On the Stability of Least Squares Learning*

Chetan Dave† James Feigenbaum‡

NYU (Abu Dhabi) Utah State University

October 9, 2010

Abstract

Econometricians choose estimators to minimize their loss function defined over estimation errors. In a model of inflation-unemployment tradeoffs, the econometrician's loss function ought to derive from the policymaker's loss function. We consider the implications for adaptive learning of a variance-adjusted least squares estimator that treats overestimates and underestimates asymmetrically. Least-squares learning is close to a ridge line of the expected loss function within this family of estimators. Asymmetric estimators can lower the average inflation rate and the policymaker's expected loss function either by reducing long run inflation or by increasing the frequency of escapes to a low-inflation outcome.

*We would like to thank John Duffy, the referee and participants at the 2005 Southern Economics Association meetings, the Federal Reserve Bank of Dallas and the Computational Economics and Finance 2010 Conference for insightful comments. All remaining errors are ours.

†Address: New York University, Department of Economics, 19 W. 4th Street, 6FL, New York, NY, 10012, USA. E-mail: cdave@nyu.edu.

‡Address: Utah State University, Jon M. Huntsman School of Business, 3565 Old Main Hill, Logan, UT., 84322. Email: j.feigen@aggiemail.usu.edu.

1 Introduction

Dynamic stochastic macroeconomic models can produce multiple outcomes depending on the equilibrium concept employed. In particular, following Kydland and Prescott (1977) and Barro and Gordon (1983), models in which a central bank sets monetary policy over time when facing a public with rational expectations can deliver one of two outcomes. The Nash equilibrium concept delivers a time-consistent high inflation outcome, while the Ramsey equilibrium concept delivers a time-inconsistent low inflation outcome. Assuming that the central bank learns the latent parameters of the structural Phillips curve by econometrically estimating a 'perceived' (and possibly misspecified) Phillips curve from the entire time series of data, it is possible to state a stability condition whereby the Nash outcome is selected as the one that is 'learnable'.[1] On the other hand, if more weight is given to recent data, Sargent (1999) and Cho et al (2002) have shown that an economy may occasionally 'escape' for brief periods to the Ramsey outcome.

The conventional assumption underlying these studies regarding the interaction between policy and data-inference is that the hypothesized econometrics is conducted using some form of a recursive least-squares estimator. In this paper we investigate what happens if the policy maker learns using an alternative estimator, for there are infinitely many estimators available. Least squares is usually employed due to its relative simplicity, but an econometrician's primary concern when choosing the estimator should be to make estimates that ultimately minimize the policymaker's loss function (see Berger (1985)).

We focus on the family of variance-adjusted least-squares estimators introduced by Varian (1975). The least-squares estimator is optimal for a loss function that is symmetric with respect to over- and underestimates of model parameters. However, if estimates in one direction are more costly than in the other direction, the optimal estimator will bias estimates in the less costly direction by an amount proportional to the variance of the least-squares

[1] See Evans and Honkapohja (1999, 2001).

estimate.[2] The family of variance-adjusted least-squares estimators is then parameterized by the proportionality constant. Conventional least squares is nested within this family but, we find that it is quite close to a ridge line of the policymaker's realized loss function. Thus if the policymaker's econometrician experiments with alternate estimators, he will quickly move away from least squares to an estimator that overestimates the (absolute) slope of the Phillips curve and underestimates the intercept, which produces better outcomes for the policymaker.

With respect to our results, the intuition for escape dynamics presented in Cho et al (2002) is particularly appealing and provides the starting point for the present analysis. Specifically, under escape dynamics, a sequence of unusual shocks causes a policy maker to lower the estimate of the trade-off between unemployment and inflation. As the trade-off diminishes, the policy maker responds by lowering the inflation target. This response produces data confirming that the trade-off has indeed diminished, and, in fact, suggests that the trade-off is even smaller than originally thought. The policy maker then further reduces the inflation target and the process continues until inflation is driven to zero and the economy replicates the Ramsey outcome. Once the targeted level of inflation is at zero, the negative correlation between inflation and unemployment becomes apparent again. Witnessing this return through their recursive symmetric least squares learning process, the government begins ratcheting up inflation, and the economy heads back toward the high-inflation Nash outcome.

Within the larger class of variance-adjusted estimators, we can generalize the above mechanism of Cho et al (2002). Most of the time, the economy is in a state of high inflation, the level of which we refer to as long run inflation. Sporadically, an econometrician with a variance-adjusted least-squares estimator will see data that produces an estimate close to zero for the Phillips curve slope, which leads the economy to escape to a state of low inflation.

[2] This mechanism is quite similar to the precautionary-saving mechanism of Leland (1968) and Sandmo (1970), which also depends on an asymmetric valuation of high consumption shocks versus low consumption shocks in the utility function.

Afterwards, the slope estimates gradually return to their average level and inflation returns to its long run rate.

Now, if the econometrician biases his slope estimates upward, this reduces the long run inflation since he believes a small increase in inflation will produce a big decrease in unemployment. It also reduces the frequency of escapes since zero slope estimates will be more infrequent. In contrast, if our econometrician biases his slope estimates downward, this increases the long run level of inflation and increases the frequency of escapes. Within the neighborhood of zero slope bias, i.e. conventional least squares, this will increase the policymaker's loss function. However, the loss function exhibits a non-monotonic dependence on the size of the bias. For a sufficiently large downward bias, the escape frequency will be so high that the economy hardly returns to its high long run level, and average inflation will be very low. With a large enough bias in either direction, the policymaker can achieve average outcomes close to the Ramsey outcome. This can occur with a smaller slope bias if, in addition, the econometrician biases down the intercept estimate, which also leads the policymaker to target inflation at a lower level.

Note that, while we are introducing new parameters, we are not modifying the economic model that delivers the time-inconsistency tension, so this does not cut against Occam's Razor. The more complex dynamics we present are not obtained by expanding the economic framework; we are focused on examining parameter values for the tools employed to analyze the model, rather than parameters of the model itself. Indeed, the space of possible time-series estimators is always infinite-dimensional, regardless of whether an econometrician ties his hands by only considering the least-squares estimator in this space. With the asymmetric (or 'precautionary') least squares learning described here, the econometrician is efficiently processing the data that comes out of the model given that the symmetric least squares case is nested within the broader framework we consider. The result is that least-squares learning is unstable to any experimentation with other estimators. Indeed allowing the econometrician to broaden the class of estimators may provide central banks an escape from the tragedy of

inevitable high inflation described in Sargent (1999).

Within the broader learning literature, it is not uncommon to examine the effects of alternate learning hypotheses such as ours. For example, Evans and Ramey (1998) examine the notion of calculation costs associated with updating in a Lucas (1972) model. We simply investigate the effects of asymmetric (or 'precautionary') least squares (which nests the conventional estimator) given that once one allows for expectations to be replaced with regressions, examining the effects of alternate estimators is of interest.

The paper is structured as follows. In Section 2 we introduce the variance-adjusted least squares estimator. Section 3 provides detail with respect to the economic environment in which learning takes place both in the standard case and in the case that the econometrician employs a variance-adjusted least squares estimator. We argue that the policy-maker, in the time inconsistency environment, is comfortable receiving the variance-adjusted least squares estimates that feed into policy-making as those estimates help lower the economic loss to the policy-maker. An almost willful ignorance of the estimates being received from the econometrician serves to lower long run inflation, or increase the frequency of escapes, and thus lowers the loss experienced by the policy-maker.[3] In Section 4 we show simulation evidence that generalizes the intuition of Cho et al. (2002) as discussed above. Section 5 concludes.

2 The Variance-Adjusted Least Squares Estimator

Following Berger (1985), let us suppose that a statistical decision maker (SDM) estimating a parameter θ in the space $\Theta \subseteq \mathbf{R}^m$ has a loss function $\tilde{L}(\theta, \widehat{\theta})$, where $\widehat{\theta}$ is his estimate of θ. A rational SDM will then specify an estimator $\widehat{\theta}(y)$, which is a function of the data y,

[3] While modeling the broader potential principal-agent tension between the policy-maker and the econometrician who delivers the estimates is beyond the scope of this paper, we do note that the notion of willful ignorance is not entirely without theoretical precedence, see Carrillo and Mariotti (2000).

with the objective of minimizing the expected loss function

$$E\left[\tilde{L}(\theta,\widehat{\theta}(y))\right], \tag{1}$$

where the expectation is taken over y.[4]

For now, let us assume for a given θ that $\tilde{L}$ is minimized when $\widehat{\theta}=\theta$, so the SDM has an incentive to find the truth.[5] If $\tilde{L}$ is C^4, then

$$\frac{\partial\tilde{L}}{\partial\widehat{\theta}_i}(\theta,\theta)=0 \tag{2}$$

for $i=1,...,m$, and

$$\begin{aligned}\tilde{L}(\theta,\widehat{\theta}) &= \tilde{L}(\theta,\theta)+\frac{1}{2}\sum_{i=1}^{m}\sum_{j=1}^{m}\frac{\partial^2\tilde{L}}{\partial\widehat{\theta}_i\widehat{\theta}_j}(\theta,\theta)(\widehat{\theta}_i-\theta_i)(\widehat{\theta}_j-\theta_j)\\ &\quad+\frac{1}{6}\sum_{i=1}^{m}\sum_{j=1}^{m}\sum_{k=1}^{m}\frac{\partial^3\tilde{L}}{\partial\widehat{\theta}_i\widehat{\theta}_j\widehat{\theta}_k}(\theta,\theta)(\widehat{\theta}_i-\theta_i)(\widehat{\theta}_j-\theta_j)(\widehat{\theta}_k-\theta_k)+O\left(\left\|\widehat{\theta}-\theta\right\|^4\right).\end{aligned} \tag{3}$$

The first-order condition for the SDM's problem is then

$$0\approx\sum_{j=1}^{m}\frac{\partial^2\tilde{L}}{\partial\widehat{\theta}_i\widehat{\theta}_j}(\theta,\theta)(E[\widehat{\theta}_j(y)]-\theta_j)+\frac{1}{2}\sum_{j=1}^{m}\sum_{k=1}^{m}\frac{\partial^3\tilde{L}}{\partial\widehat{\theta}_i\widehat{\theta}_j\widehat{\theta}_k}(\theta,\theta)E\left[(\widehat{\theta}_j(y)-\theta_j)(\widehat{\theta}_k(y)-\theta_k)\right]. \tag{4}$$

Typically one assumes third derivatives of $\tilde{L}$ at $\widehat{\theta}=\theta$ are negligible, in which case (4) simplifies to

$$\sum_{j=1}^{m}\frac{\partial^2\tilde{L}}{\partial\widehat{\theta}_i\widehat{\theta}_j}(\theta,\theta)(E[\widehat{\theta}_j(y)]-\theta_j)=0. \tag{5}$$

If the Hessian

$$H_{ij}=\frac{\partial^2\tilde{L}}{\partial\widehat{\theta}_i\widehat{\theta}_j}(\theta,\theta) \tag{6}$$

[4] We assume the SDM follows the classical or frequentist paradigm.

[5] In our inflation-unemployment model, the loss function will not be minimized for $\widehat{\theta}=\theta$ because of time inconsistency, but presumably our naive econometrician does not know this, else he would not employ a least-squares estimator.

of $\tilde{L}$ is nonsingular at $\widehat{\theta} = \theta$, as it will be if H is positive definite (so $\tilde{L}$ is strictly concave), then $E[\widehat{\theta}(y)] = \theta$ is the unique solution to (5), and the loss function will only be minimized for an unbiased estimator, including a least-squares estimator.

In general, however, the third derivatives of $\tilde{L}$ at $\widehat{\theta} = \theta$ will be non negligible. Let

$$\Omega_{jk} = E\left[(\widehat{\theta}_j(y) - \theta_j)(\widehat{\theta}_k(y) - \theta_k)\right] \tag{7}$$

for $j, k = 1, ..., m$. Then a solution to (4) in this wider case must satisfy

$$E[\widehat{\theta}_i(y)] = \theta - \frac{1}{2}\sum_{j=1}^{m}\sum_{k=1}^{m}\sum_{l=1}^{m} H_{ij}^{-1}\frac{\partial^3 \tilde{L}}{\partial\widehat{\theta}_j\widehat{\theta}_k\widehat{\theta}_l}(\theta,\theta)\Omega_{kl} \tag{8}$$

for $i = 1, ..., m$. Thus a rational SDM will choose a biased estimator. This reflects the fact that the loss function increases more sharply if the estimator is wrong in one direction than in the opposite direction. The bias in (8) is proportional to the variance of the estimator $\widehat{\theta}$ relative to the true value θ since the mean squared error of the estimate is given by Ω.

Suppose that $\widehat{\theta}^u$ is an unbiased estimator with variance-covariance matrix Ω^u. Then (8) suggests the SDM ought to replace the unbiased estimator with the variance-adjusted estimator defined by

$$\widehat{\theta}_i^v(y) = \widehat{\theta}_i^u(y) - \frac{1}{2}A_{ij}\Omega_{jk}^u, \tag{9}$$

where A_{ij} is an appropriate estimate of

$$A_{ij} = \sum_{j=1}^{m}\sum_{k=1}^{m}\sum_{l=1}^{m} H_{ij}^{-1}\frac{\partial^3 \tilde{L}}{\partial\widehat{\theta}_j\widehat{\theta}_k\widehat{\theta}_l}(\theta,\theta). \tag{10}$$

For the special case of Varian's (1975) LINEX loss function, the variance-adjusted estimator will minimize the loss function exactly (see also Zellner (1976)).

3 A Monetary-Policy Environment

3.1 The Model

As an example of how an SDM would choose an estimator in a learning environment, we consider the following model of inflation and unemployment. Let u_t represent unemployment and π_t represent inflation. The monetary authority chooses its target inflation rate x_t to minimize the loss function

$$L(x_t) = E[\pi_t^2 + \alpha u_t^2], \tag{11}$$

where u_t is governed by

$$u_t = u_{NR} - \omega(\pi_t - \widehat{x}_t) + \sigma_1 W_{1t} \tag{12}$$

and

$$\pi_t = x_t + \sigma_2 W_{2t}. \tag{13}$$

The shocks $(W_{1t}, W_{2t})\prime$ are i.i.d. and normally distributed with zero means and identity covariance matrices. The standard deviations σ_1, σ_2 are positive. In the Phillips curve (12), $\widehat{x}_t$ is the private sector's expectation of inflation at t and the parameters $u_{NR} > 0$ and $\omega > 0$ are respectively the natural rate of unemployment and the slope.

3.2 Equilibria Under Complete Information

Prior to discussing adaptive learning mechanisms, and our innovation to those mechanisms, we first describe the outcomes that arise when the monetary authority knows the parameters u_{NR} and ω (as in Kydland and Prescott (1977)). If the monetary authority takes $\widehat{x}_t$ as given, it will choose x_t to minimize

$$L_N^{ALM}(x_t) = x_t^2 + \alpha[u_{NR} - \omega(x_t - \widehat{x}_t)]^2 + \alpha\sigma_1^2 + (1 + \alpha\omega^2)\sigma_2^2. \tag{14}$$

The solution for the optimal inflation target is then given by

$$x_t(\widehat{x}_t) = \frac{\alpha\omega}{1+\alpha\omega^2}\left(u_{NR} + \omega\widehat{x}_t\right). \tag{15}$$

Under rational expectations, $\widehat{x}_t = x_t(\widehat{x}_t)$, so the equilibrium policy is

$$x_t^* = \alpha\omega u_{NR}. \tag{16}$$

This corresponds to the time-consistent Nash equilibrium of Kydland and Prescott (1977), for which $E[U_t] = u$, $E[\pi_t] = \alpha\omega u$, and

$$L_N^{ALM}(x_t^*) = \alpha\sigma_1^2 + (1+\alpha\omega^2)(\alpha u_{NR}^2 + \sigma_2^2). \tag{17}$$

If, on the other hand, the monetary authority does not try to exploit the private sector's expectations and assumes the private sector will have rational expectations such that $\widehat{x}_t = x_t$, then the enlightened policymaker's loss function is

$$L_R^{ALM}(x_t) = x_t^2 + \sigma_2^2 + \alpha\left[u_{NR}^2 + \sigma_1^2 + \omega^2\sigma_2^2\right]. \tag{18}$$

The policymaker sees no tradeoff between inflation and unemployment, so the optimal policy is $x_t = 0$. In this Ramsey equilibrium, $E[U_t] = u_{NR}$, $E[\pi_t] = 0$, and

$$L_R^{ALM}(0) = \alpha\sigma_1^2 + \alpha u_{NR}^2 + (1+\alpha\omega^2)\sigma_2^2 < L_N^{ALM}(x_t^*). \tag{19}$$

This enlightened policy produces a smaller value for the loss function, but it requires commitment on the part of the monetary authority. If x_t is chosen after $\widehat{x}_t$, an inflation target of

$$x_t(0) = \frac{\alpha\omega u_{NR}}{1+\alpha\omega^2} \tag{20}$$

would produce an even smaller loss

$$L_N^{ALM}(x_t(0)) = \alpha\sigma_1^2 + \frac{\alpha}{1+\alpha\omega^2}u_{NR}^2 + (1+\alpha\omega^2)\sigma_2^2 < L_R^{ALM}(0) \tag{21}$$

if the monetary authority catches the public off guard. Thus the enlightened Ramsey policy is time-inconsistent.

3.3 Adaptive Learning with a Misspecified Model

The literature on adaptive learning relevant to this policy problem assumes that the monetary authority believes in a misspecified version of the Phillips curve:

$$u_t = u_{NR} - \omega\pi_t + \eta_t, \tag{22}$$

where η_t is a mean-zero noise variable uncorrelated with the other variables determined at t. This perceived law of motion (PLM) assumes that unemployment is inversely related to realized inflation, as opposed to the actual law of motion (the ALM in (12)) in which unemployment is inversely related to unexpected inflation. The policymaker then uses estimates of u_{NR} and ω to set a target inflation rate x_t that minimizes (11) subject to the constraint (22). The current estimate of u_{NR} and ω are determined by some function of previous data on inflation and unemployment. What makes this an interesting model is the fact that future estimates of u_{NR} and ω are endogenous, depending on the policymaker's choice of the inflation target and its attendant effect on unemployment via the ALM.

Formally, the policymaker will choose x_t to minimize (11) subject to the constraints (22) and (13). This obtains the loss function

$$L^{PLM}(x_t) = x_t^2 + \alpha(u_{NR} - \omega x_t)^2 + \alpha\left[\omega^2\sigma_2^2 + V[\eta_t^2]\right] + \sigma_2^2. \tag{23}$$

Since $L^{PLM}(x_t)$ only differs by an additive constant from $L_N^{ALM}(x_t)$ when $\widehat{x}_t = 0$, the two

problems have the same solution:

$$x_t^{PLM} = \frac{\alpha\omega_t u_{NR,t}}{1+\alpha\omega_t^2}. \tag{24}$$

Of course, the actual unemployment rate is still governed by (12). We assume that the public has the same information as the policy-maker and that the public knows how the policy-maker makes its decisions. Consequently, we still have $\widehat{x}_t = x_t$.

3.4 Least-Squares Learning

Where we depart from previous work is in the determination of ω and u_{NR}. The existing literature has focused primarily on least-squares estimation or some variation thereof (Carceles-Poveda and Giannitsarou (2007)), where the policy-maker computes a regression to estimate the parameters of the PLM. Under the conventional approach, as in Evans and Honkapohja (2001), equal weight is put on the entire data set. This is consistent with the belief that the Phillips curve is constant over time, so data from, say 100 years ago, are just as informative as data observed today. A Kalman filter is used to construct a recursive formulation of the least-squares estimator.

Let $\xi_t = (u_{NR_t}, \omega_t)$ be the current estimates of the Phillips curve parameters given the data $\{(u_i, \pi_i)\}_{i=0}^{t-1}$, which are given by the least-squares formula

$$\xi_t = (X_t' X_t)^{-1} X_t' U_t, \tag{25}$$

where we define $U_t = (u_0, ..., u_{T-1})'$ and

$$X_t = \begin{bmatrix} 1 & \pi_0 \\ \vdots & \vdots \\ 1 & \pi_{t-1} \end{bmatrix}. \tag{26}$$

Let us also define

$$R_t = \frac{1}{t} X_t' X_t, \tag{27}$$

which satisfies the matrix difference equation

$$R_t = R_{t-1} + \frac{1}{t}\left(x_{t-1}x_{t-1}' - R_{t-1}\right), \tag{28}$$

where $x_t = (1, \pi_t)'$. Inserting (27) into (25), we obtain

$$\xi_t = t^{-1} R_t^{-1}\left((t-1)R_{t-1}\xi_{t-1} + x_{t-1}u_{t-1}\right).$$

Applying (28), this simplifies to

$$\xi_t = \xi_{t-1} + t^{-1} R_t^{-1} x_{t-1}(u_{t-1} - x_{t-1}'\xi_{t-1}). \tag{29}$$

Note that the forecast error $u_{t-1} - x_{t-1}'\xi_{t-1}$ is weighted by the gain factor t^{-1}. Since the data set increases with each new observation, this is an example of a decreasing-gain learning algorithm. Under this learning process, the economy will converge to the Nash equilibrium with its high value for the policymaker's loss function $L_N^{ALM}(x_t^*)$.

Cho et al (2002) consider instead a constant-gain generalization of the least-squares estimator, where each new datum is given the same weight in the Kalman filter. Although procedurally different, this produces results similar to what one would obtain if one computed a regression on only the last T observations for some fixed T. More precisely, we replace (28) and (29) by the difference equations

$$\xi_t = \xi_{t-1} + T^{-1} R_t^{-1} x_{t-1}(u_{t-1} - x_{t-1}'\xi_{t-1}) \tag{30}$$

and

$$R_t = R_{t-1} + \frac{1}{T}\left(x_{t-1}x_{t-1}' - R_{t-1}\right), \tag{31}$$

where T is fixed. This is appropriate if the policymaker is concerned about structural change to the Phillips curve and believes data from 100 years ago is not informative about what is happening today. The constant-gain recursive estimator departs from its decreasing-gain counterpart in that escape dynamics can arise. If by chance the Phillips curve in the last T periods appears flat, the inflation rate can temporarily drop to the low value of the Ramsey equilibrium with its lower loss $L_R^{ALM}(0)$. It is this constant-gain case that we generalize upon in the following section.

3.5 Variance Adjusted Learning

Now let us suppose that our policymaker engages an econometrician who supposes, correctly as it turns out, that estimation errors in one direction might be more costly with respect to the policymaker's loss function (11) than errors in the other direction. Following Varian (1975) and our discussion in Section 2, he considers the larger space of estimators

$$\xi_t^a = \xi_t + \frac{1}{2}\Sigma_t a, \tag{32}$$

where Σ_t is the variance-covariance matrix of the least-squares estimator and $(a_{u_{NR}}, a_\omega)'$ measures the strength of the 'precautionary motive' (or asymmetry) against underestimating the two Phillips curve parameters u_{NR} and ω. If either $a_{u_{NR}}$ or a_ω is negative then the precautionary motive for that variable is against overestimation. For the least-squares estimator, the variance-covariance matrix is

$$\Sigma_t = t^{-1} R_t^{-1} \sigma_t^2, \tag{33}$$

where

$$\sigma_t^2 = \frac{1}{t-1} \sum_{i=0}^{t-1} (u_i - x_i' \xi_t)^2 \tag{34}$$

is the most recent estimate of the variance of η_t. We define

$$S_t = \frac{1}{t}\sum_{i=0}^{t-1} u_i x_i', \tag{35}$$

which satisfies the difference equation

$$S_t = S_{t-1} + \frac{1}{t}(u_{t-1}x_{t-1}' - S_{t-1}). \tag{36}$$

Then we can rewrite σ_t^2 as

$$\sigma_t^2 = \frac{1}{t-1}\sum_{i=0}^{t-1}\left[(u_i - x_i'\xi_{t-1})^2 + 2(u_i - \xi_{t-1}'x_i)x_i'(\xi_{t-1} - \xi_t) + (\xi_{t-1} - \xi_t)\prime x_i x_i'(\xi_{t-1} - \xi_t)\right],$$

which simplifies to

$$\begin{aligned}\sigma_t^2 = \sigma_{t-1}^2 + \frac{1}{t-1}[&(u_{t-1} - x_{t-1}'\xi_{t-1})^2\left(1 + t^{-1}x_{t-1}'R_t^{-1}x_{t+1}\right) \\ &+ 2\left(\xi_{t-1}'R_t - S_t\right)R_t^{-1}x_{t-1}(u_{t-1} - x_{t-1}'\xi_{t-1})].\end{aligned} \tag{37}$$

The preceding derivation is exact and the estimator ξ_t^a with σ_t^2 and S_t defined above is a variance-adjusted generalization of the decreasing gain least-squares estimator described in 3.4. Not surprisingly, the dynamics generated with variance-adjusted decreasing-gain learning algorithm converge to the same Nash equilibrium as the ordinary least-squares algorithm. In the long run, σ_t^2 will converge to some steady-state variance and the precautionary adjustment for the variance will disappear.

To generate interesting dynamics as in Cho et al (2002), we must again use a constant-gain version of the learning algorithm. As in the case of the least-squares learning algorithm, we replace the time t by a fixed horizon T to obtain the difference equations

$$S_t = S_{t-1} + \frac{1}{T}(u_{t-1}x_{t-1}' - S_{t-1}) \tag{38}$$

$$\sigma_t^2 = \sigma_{t-1}^2 + \frac{1}{T-1}[(u_{t-1} - x'_{t-1}\xi_{t-1})^2 \left(1 + T^{-1}x'_{t-1}R_t^{-1}x_{t+1}\right)$$
$$+ 2\left(\xi'_{t-1}R_t - S_t\right) R_t^{-1}x_{t-1}(u_{t-1} - x'_{t-1}\xi_{t-1})], \quad (39)$$

which together with (30) and (39) define the constant-gain version of the variance-adjusted learning algorithm.

4 Simulation Results

4.1 Loss Function Surfaces

Given an environment in which agents adaptively learn using least squares, the usual procedure is to stochastically approximate the system of equations that characterize model dynamics, including the Ricatti equations delivered by recursive least squares regressions. Such an approximation for 'mean' dynamics delivers a system of differential equations that can be examined for stability of model equilibria. Cho et al . (2002) went further and defined escape dynamics as well for the case of recursive symmetric least squares learning under constant gain. Since this analytical procedure is not easily generalized to account for the dynamics of the variance, we instead opt to provide evidence via simulation on whether Ramsey or Nash inflation is obtained under variance-adjusted learning. The simulation results prompt the intuition described above: variance adjusted learning on the part of a monetary authority can lead to a higher frequency of escape and a lowering of the long run average rate of inflation. The key of course to this intuition is the nature of the econometrician's loss function surface and so we begin with three plots of these surfaces in Figures 1, 2 and 3 below.

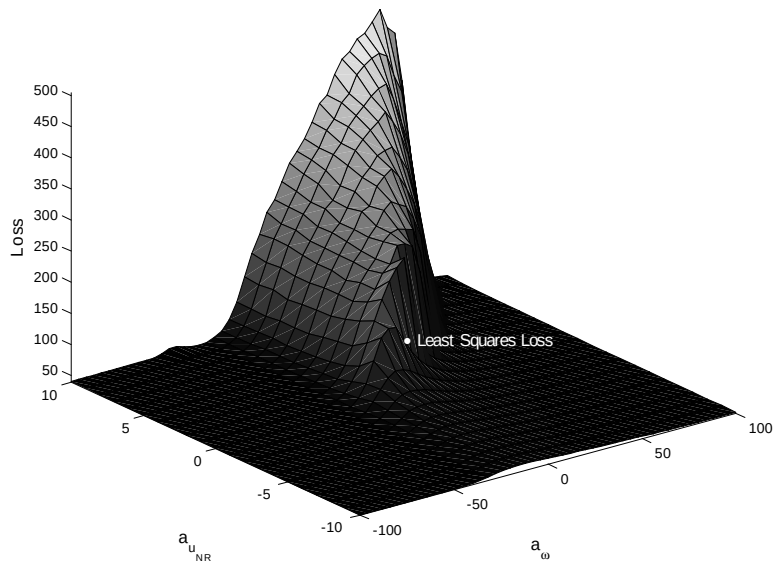

Figure 1. Econometricians' Loss Function (0.05 gain).

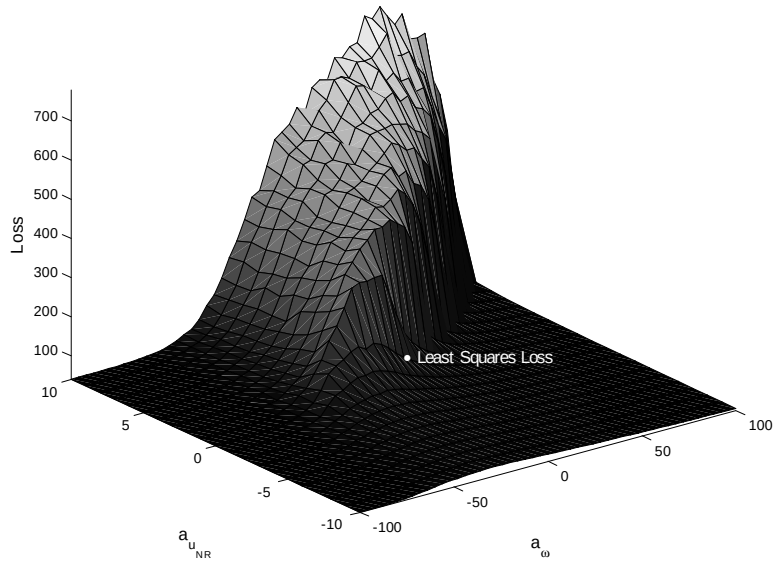

Figure 2. Econometricians' Loss Function (0.0025 gain).

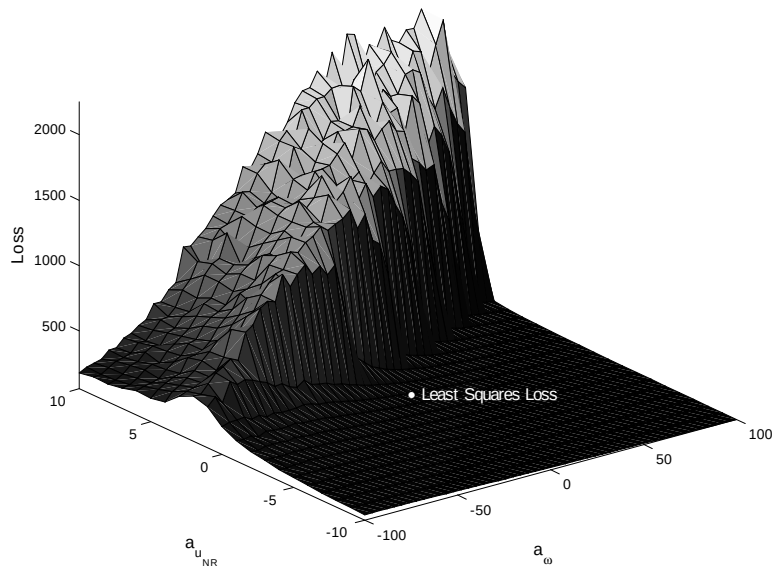

Figure 3. Econometricians' Loss Function (0.1 gain).

We generate each of these figures (and all reported simulations below) having fixed α at 1, ω at 2 and u_{NR} at 5 so that the Nash level of inflation from (16) is 10 while the Ramsey equilibrium inflation rate is, of course, zero. In the self-confirming equilibrium that a decreasing-gain learning algorithm will converge to, the Phillips curve is perceived to have a slope of 2, which is the actual value, and an intercept of 25, which adds to u_{NR} the effect of the public's expectations about inflation. The standard deviations of the shocks, σ_1 and σ_2, are set to 1.5. The econometrician starts each simulation looking at data generated by a predecessor policy maker who targeted the inflation rate at its Nash value.

In our baseline calibration, we set the horizon for the constant gain algorithm to 20 periods, implying a gain of 0.05. This is the loss surface shown in Figure 1, and we also use this calibration for the time series plots below. Recall that a high gain translates into a shorter time span of data employed by the econometrician. Thus the probability of observing a history in which the Phillips curve appears to be flat is high, and consequently the probability of escape is also high. Since we are interested in examining dynamics in the variance-adjusted learning environment a high gain of 0.05 is a natural choice. To show the robustness of our main result we also show the loss surface for the low value of 0.0025 in

Figure 2 and at the even higher value of 0.1 in Figure 3.[6]

In each of Figures 1-3, we see that least-squares learning is close to a ridge line of the more general loss function. If our econometrician experiments by allowing for a small degree of asymmetry in the estimation of either Phillips curve parameter, she will quickly discern that better policy results from overestimating the slope and underestimating the intercept relative to least squares. However, movement away from least squares in this direction is a purely local phenomenon. Whereas slight underestimates of the slope lead to a worsening of policy, large underestimates improve policy, and likewise for overestimates of the intercept. The econometrician will be nearly indifferent between giving the policymaker very high or very low estimates of the parameters as long as she does not give them the least-squares estimates of the parameters.

4.2 Time-Series Simulations

Why do biased estimates of the Phillips curve parameters improve policy? Do they cause the policymaker to target rates of inflation lower than those consistent with a Nash outcome? We investigate this question by examining the time series generated by our model in simulations.

For our baseline parameters, the loss function for the Ramsey equilibrium will average to $L_R^{ALM}(0) = 38.5$ from (19). In contrast, the loss function for the Nash equilibrium will average to the much higher $L_N^{ALM}(x_t^*) = 138.5$ from (17).

[6] Loss function surfaces are generated by averaging across simulations for 10000 periods.

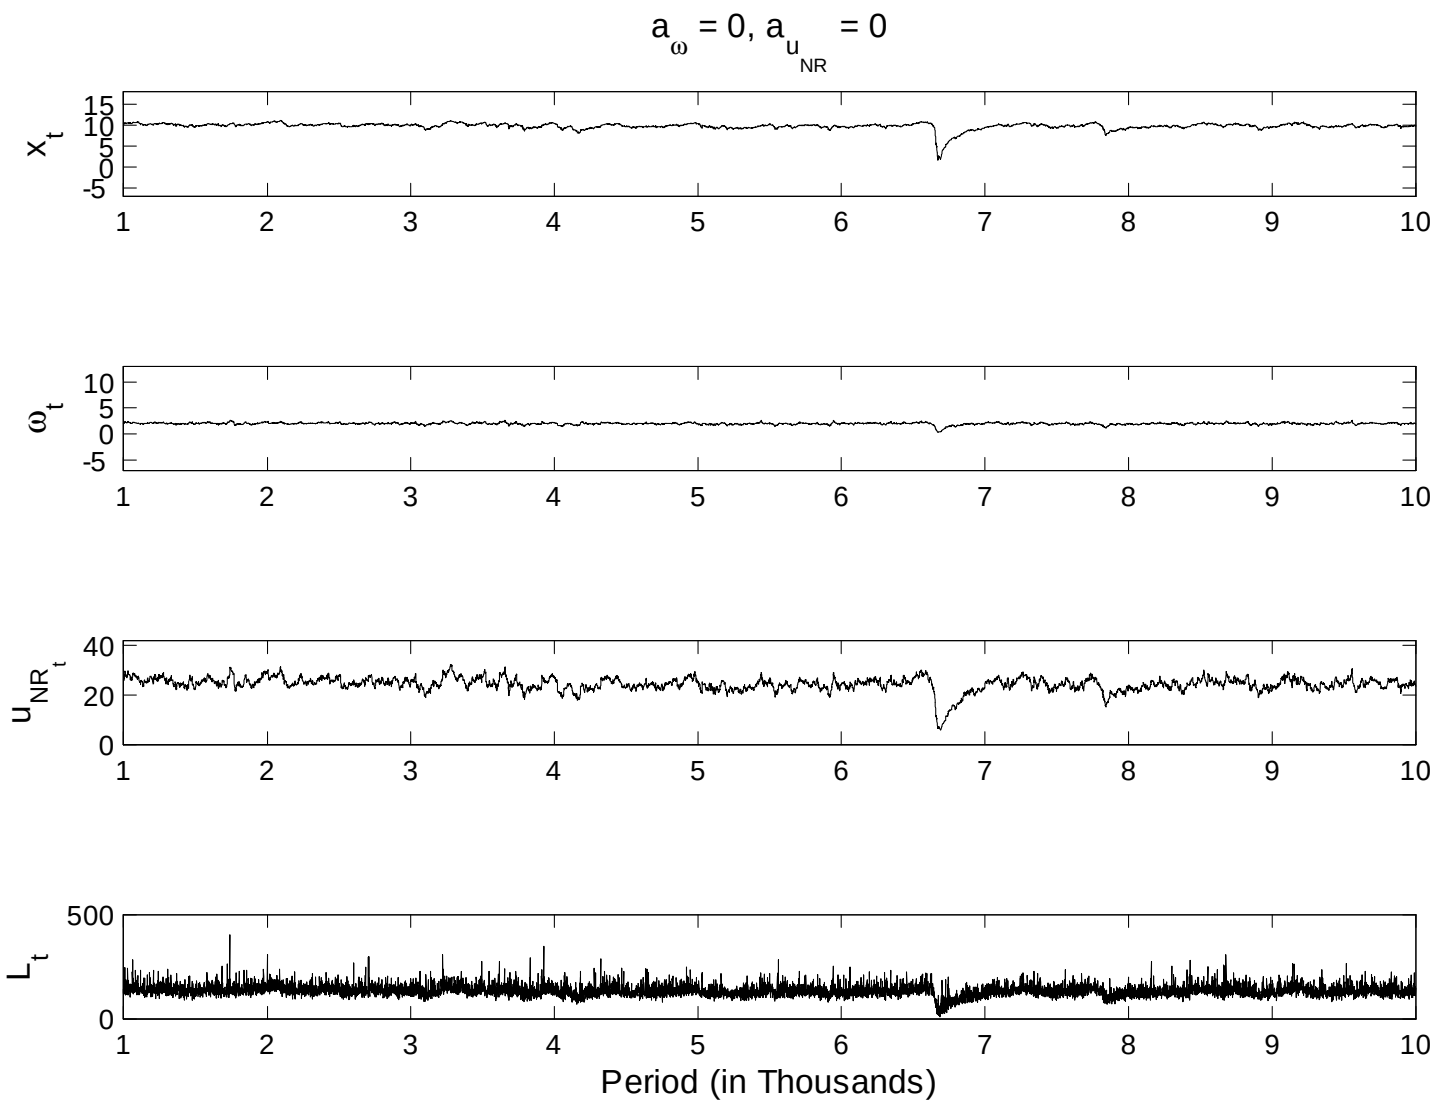

Figure 4. Targeted Inflation, Slope, Intercept and Economic Loss Estimates (Symmetric Least Squares).

First, let us consider what happens for the familiar case of least-squares estimation, where $a = (0, 0)$. The time series for targeted inflation, the estimated Phillips curve parameters, and the period loss function are shown in Figure 4. Consistent with Cho et al. (2002), the inflation target, shown in the top panel, remains around the Nash equilibrium value of 10 most of the time but there is an escape where the target falls close to the Ramsey level of zero. The plots of ω_t and u_{NR_t} in the middle two panels reflect the intuition of Cho et al. (2002) as follows. The average estimate of ω is 1.97 and of u_{NR} is 24.3, both close to the values that would be obtained by regressing u_t on π_t according to (22) in the Nash equilibrium.[7] However, after the 6500th observation a string of data occurs where the correlation of u_t and π_t falls, resulting in smaller estimates of both ω and u_{NR}. The policymaker thinks there is a smaller tradeoff between unemployment and inflation, so according to (24) she reduces target inflation. As a consequence of the lower target, the slope and intercept estimates remain depressed, so it takes time for the system to return to Nash equilibrium values. The period loss function falls to an average value of around 20 during the escape but most of the time is around the Nash equilibrium value of 138.

Intuitively, one might expect the policymaker could achieve better results if he underestimated the slope since that will increase the probability of finding there is no tradeoff between unemployment and inflation, which will produce a Ramsey-like outcome. The time series that result from setting $a_\omega = -10$ are shown in Figure 5. Consistent with our expectation, escapes occur more frequently in this case. However, the average loss function increases to 153. This is because the inflation rate is higher when the economy is not escaping; the policy maker is overshooting.

[7] The misspecification of the model will not bias estimates of the slope but will produce consistent estimates of $u_{NR} + \theta x^*$ rather than u_{NR}.

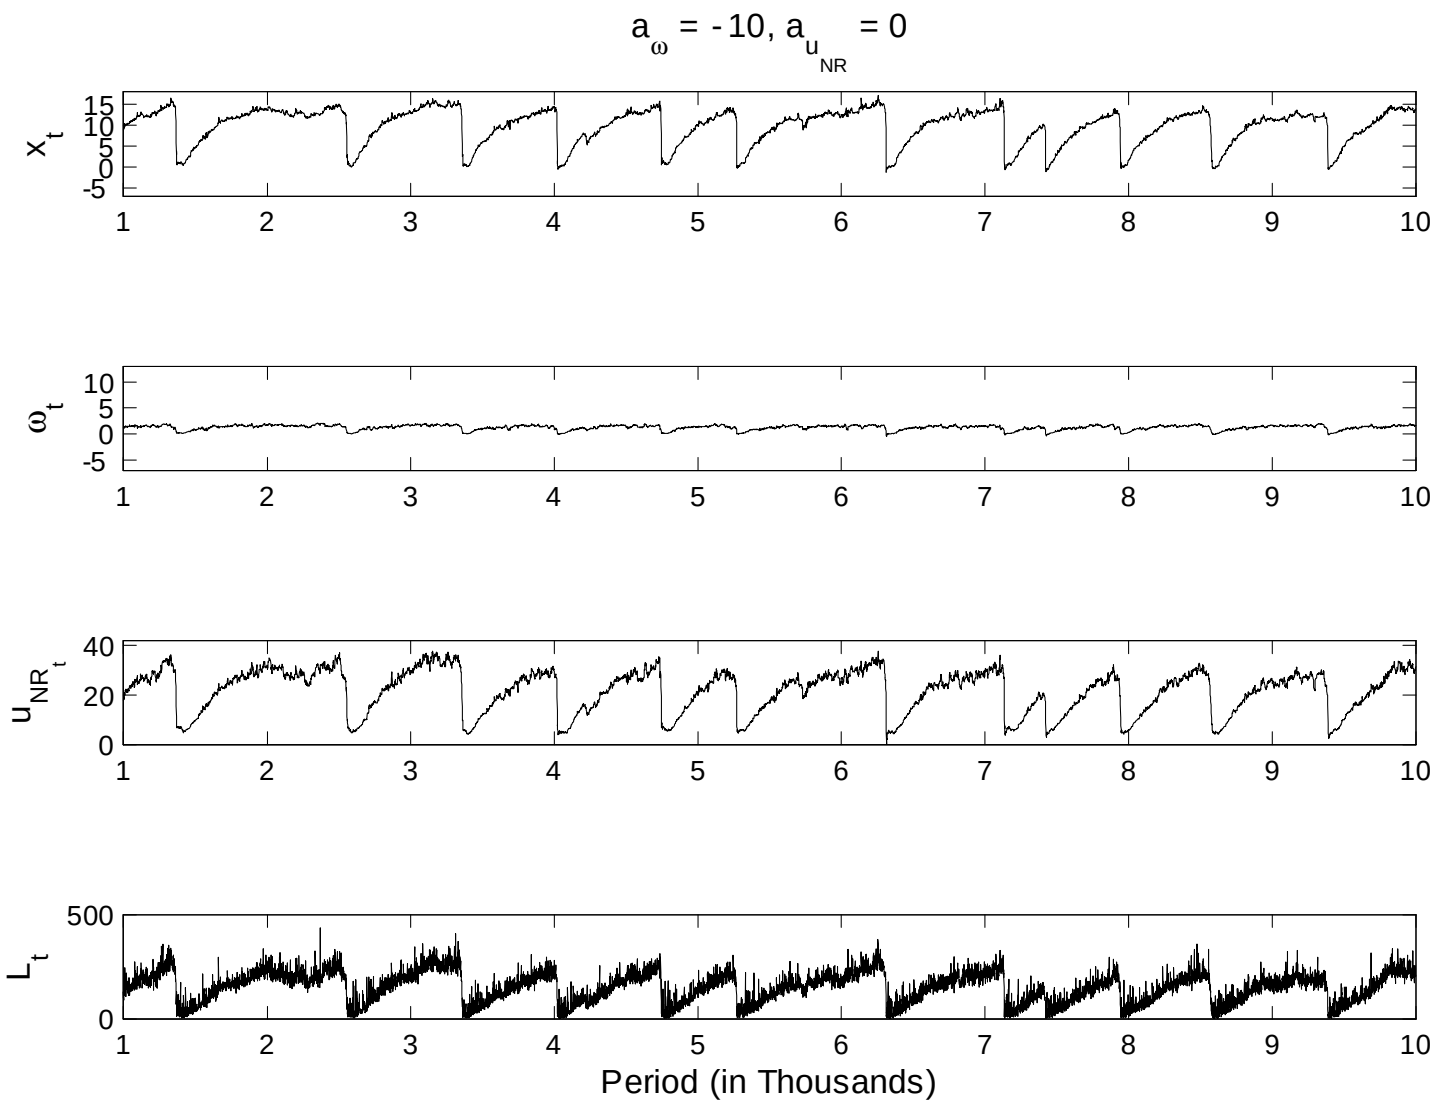

Figure 5. Targeted Inflation, Slope, Intercept and Economic Loss Estimates (Slope Asymmetry −10).

Whether policy will improve if a_ω is increased above zero or decreased below zero depends on the parameters of the model. Differentiating (24) by ω, we find

$$\frac{\partial x_t^{PLM}}{\partial \omega_t} = \frac{\alpha u_{NR,t}(1 - \alpha\omega_t^2)}{(1 + \alpha\omega_t^2)^2}, \tag{40}$$

which is of ambiguous sign. For our parameters, $\alpha\omega^2 = 4 > 1$, where ω is the actual slope of the Phillips curve. Thus a small decrease in the slope estimate will actually increase the target inflation rate at normal times.

In this case, overestimating the slope leads to better policy. To see this, we show time series for $a_\omega = 10$ in Figure 6. Because the policy maker is overestimating the Phillips curve slope, he thinks the tradeoff between unemployment and inflation is so large that he does not need so much inflation to achieve the desired low unemployment. Escapes never happen in this simulation because he never convinces himself the tradeoff has gone away.

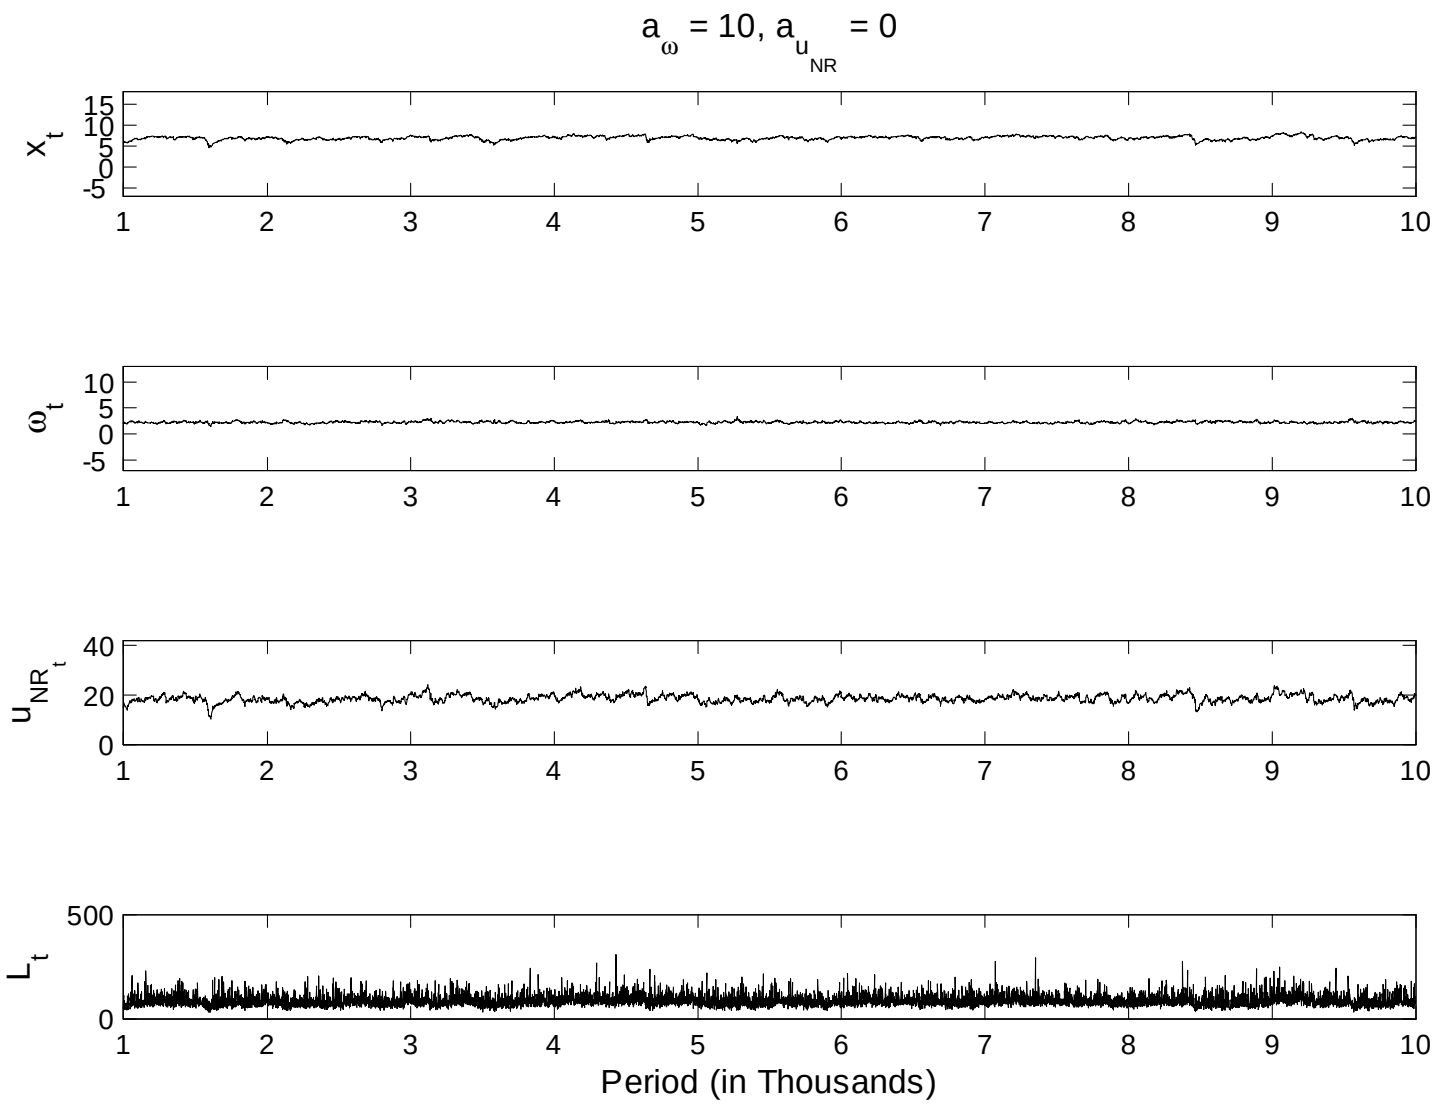

Figure 6. Targeted Inflation, Slope, Intercept and Economic Loss Estimates (Slope Asymmetry +10).

In the bottom panel of Figure 5, we see that the period loss function can fall almost to zero during the many escapes that occur when the Phillips curve slope is underestimated, but these gains are dominated by the higher values around 200 of the period loss function that arises during normal periods when inflation is targeted around 15 instead of 10. In contrast, when the slope is overestimated, the period loss function remains steady with an average of 87 and an average inflation target of 7.

If we continue to increase a_ω, we can achieve additional improvements to policy by reducing the target inflation rate even more. This is shown in Figure 7, where $a_\omega = 20$, target inflation averages to 5.3, and the period loss function averages to 67. If we increase a_ω all the way to 100, as in Figure 8, the average target inflation falls to 1.9, and the period loss function averages to 42, which is nearly the value of the Ramsey equilibrium.[8]

[8] This will be the minimum average value of the loss function over long periods in a model where the public has rational expectations.

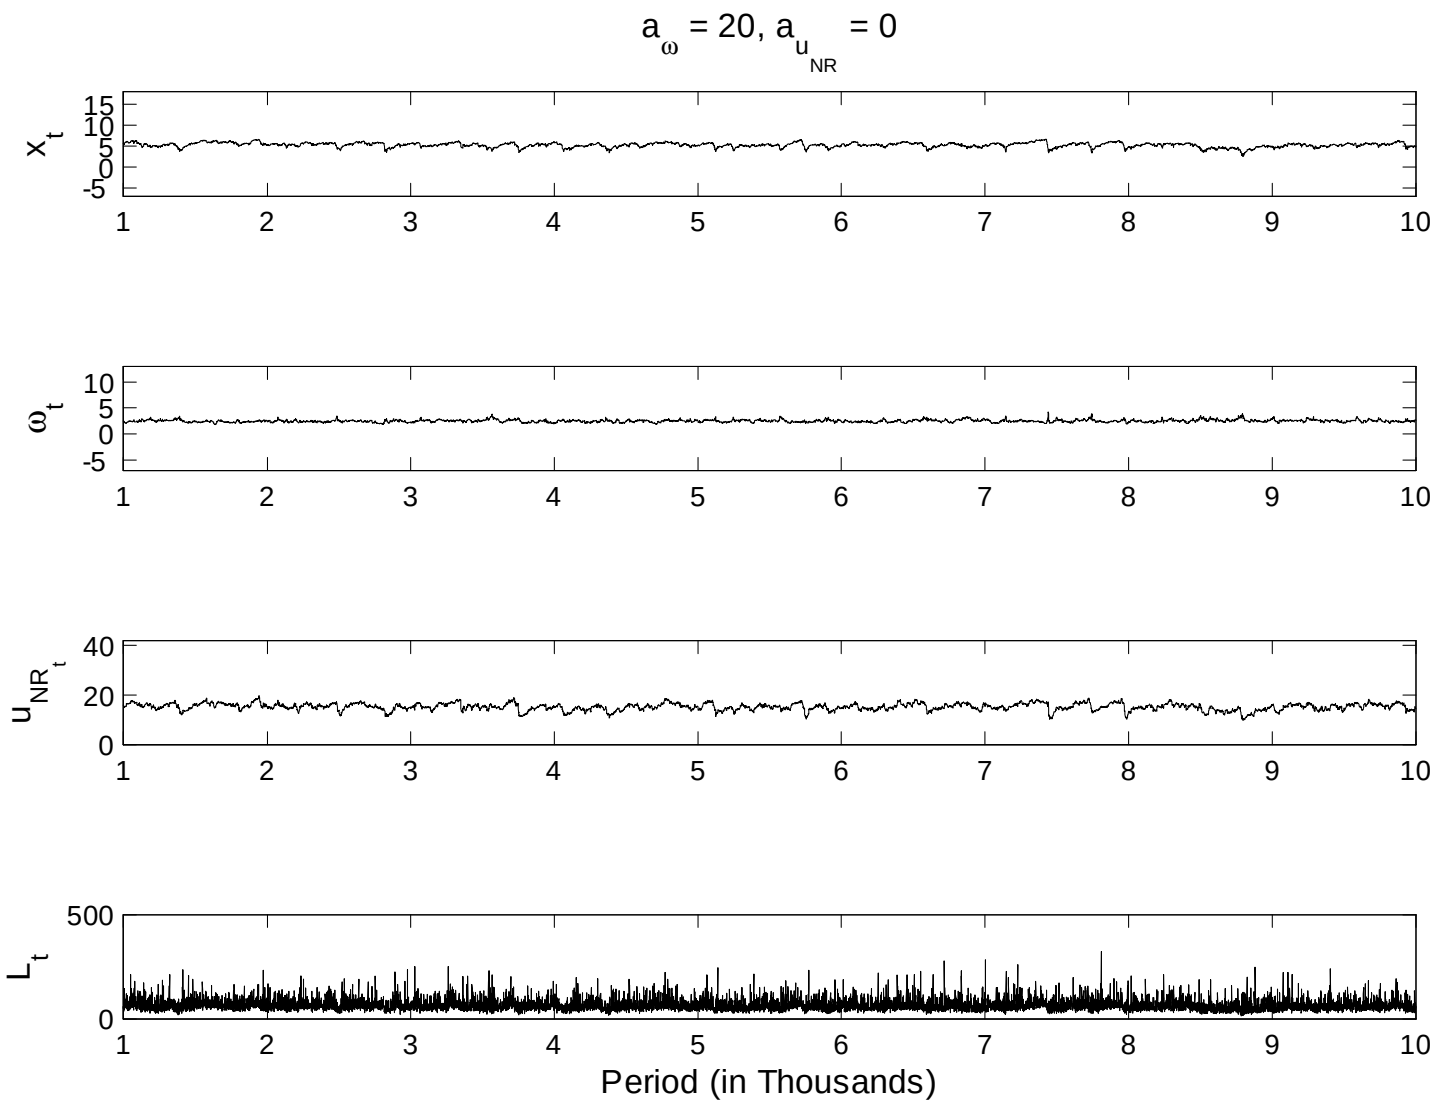

Figure 7. Targeted Inflation, Slope, Intercept and Economic Loss Estimates (Slope Asymmetry +20).

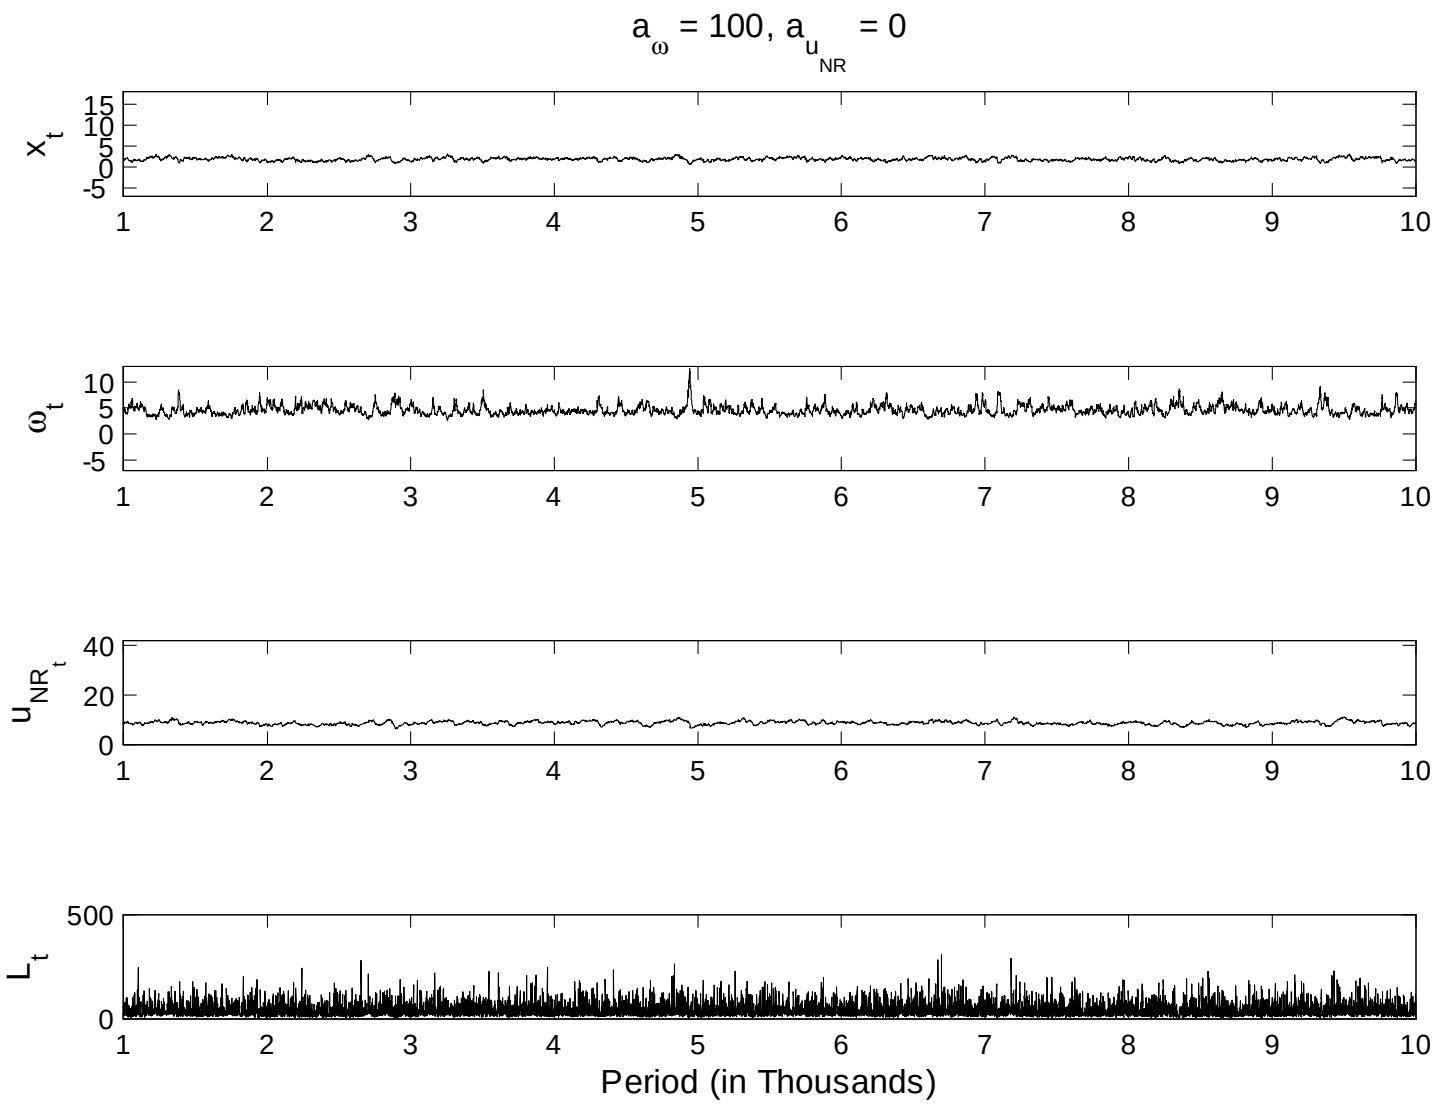

Figure 8. Targeted Inflation, Slope, Intercept and Economic Loss Estimates (Slope Asymmetry +100).

What happens if we continue to vary a_ω in the opposite direction from -10? If $a_\omega = -20$ with $a_{u_{NR}}$ still equal to zero, we will have passed over to the other side of the ridge line in the loss function surface shown in Figure 1. The time series for this case are shown in Figure 9.

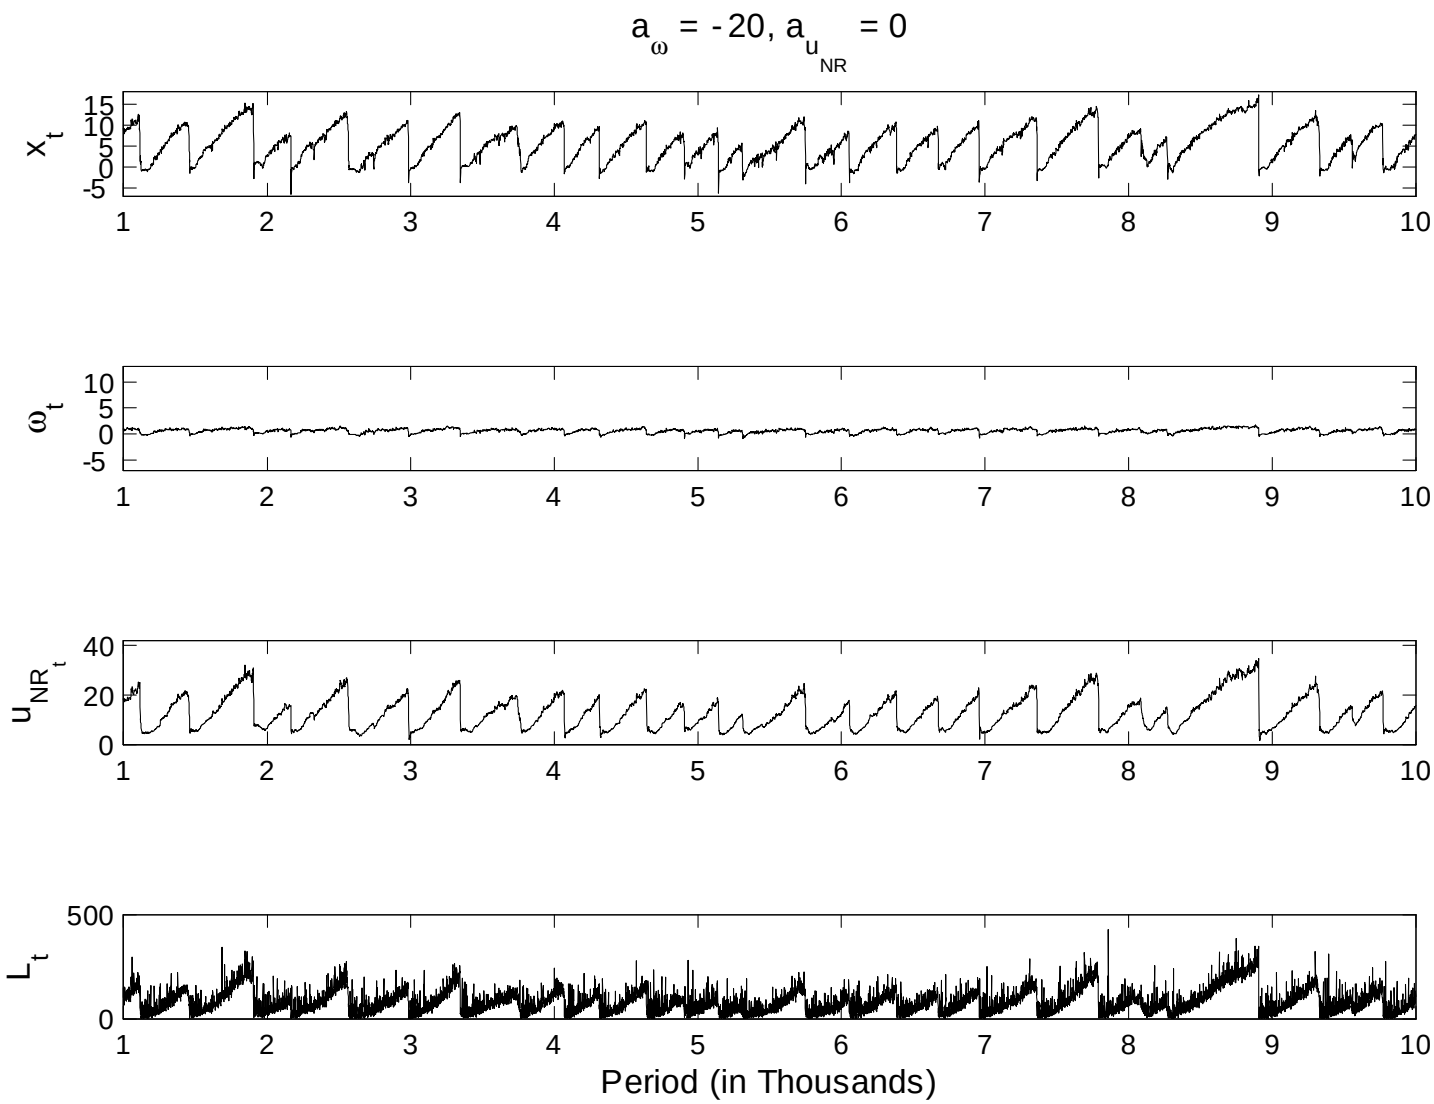

Figure 9. Targeted Inflation, Slope, Intercept and Economic Loss Estimates (Slope Asymmetry −20).

With such a large negative bias, the economy is almost always in an escape. In the rare times when the economy returns to Nash-like behavior, target inflation and the period loss function are very high. But most of the time, the return to Nash-like behavior is preempted by a new escape. Thus the loss function averages only 83. If we decrease a_ω even further to -100, as in Figure 10, the behavior is very similar to what happens with $a_\omega = +100$ (Figure 8).

Target inflation is nearly constant at -1.1 with a tiny standard deviation of 0.4. The period loss function averages to 40, which is even lower than what we obtained with $a_\omega = +100$.

The econometrician can also apply a precautionary bias to estimation of the intercept of the PLM, i.e. the natural rate of unemployment. What happens then? What of variance-adjusted learning on the intercept of the PLM? Figure 11 shows what happens if we overestimate the intercept with $a_{u_{NR}} = 1$ and Figure 12 what happens if we underestimate the intercept with $a_{u_{NR}} = -1$.

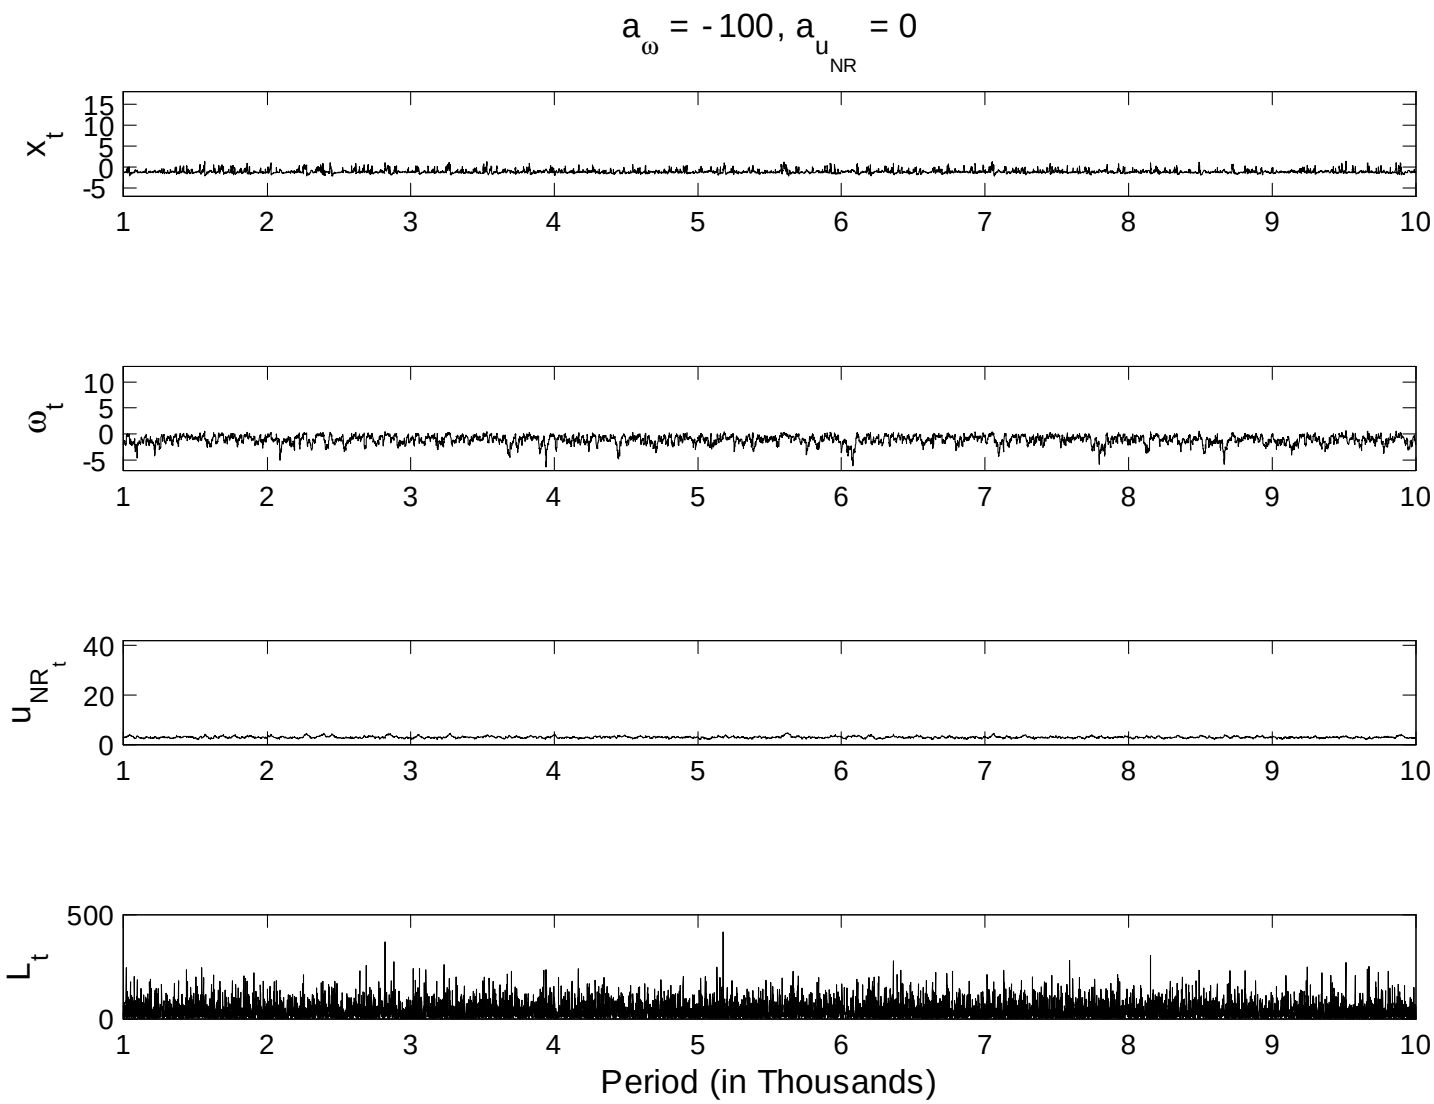

Figure 10. Targeted Inflation, Slope, Intercept and Economic Loss Estimates (Slope Asymmetry −100).

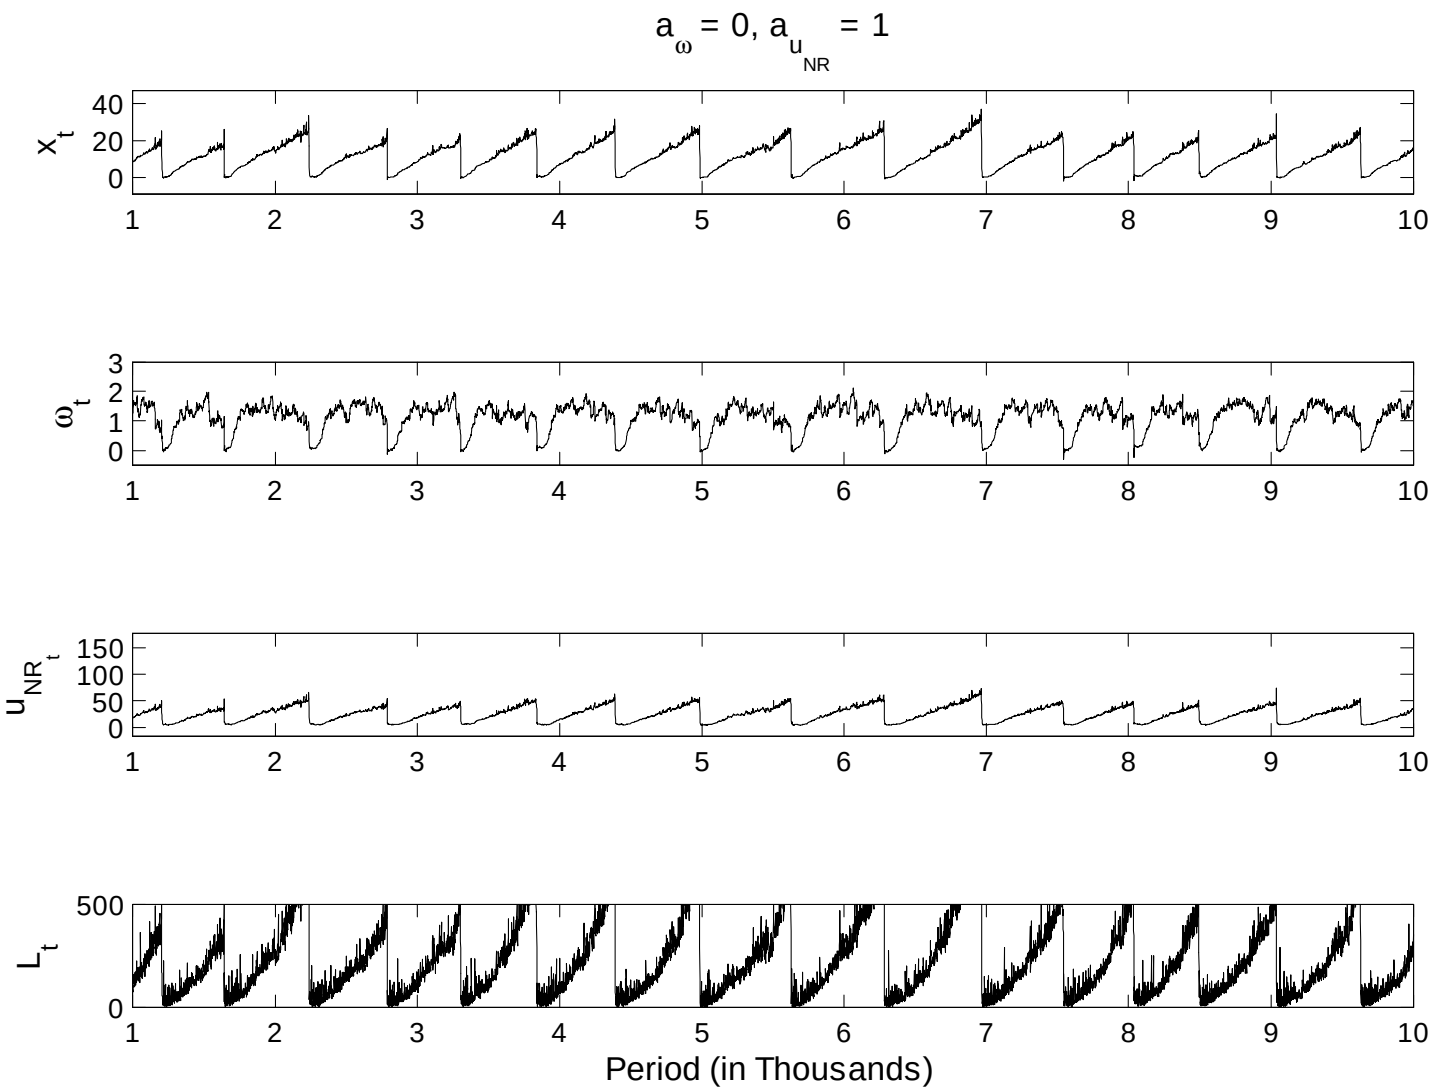

Figure 11. Targeted Inflation, Slope, Intercept and Economic Loss Estimates (Intercept Asymmetry +1).

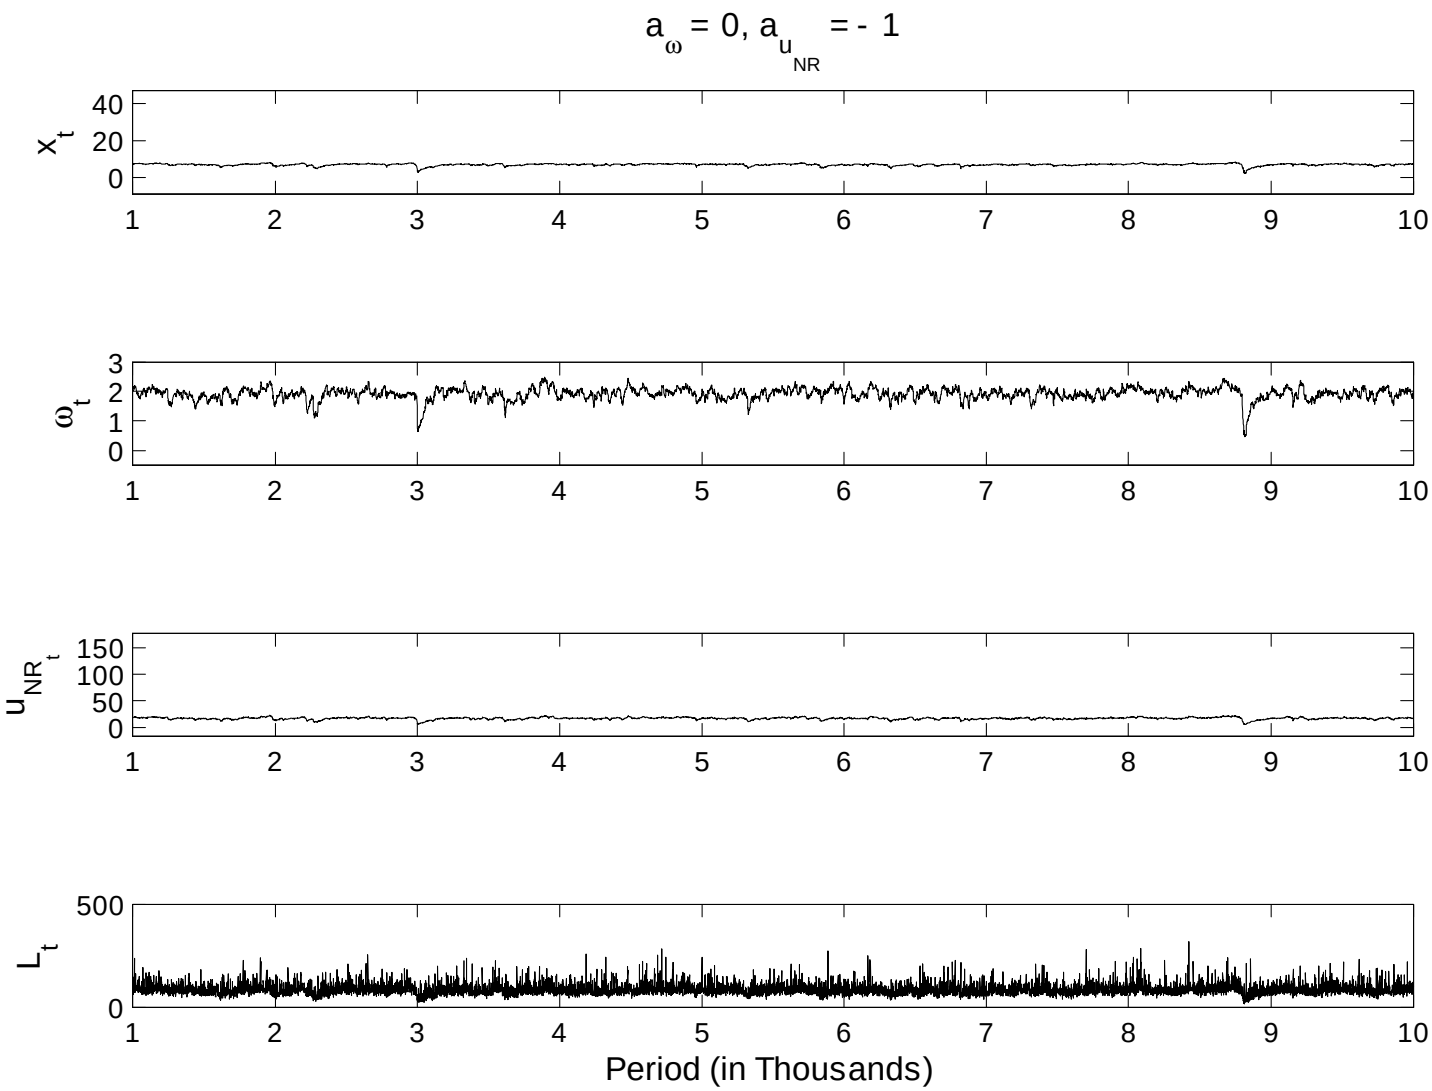

Figure 12. Targeted Inflation, Slope, Intercept and Economic Loss Estimates (Intercept Asymmetry -1).

In the second panel of Figure 11, we see that overestimation of the intercept leads indirectly to lower values of the slope estimate. As a consequence, the probability of escape increases. However, if we differentiate (24) by u_{NR} , we find

$$\frac{\partial x_t^{PLM}}{\partial u_{NR,t}} = \frac{\alpha\omega_t}{1+\alpha\omega_t^2} > 0, \tag{41}$$

so overestimation of the intercept also produces higher values for target inflation. The benefit of more frequent escapes will be overwhelmed by the cost of higher inflation during normal times for small $a_{U_{NR}}$, and the period loss function averages to 219 for $a_{U_{NR}} = 1$. In contrast, Figure 12 shows that escapes are eliminated when we underestimate the intercept, but target inflation averages to a lower value of 6.9 when $a_{U_{NR}} = -1$. In that case, the period loss function averages to 87. Thus experimentation around $a_{u_{NR}} = 0$ will cause the econometrician to slightly underestimate the intercept. This pattern persists if we increase $|a_{U_{NR}}|$ even future. Figure 13 shows the time series for $a_{U_{NR}} = 10$, and Figure 14 shows the time series for $a_{u_{NR}} = -10$.

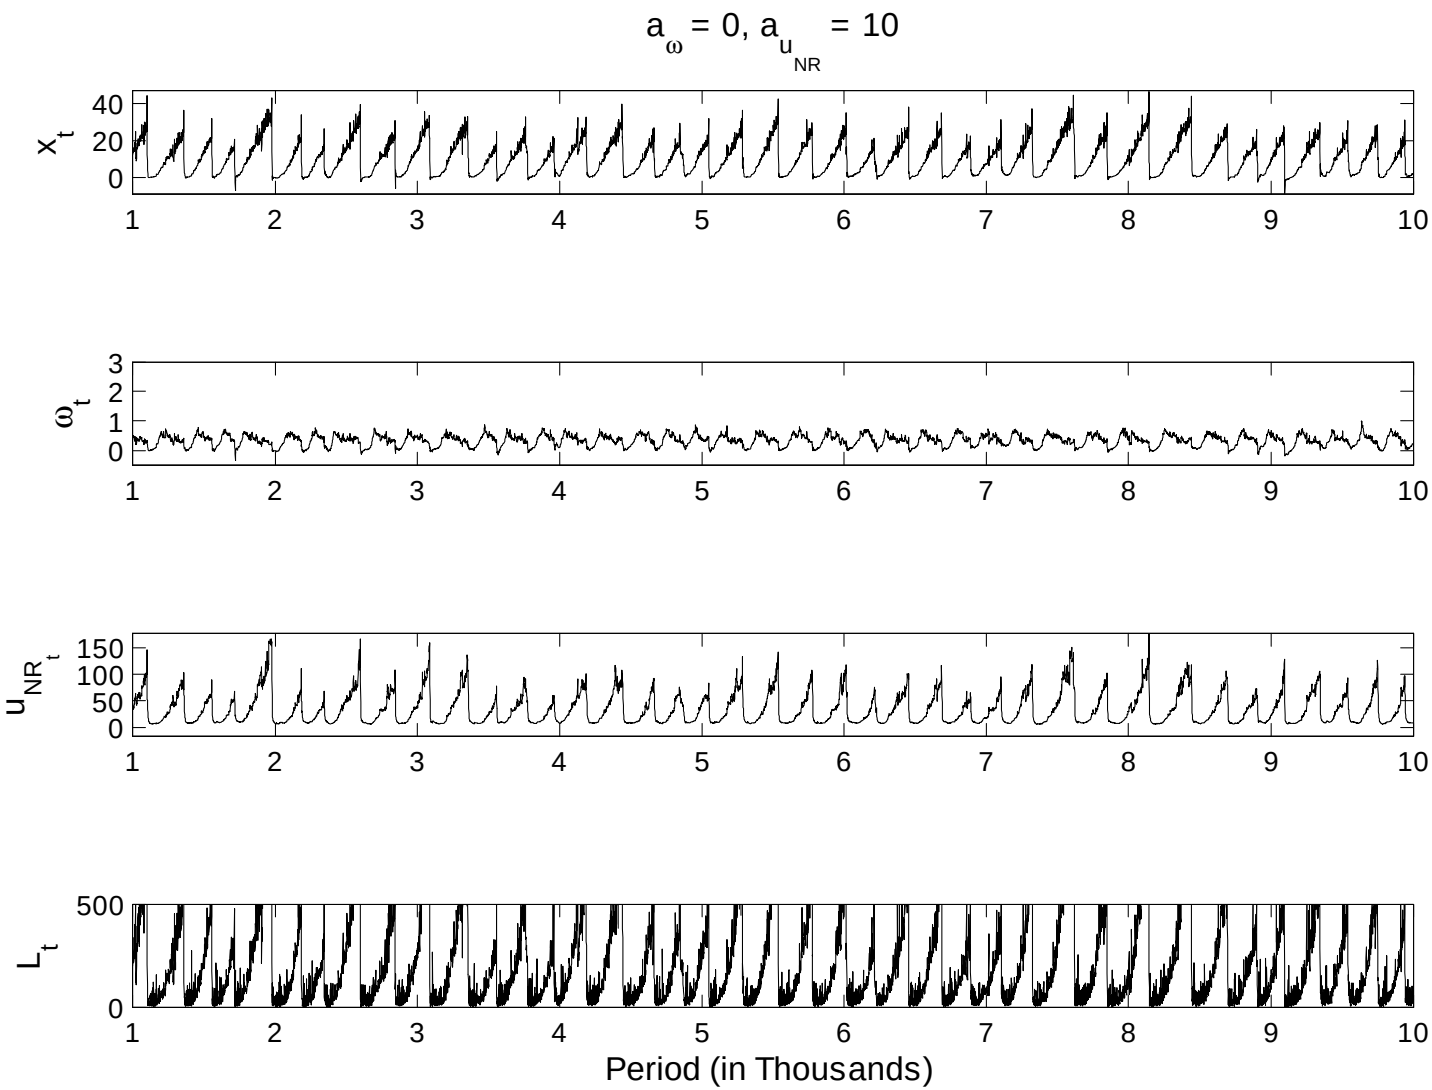

Figure 13. Targeted Inflation, Slope, Intercept and Economic Loss Estimates (Intercept Asymmetry +10).

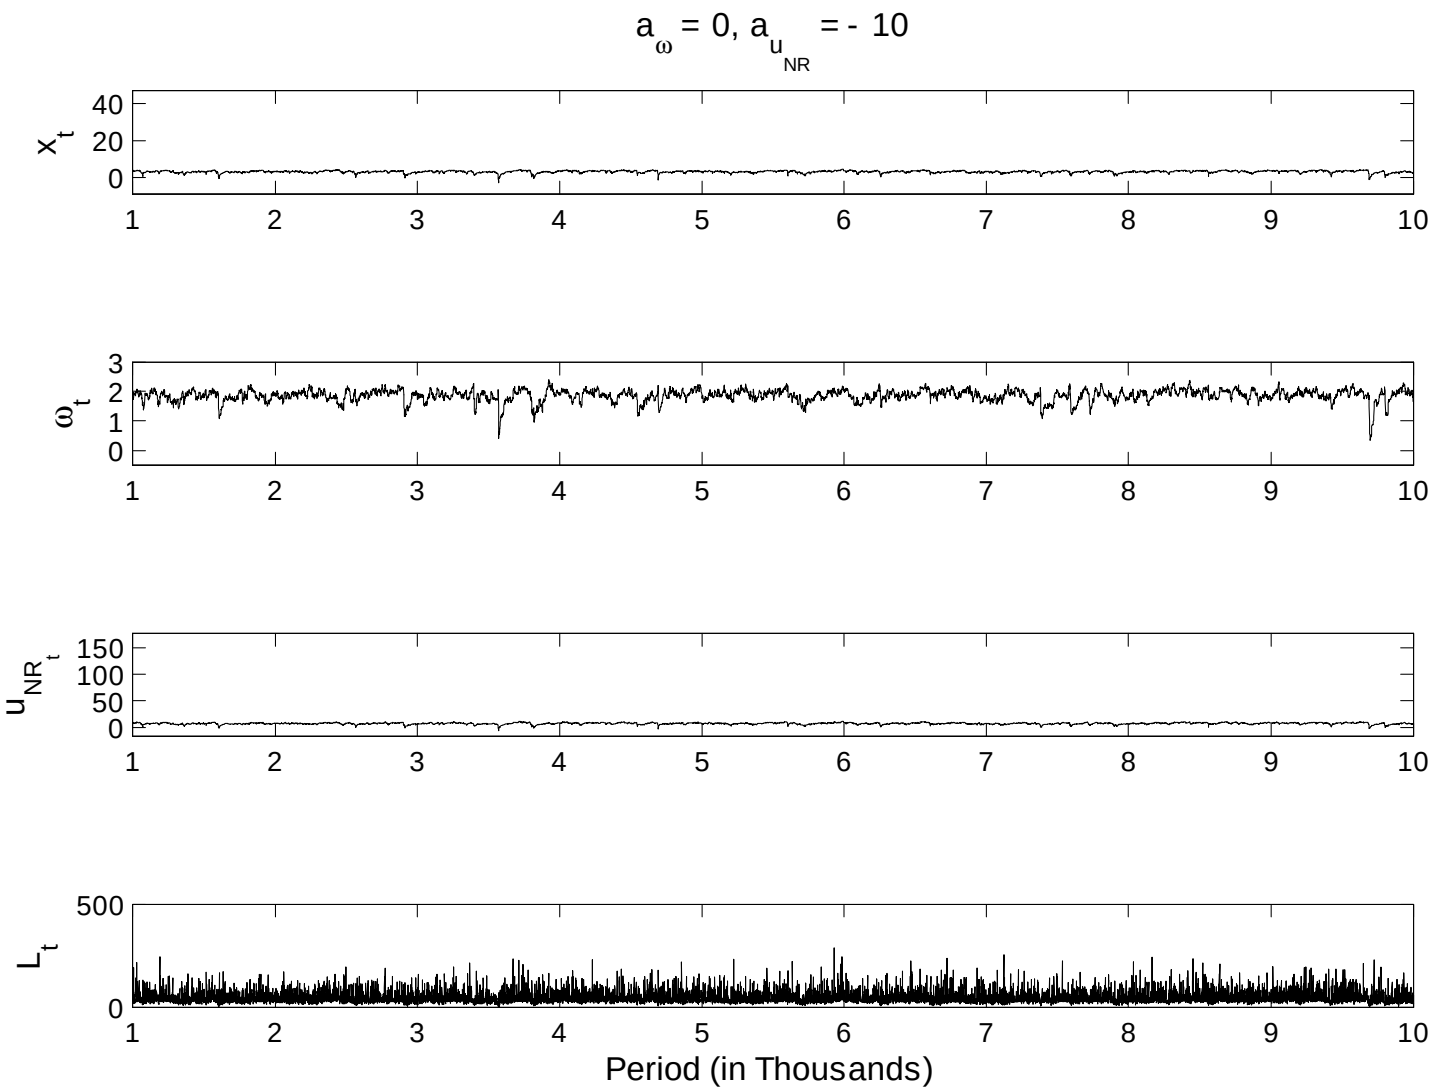

Figure 14. Targeted Inflation, Slope, Intercept and Economic Loss Estimates (Intercept Asymmetry -10).

Finally, Figure 15 shows the time series with the parameters $(a_{U_{NR}}, a_\omega) = (-10, -76)$ that produce the minimum loss function value of 39 for the region of the parameter space over which we searched. Although in the vicinity of $a_\omega = 0$, the loss function only decreases if we overestimate the slope with $a_\omega > 0$, for large $|a_\omega|$ underestimates reduce the loss function slightly more than overestimates.

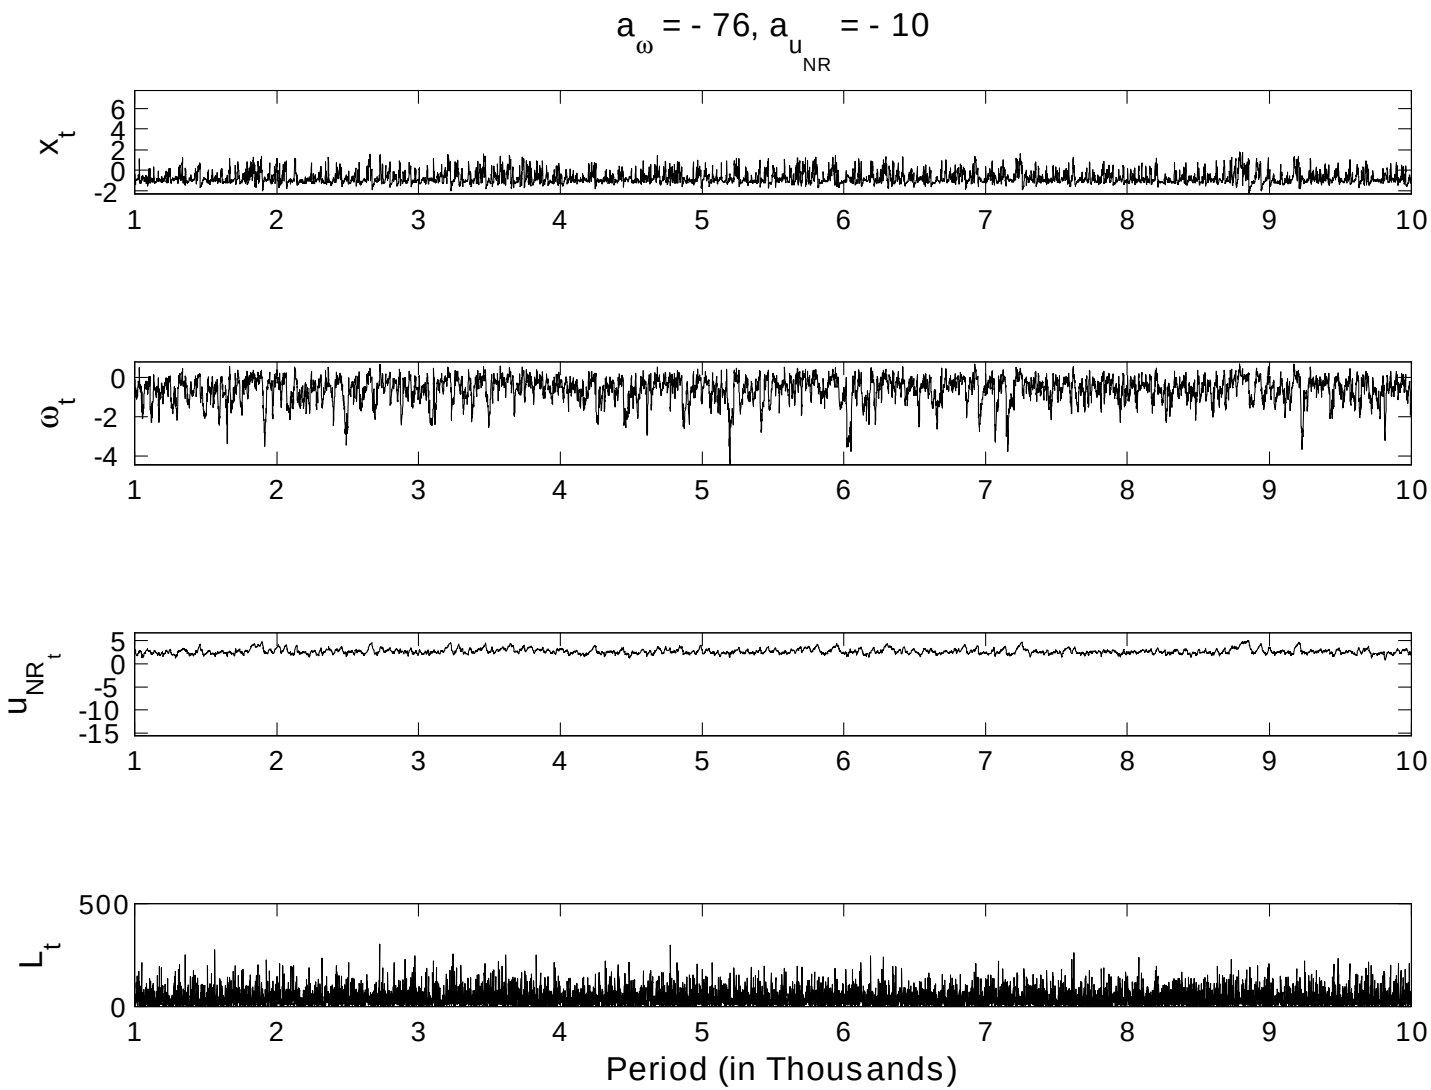

Figure 15. Targeted Inflation, Slope, Intercept and Economic Loss Estimates ($a_\omega = 76$, $a_{u_{NR_t}} = -10$).

5 Conclusion

The adaptive learning approach is a rich environment in which stability, perturbation, and other issues can be analyzed with respect to models with multiple possible outcomes. In this paper the approach has been to analyze learning dynamics in the event that the econometrician chooses the structural estimator to optimize the policymaker's loss function, which a responsible econometrician ought to do. Here we experiment with estimators that are appropriate if estimation errors in one direction are particularly costly. Our simulations demonstrate complex learning dynamics rich enough to warrant further analytical investigation. Unfortunately, we are not aware of any method by which stochastic approximation can be conducted for the constant gain case under our notion of variance-adjusted learning, so we limit our analysis to presenting the simulation results.

The simulations indicate that if there is overestimation of the slope relative to least squares, we see an increased frequency of the escapes analyzed by Cho et al. (2002). This is because estimates of the slope of the perceived law of motion (Phillips curve) fall often, along with wide variation in estimates of the intercept. This leads to overshooting-type dynamics in the targeted rate of inflation. However, this behavior is not monotonic. For high degrees of slope underestimation, the econometrician sees negative slope estimates and pushes down his estimates of the natural rate to near zero leading to low values of the targeted rate of inflation. These results suggest a possible tension between the pull of time-consistency and the degree to which a statistical decision maker is cautious about interpreting information (reflected in the estimator employed).

In summary, our analyses add to the ongoing investigation of various learning hypotheses advanced within the literature that employs adaptive filtering as the mechanism via which expectations are formed. This literature replaces the usual signal processing assumption (effectively, analog noise cancellation) with that of adaptive filtering. Doing so allows researchers to examine the effects of various types of regressions that can be performed as part of the filtering algorithm; we examine the effects of asymmetric least squares regressions

which nest the usual least squares algorithms, in the constant gain case. Our results point towards situations in which 'precaution' on the part of the learner could result in the lowering of the long run rate of inflation in an economy characterized by the time inconsistency of discretionary policy.

References

[1] Barro, R. J., Gordon, D. R., 1983. Rules, Discretion and Reputation in a Model of Monetary Policy. Journal of Monetary Economics 12, 101–121.

[2] Berger, J. O., 1985. Statistical Decision Theory and Bayesian Analysis (2nd. ed.). Springer.

[3] Carceles-Poveda, E., Giannitsarou, C., 2007. Adaptive Learning in Practice. Journal of Economic Dynamics and Control 31, 2659-2697.

[4] Cho, I-K., Sargent, T. J., Williams, N., 2002. Escaping Nash Inflation. Review of Economic Studies 69, 1-40.

[5] Cukierman, A., 2002. Are Contemporary Central Banks Transparent about Economic Models and Objectives and What Difference Does it Make?. Federal Reserve Bank of St. Louis Review 84, 15-35.

[6] Evans, G., Honkapohja, S., 1999. Learning Dynamics. Handbook of Macroeconomics (v. 1A). Elsevier.

[7] Evans, G., Honkapohja, S., 2001. Learning and Expectations in Macroeconomics. Princeton University Press.

[8] Kydland, F., Prescott, E. C., 1977. Rules Rather Than Discretion: The Inconsistency of Optimal Plans. Journal of Political Economy 85, 473-491.

[9] Leland, H. E., 1968. Saving and Uncertainty: The Precautionary Demand for Saving. The Quarterly Journal of Economics 82, 465-473.

[10] Ruge-Murcia, F. J., 2003. Inflation Targeting under Asymmetric Preferences. Journal of Money, Credit and Banking 35, 763-785.

[11] Sandmo, A., 1970. The Effect of Uncertainty on Saving Decisions. Review of Economic Studies 37, 353-360.

[12] Sargent, T. J., 1999. The Conquest of American Inflation. Princeton University Press.

[13] Varian, Hal R., (1975), "A Bayesian Approach to Real Estate Assessment," in *Studies in Bayesian Econometrics and Statistics in Honor of Leonard J. Savage*, eds. Stephen E. Fienberg and Arnold Zellner (Amsterdam: North-Holland).

[14] Zellner, A., 1976. Bayesian Estimation and Prediction Using Asymmetric Loss Functions. Journal of the American Statistical Association 81, 446-451.