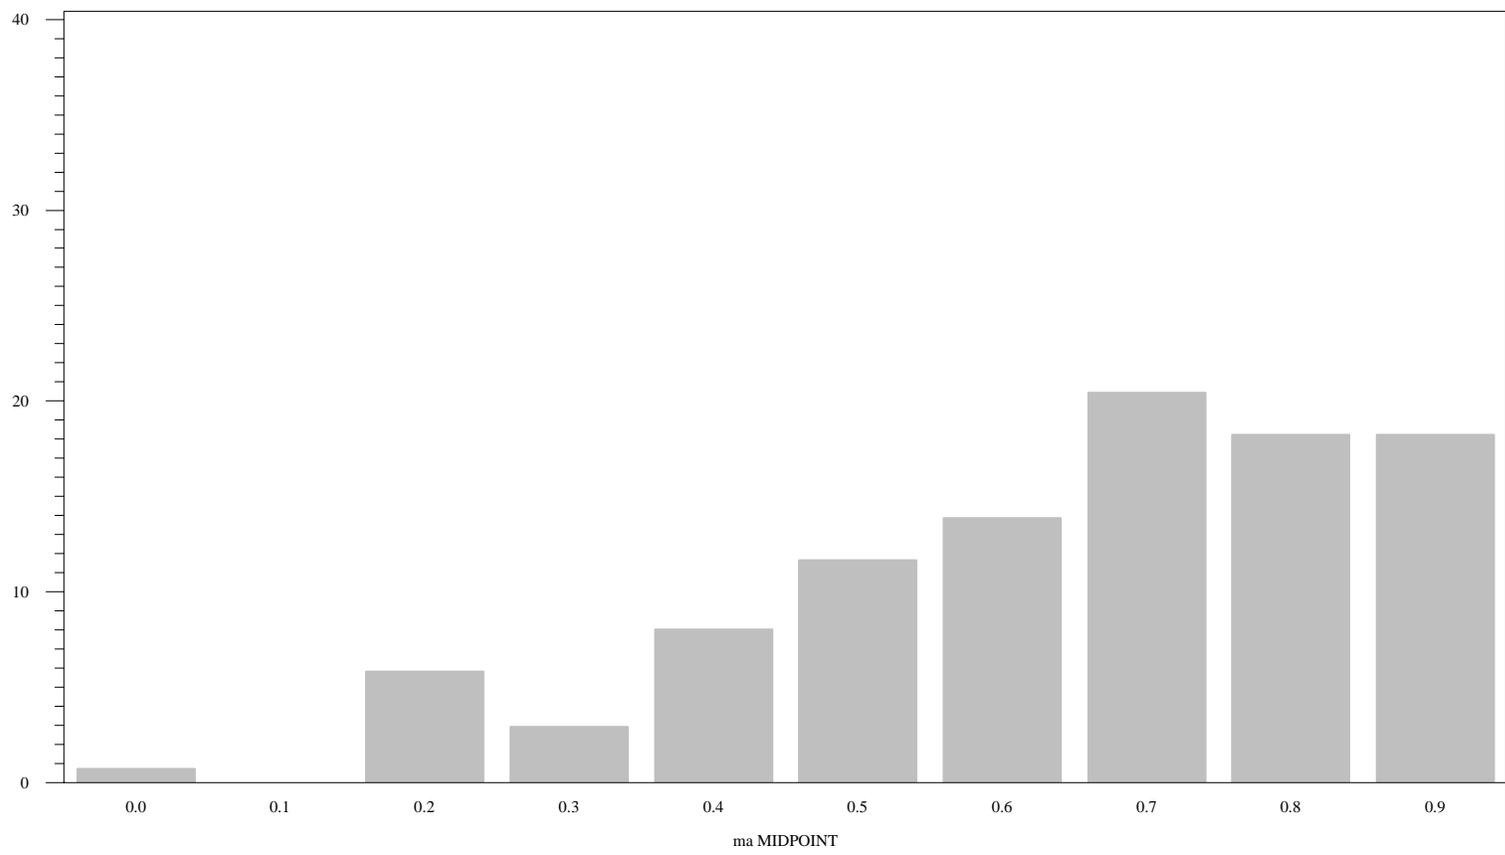
Thanks to Bill Bell, Tucker McElroy, Danny Pfeffermann, and Michail Sverchkov for helpful comments. Thanks to Brian Monsell and Bill Bell for a convenient updated version of REGCOMPNT.

References

- Bell, William R. 2004. On RegComponent time series models and their applications, in *State Space and Unobserved Component Models: Theory and Application*, A. C. Harvey, S. J. Koopman, & N. Shephard, eds. Cambridge, UK, Cambridge University Press. 248-283.
- Bell, William R. and Steven C. Hillmer 1990. The time series approach to estimation for repeated surveys. *Survey Methodology* 16. 195-215.
- Bell, William R. and Matthew Kramer 1999. Toward variances for X-11 seasonal adjustments. *Survey Methodology* 25. 13-29.
- Gomez, Victor and Agustin Maravall 2001. Seasonal adjustment and signal extraction in economic time series. In Daniel Peña, George C. Tiao, and Ruey S. Tsay, eds., *A Course in Time Series Analysis*. John Wiley & Sons, New York.
- Granger, C. W. J. and M. J. Morris 1976. Time series modelling and interpretation. *Journal of the Royal Statistical Society Series A* 139. 246-257.
- Harvey, Andrew C. 1989. *Forecasting, Structural Time Series Models, and the Kalman Filter*. Cambridge, UK: Cambridge University Press.
- Hillmer, S. C. and G. C. Tiao 1982. An ARIMA-model-based approach to seasonal adjustment. *Journal of the American Statistical Association* 77. 63-70.
- Tiller, Richard B. 1992. Time series modeling of sample survey data from the U.S. Current Population Survey. *Journal of Official Statistics* 8. 149-166.

The opinions expressed in this paper are those of the author and do not constitute policy of the Bureau of Labor Statistics.

Figure 1. Distribution of 137 Seasonal MA Parameters
 a. Observed Employment Series



b. Signal Component - from Signal Extraction

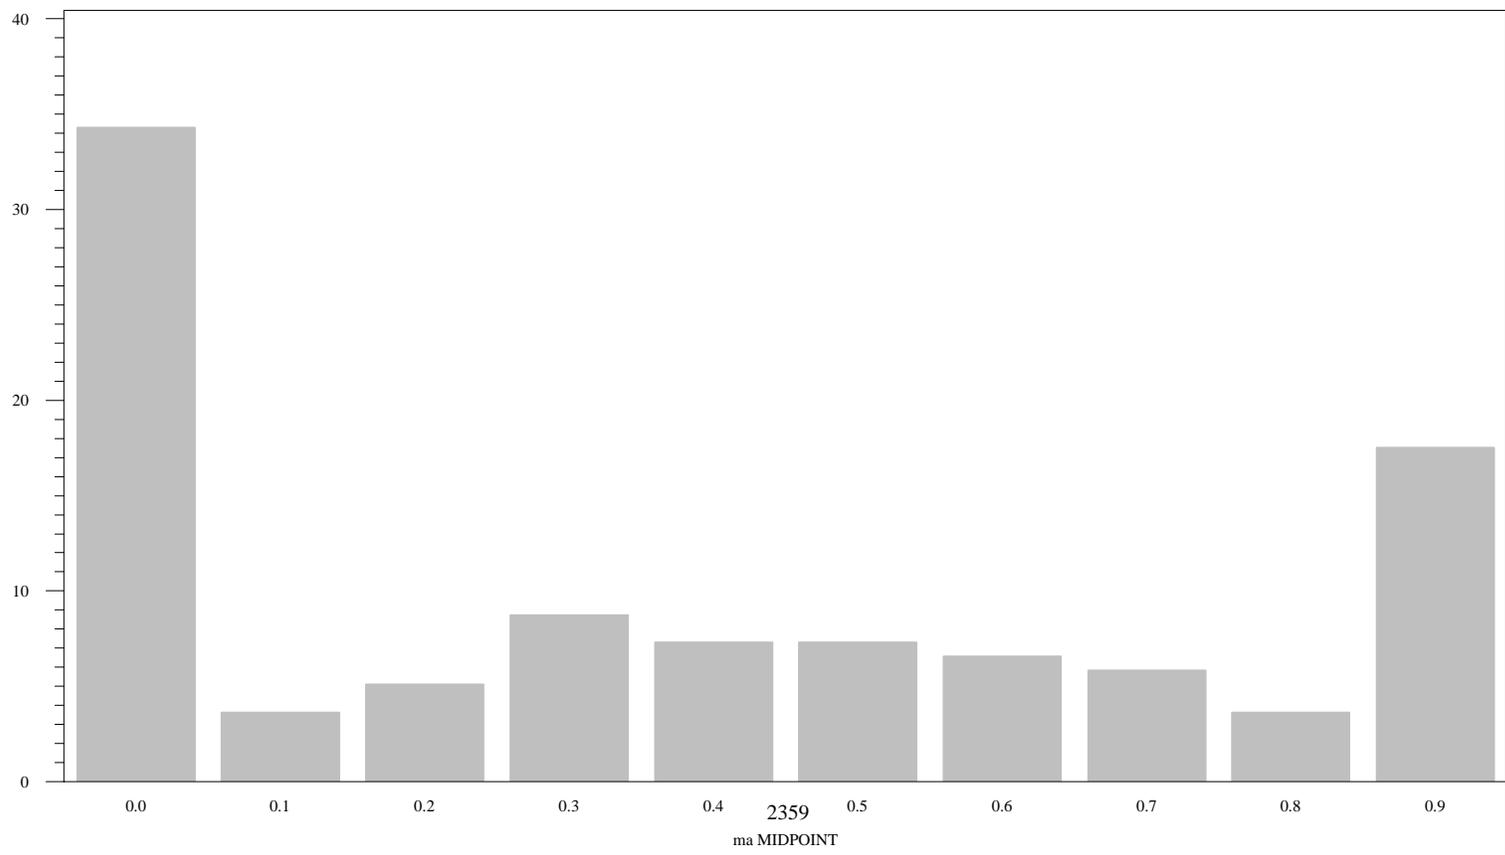
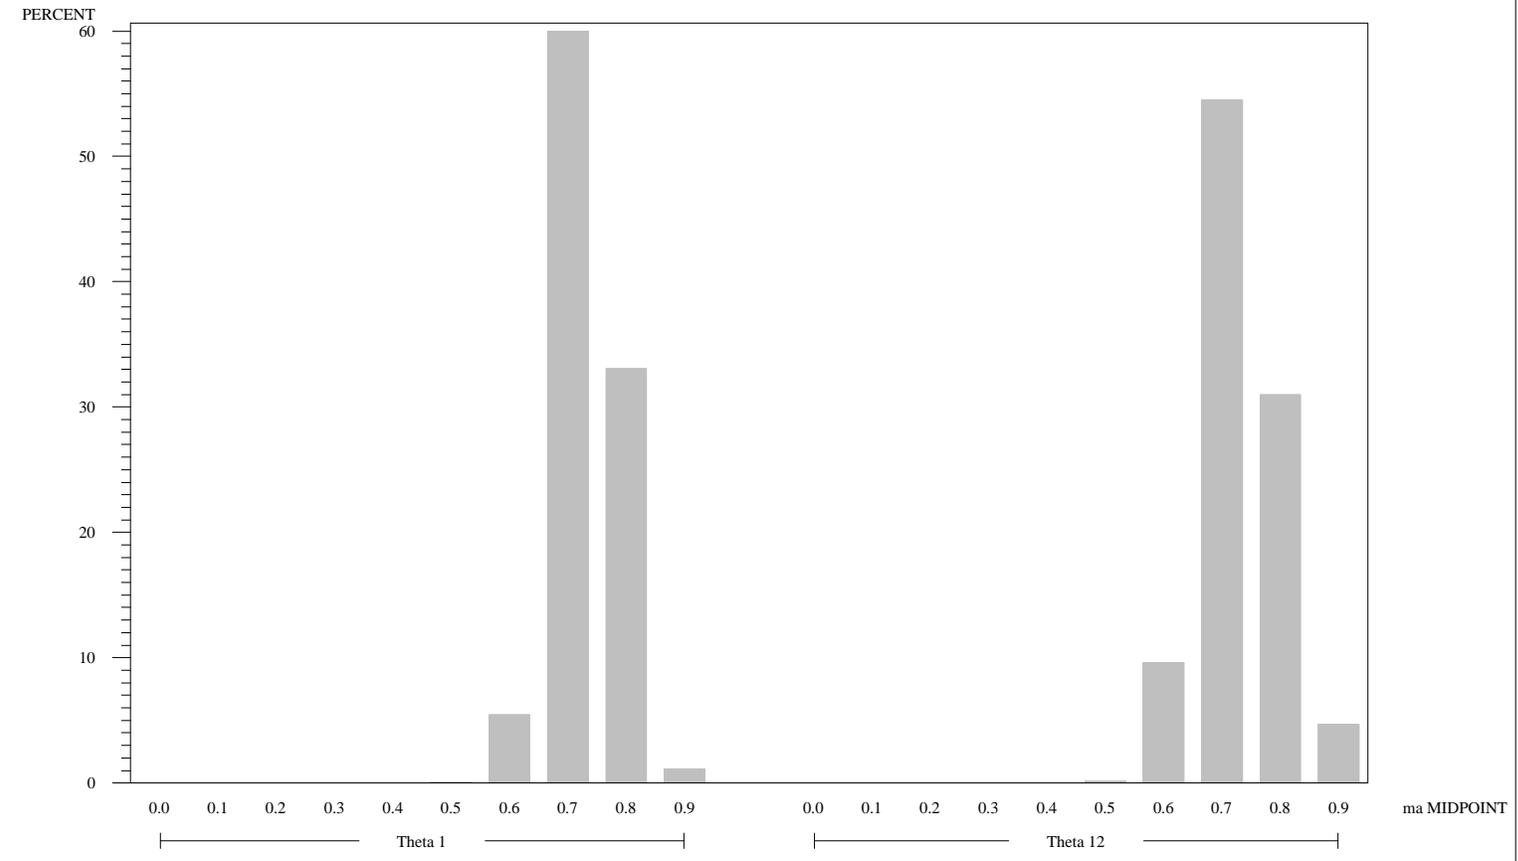


Figure 2. Distribution of MA Parameters - Observed Series
a. 18-Year Series



b. 10-Year Series

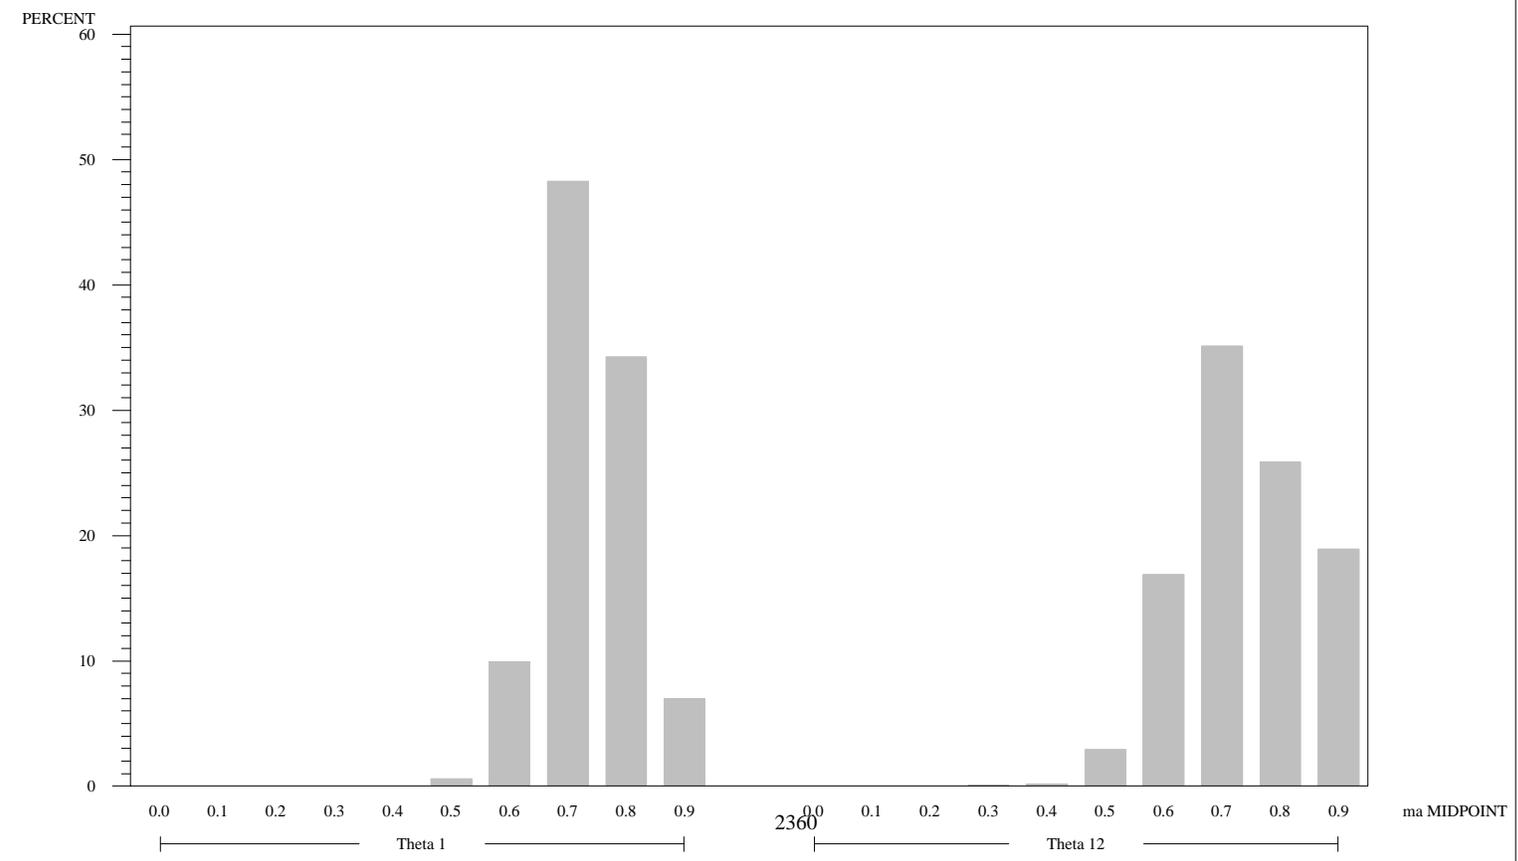
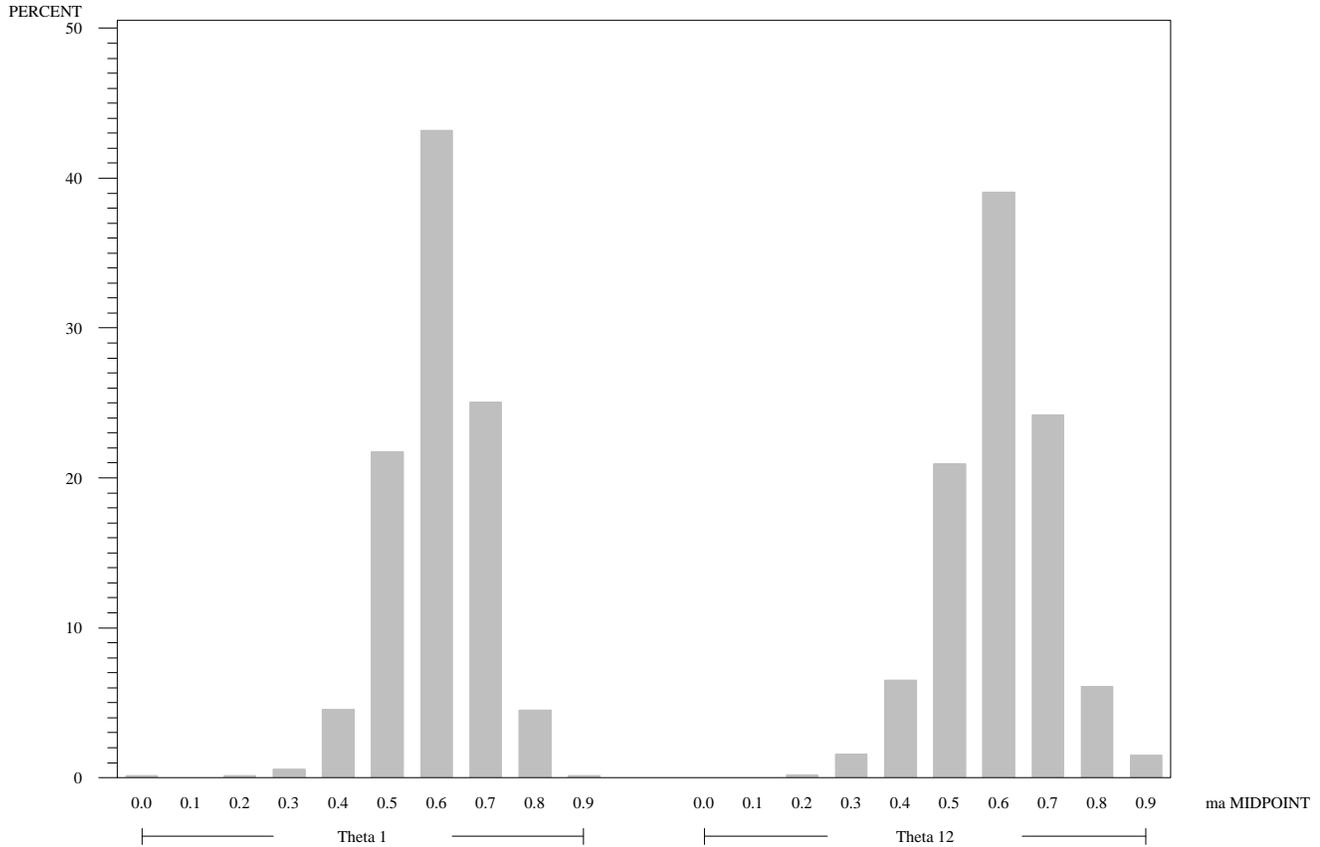


Figure 3. Distribution of MA Parameters - Signal Extraction

a. 18-Year Series, SE Var = 100



b. 10-Year Series, SE Var = 100

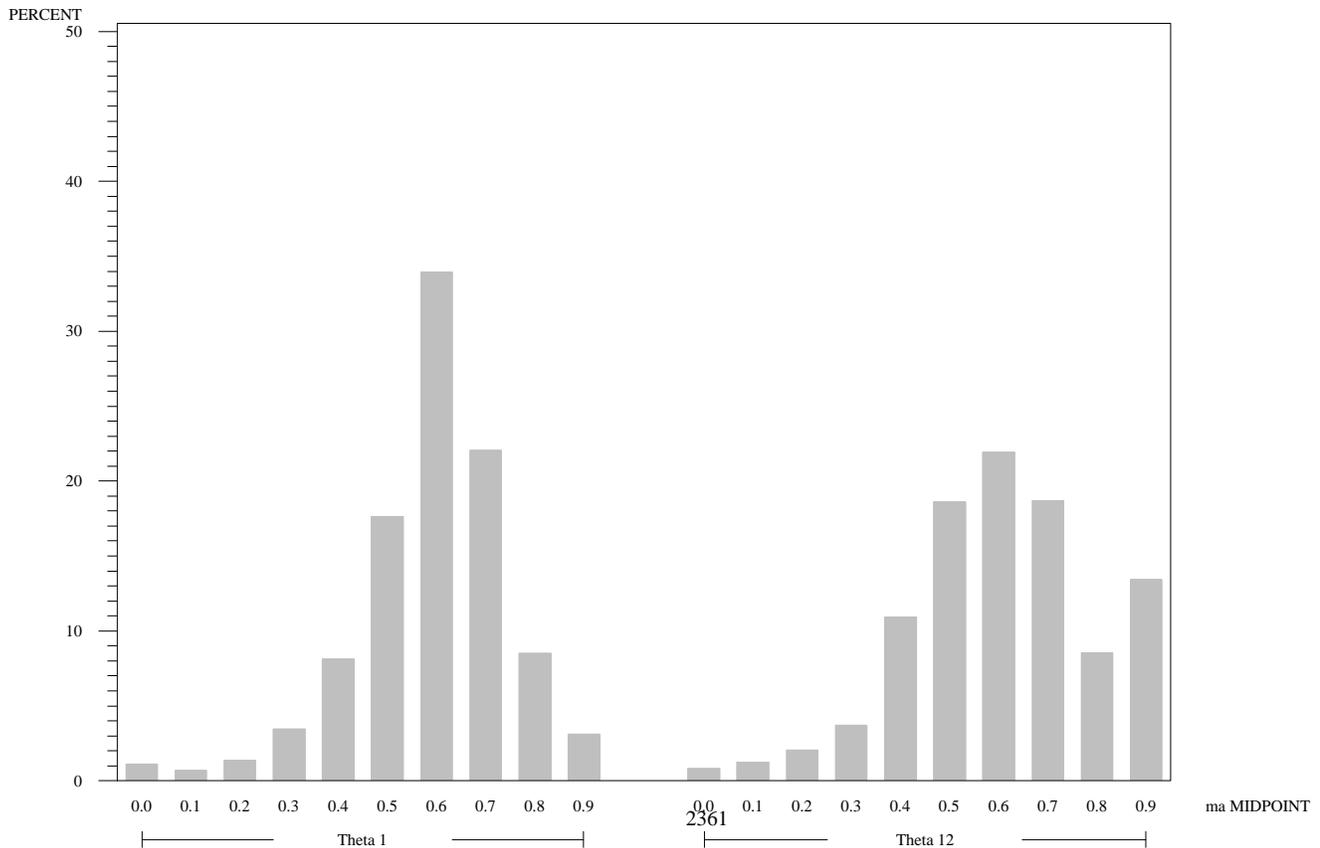
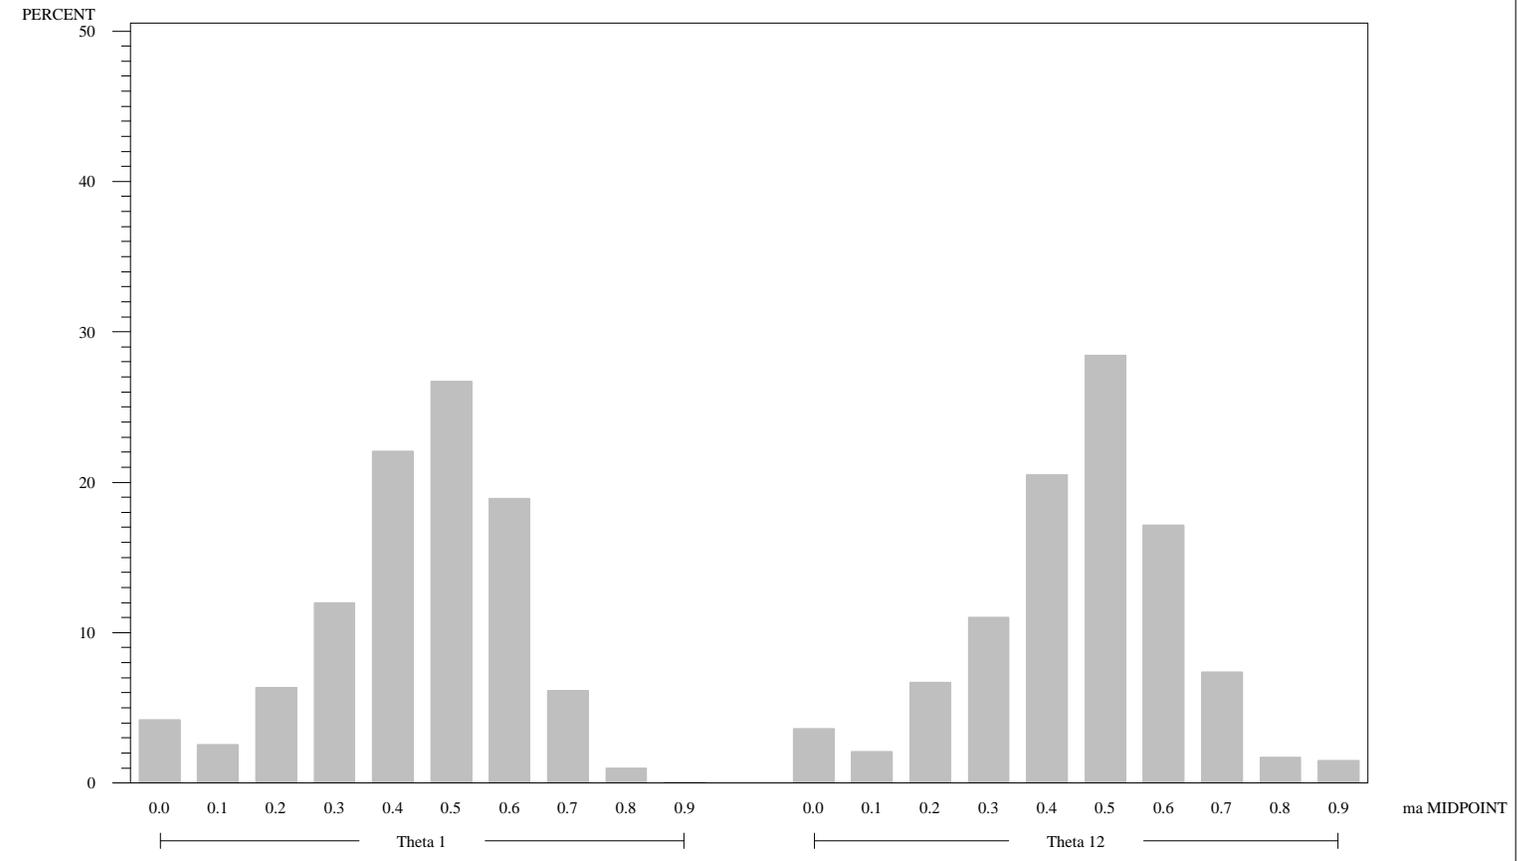


Figure 4. Distribution of MA Parameters - Signal Extraction
 a. 18-Year Series, SE Var = 144



b. 10-Year Series, SE Var = 144

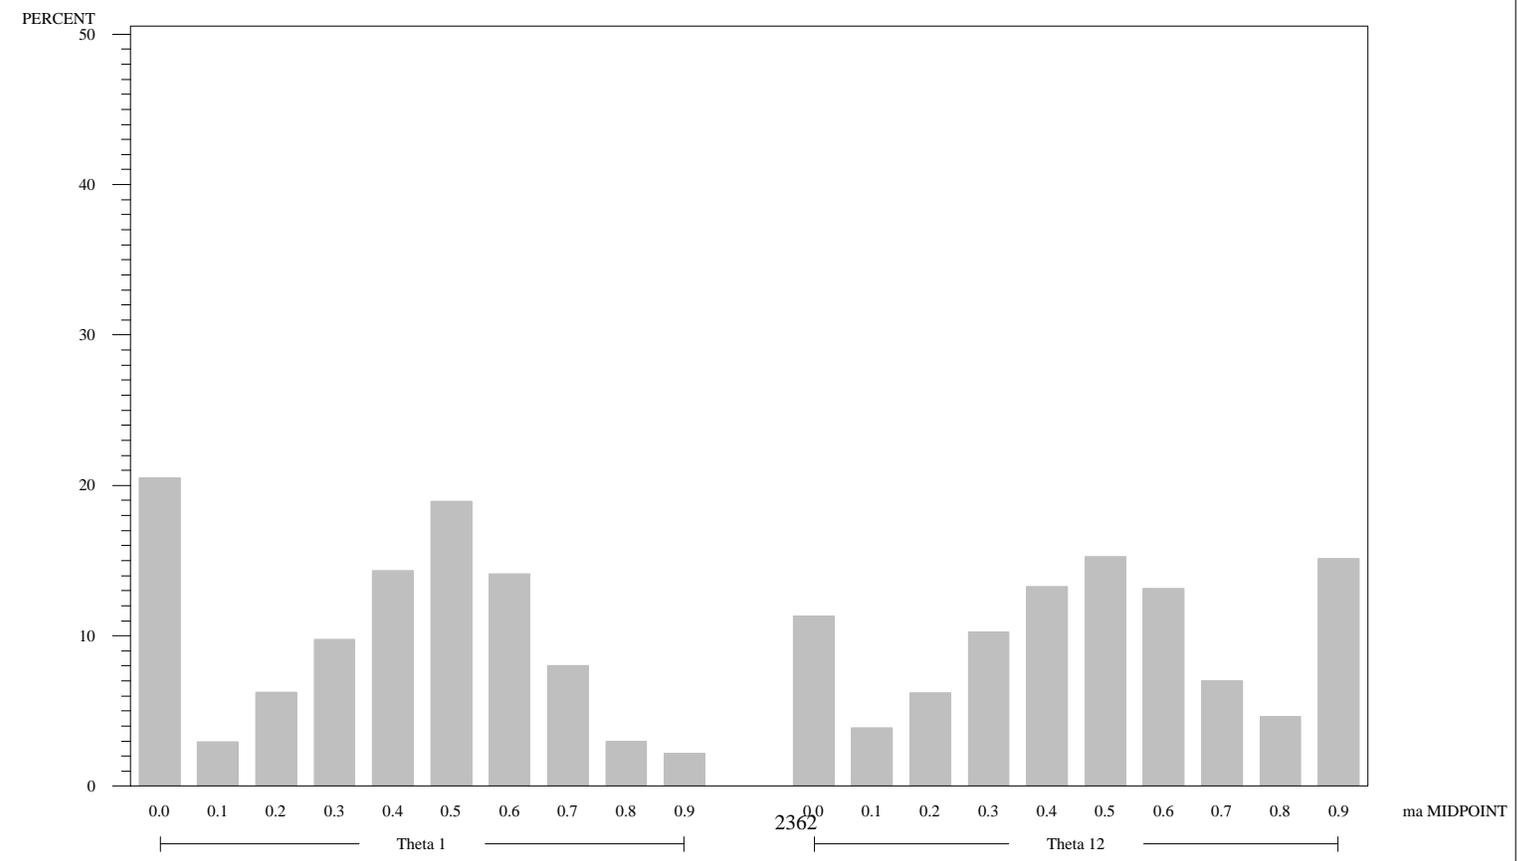
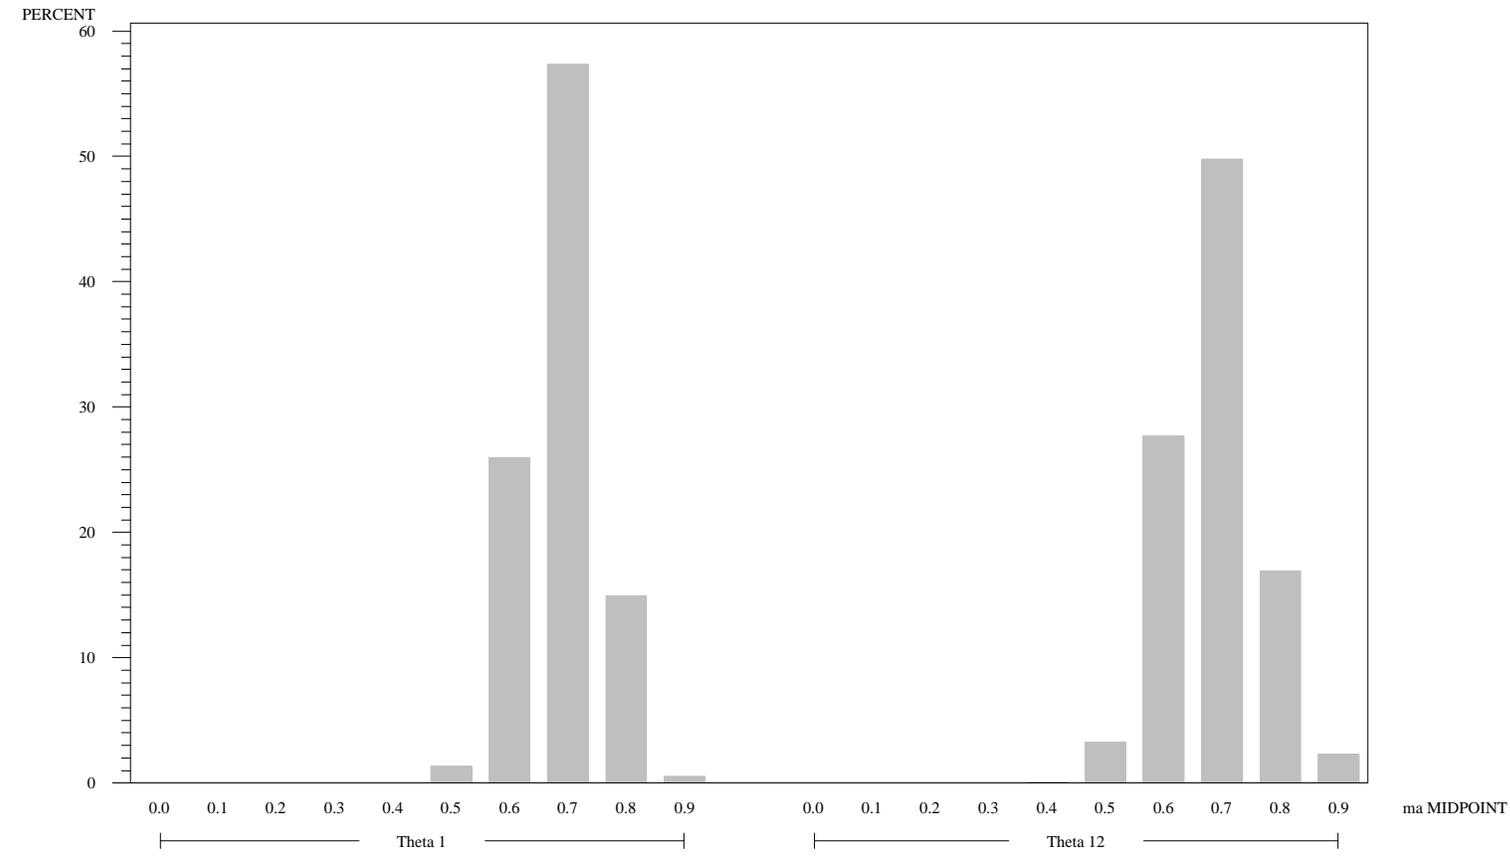


Figure 5. Distribution of MA Parameters - Signal Extraction
 a. 18-Year Series, SE Var = 49



b. 10-Year Series, SE Var = 49

