CEM
Center for Empirical Macroeconomics

Working Paper No. 60

Long Horizon versus Short Horizon Planning in Dynamic Optimization Problems with Incomplete Information

by

Herbert Dawid

University of Bielefeld
Department of Economics
Center for Empirical Macroeconomics
P.O. Box 100 131

33501 Bielefeld, Germany

Long Horizon versus Short Horizon Planning in Dynamic Optimization Problems with Incomplete Information*

Herbert Dawid†

Department of Economics, University of Bielefeld

Abstract

This paper compares the implications of short and long horizon planning in dynamic optimization problems with the structure of a standard one-sector growth model if agents have incomplete knowledge about the production function. Agents know the output and rate of return at the current capital stock and use an estimation of the production function based on this knowledge to determine current consumption. For standard utility functions without wealth-effects both long and short planning horizons yield convergence to the steady state – however at a faster rate than optimal –, or fluctuations around the steady state, and in both cases, long horizon planning yields a policy which locally at the steady state is closer to the optimal one than short horizon planning. On the other hand, for preferences with wealth effects where the intertemporal optimal path exhibits fluctuations, long horizon planning destabilizes the path and short horizon planning can generate paths which are qualitatively closer to the optimal one and yield higher discounted utility.

JEL Classification: C61, D83, D90

Keywords: Dynamic Optimization, Heuristics, Local Stability

*The author would like to thank Richard Day for numerous stimulating discussions which led to this article and an anonymous referee for helpful comments.

†Department of Economics, University of Bielefeld, 33651 Bielefeld, Germany, email: hdawid@wiwi.uni-bielefeld.de.

1 Introduction

Problems of intertemporal optimization appear frequently in numerous fields of economic activity and there is a large amount of economic research dedicated to the analysis of such decision problems. Among the most prominent intertemporal optimization models studied is the standard one sector growth model introduced by Ramsey (1928) and extensively studied since. Typically, in the analysis of this type of model – which is a representative agent model – it is assumed that the decision maker has complete knowledge about all aspects of the problem, in particular about the production technology, and that he solves the resulting problem using dynamic programming techniques. It is well known that intertemporally optimal capital accumulation paths are monotonous in the standard one-sector model. To explain empirically observable fluctuations of capital stocks the standard model has to be modified by either introducing stochastic shocks (see e.g. Kydland and Prescott (1982)), the existence of wealth effects (see e.g. Majumdar and Mitra (1994)) or non-concavities in the utility function (see e.g. Dawid and Kopel (1997)).

Recently, Noussair and Matheny (2000) have carried out laboratory experiments with U.S. and Japanese subjects who had to determine their consumption paths in the framework of the standard one-sector growth model. They find systematic deviations from the intertemporally optimal paths in all their experiments, in particular too fast approaches towards the long run state and fluctuating paths. They also identify several heuristic decision rules of the subjects which seem to be based on myopic short run considerations rather than long horizon intertemporal optimization considerations. This is also in accordance with findings by Rust (1994) who studies a large set of different classes of dynamic optimization problems and compares the optimal solutions of the problems with different data sets. He finds that especially for problems with a continuous decision space the explanatory power of the optimal policies is very weak and argues that *[C]onsiderations of computational complexity may also force individuals to adopt simple heuristic decision rules that might be very different than the optimal decision rules calculated by dynamic programming.[...] people make decisions similar to the way a good AI program plays chess: it assigns more or less heuristic weights to provide summary evaluations of board positions many moves in the future, while using "brute force" calculations to systematically evaluate the consequences of intermediate moves.* Certainly computational complexity is an important motivation for heuristic decision making and the use of short planning horizons (see e.g. Simon (1982)), however we will argue in this paper that in situations, where the decision maker is not able to fully understand his environment and re-estimates a simplifying model of his environment every period, a case for short horizon planning can be made even without taking into account costs of computational complexity[1]. The intuitive argument here might be that if the decision maker's model of his environment is only locally correct, he might aggregate and magnify estimation errors when trying to predict future payoffs over a long time horizon and hence decrease rather than increase his actual payoff stream when planning far into the future. The purpose of this paper is to make this intuition more precise and characterize scenarios where it is valid and short run planning actually can be beneficial for the decision maker even if computational costs are ignored.

To do this we adopt the framework used in Day (2000). He considers a standard one-sector growth model and assumes that the decision makers do not have perfect knowledge about the production function but have to rely on linear estimations based on the current rate of return. Furthermore, the decision maker in Day's model is not willing or able to solve the infinite horizon dynamic optimization problem. Rather, he considers the effect of current consumption on current utility and the next period capital stock and uses a function to evaluate the future value of this capital stock (much in the spirit of Rust's quote above). However, this evaluation function is not the value function generated by infinite horizon optimization but a heuristic based on the assumption that the capital stock stays constant in the future. Day (2000) calls this kind of behavior adaptive economizing and shows that such behavior generates persistent oscillations in the standard one-sector Ramsey model. In this paper we take this adaptive economizing model as a representative example of a two-period horizon heuristic[2] and pose the question for which characteristics of the preferences and the production function the capital accumulation paths generated by this heuristic are 'closer' to the intertemporally optimal ones than the ones generated by infinite horizon planning every period. Whereas a comparison of global properties by analytical means is not feasible, locally at the steady states both the adaptive economizing policy and the policy generated by infinite horizon optimization

[1]It should be pointed out here that in the experiments by Noussair and Matheny (2000), the subjects did not have information about the actual functional form of the production function but could only observe discrete points on the production function.

[2]Clearly, there are many other sensible heuristics with this planning horizon and this choice is in some sense arbitrary. This is a general problem of the bounded rationality approach and our motivation here is that this is a model which has been well analyzed in the literature.

can be analytically characterized and compared. We do this for a general class of production function estimation functions which include the linear estimation functions considered in Day (2000). Furthermore, we present a numerical example which illustrates the theoretical findings and shows that adaptive economizing might indeed generate higher discounted utility values than infinite horizon planning.

The paper is organized as follows. In section 2 we introduce the model and the different types of strategies we consider. In section 3 we characterize the steady states of the policies considered and compare the stability properties of the different policies in the absence of wealth effects in section 4. We consider the model with wealth effects and discuss a numerical example in section 5 and close with a discussion of the results in section 6. All proofs are given in the Appendix.

2 General Framework

We consider an intertemporal optimization problem of the standard one-sector growth model type. Whereas we will interpret the problem in the growth theory framework it should be pointed out that models of this type appear in several different fields of economics, like resource economics or international economics (see e.g. Clark (1971) or the discussion in Noussair and Matheny (2000)). Every period output is produced using the current capital stock as the single input factor. Output can either be consumed or reinvested. The objective is to choose a consumption path which maximizes the discounted infinite horizon utility stream. Neglecting population growth we express all variables in per capita terms. The capital stock at time t is denoted by k_t and c_t is consumption. We have a continuous production function $f : \mathbb{R}_0^+ \mapsto \mathbb{R}_0^+$ which is twice continuously differentiable on $(0, \infty)$ with the standard properties

$$\begin{array}{c} f'(k) > 0, \quad f''(k) < 0 \quad \forall k > 0 \\ f(0) = 0, \quad \lim_{k \to 0} f'(k) = \infty, \quad \lim_{k \to \infty} f'(k) = 0. \end{array} \tag{1}$$

Capital accumulates according to

$$k_{t+1} = (1 - \delta)k_t + f(k_t) - c_t \qquad k \geq 0, \tag{2}$$

where δ is the depreciation factor of capital. It is assumed that investment is irreversible and only current output can be consumed, i.e.: $0 \leq c_t \leq f(k_t)$. Standard arguments show that there is a unique capital stock $k^m > 0$ with

$k^m = (1-\delta)k^m + f(k^m)$ which would be the limit point of the trajectory of capital stocks without consumption. In what follows we only consider the state space $[0, k^m]$. Per period utility may in general depend both on current consumption c_t and the current capital stock k_t. The utility function $u(k_t, c_t)$, which maps $\mathbb{R}_0^+ \times \mathbb{R}_0^+$ into $\mathbb{R}_0^+$, is assumed to be twice differentiable and everywhere satisfies the following standard conditions:

$$\begin{aligned} &u_1 \geq 0, \quad u_{11} \leq 0, \quad u_2 > 0, \quad u_{22} < 0 \\ &u_{11}u_{22} - u_{12}^2 \geq 0, \ \lim_{c \to 0} u_2(k, c) = \infty. \end{aligned} \tag{3}$$

Utility of consumption increases both with current consumption and the current capital stock, where the utility function is jointly concave in k and c and strictly concave in consumption.

The purpose of this study is to analyze the implications of myopic decision making if only limited information about the environment is available. Hence, the model we use has two central aspects. First, the decision maker does not know the exact form of the production function but estimates it based on observable information. Second, given this limited information, the decision maker might either use a heuristic based on a simplified two-period optimization problem or solve the infinite horizon problem.

Let us discuss these two aspects in more detail. Assume the decision maker does not only not know the exact form of the production function f, but also misses sufficient structural insight into the production process to be able to determine the correct functional form of the production function. On the other hand, he can observe the output for the current capital stock and the current rate of return. Using this (and maybe some additional information about the production process he receives) he constructs an estimation of the production function based on some (in general incorrect) internal model. We denote by $\hat{f}(k_1; k)$ the estimated output for a capital stock of k_1 if the current capital stock is k. For analytical convenience we assume that $\hat{f}$ is twice differentiable with respect to k_1, differentiable with respect to k and concave in k_1. Since current output and current rate of return can be observed, this estimation function fulfills

$$\hat{f}(k; k) = f(k), \quad \hat{f}'(k; k) = f'(k) \ \forall k \in [0, k^m]. \tag{4}$$

Furthermore, we assume that the estimated output always lies between the actual output and the linear estimation based on the current rate of return:

$$f(k_1) \leq \hat{f}(k_1; k) \leq f(k) + f'(k)(k_1 - k). \tag{5}$$

This assumption is for example satisfied by all piece-wise linear approximations based on information on rates of returns at different capital stocks[3].

The set of all twice differentiable estimation functions satisfying (4) and (5) is denoted by $\mathcal{F} \subseteq C^2[0, k^m]$. Given these assumptions, it is easy to derive the following properties of the estimation function by total differentiation of the equalities in (4) ($\hat{f}_k$ denotes the derivative of $\hat{f}$ with respect to the current capital stock):

$$\hat{f}_k(k;k) = 0, \quad \hat{f}'_k(k;k) = f''(k) - \hat{f}''(k;k) \quad \forall k \in [0, k^m] \tag{6}$$

Assumption (5) on the other hand implies:

$$f''(k) \leq \hat{f}''(k;k) \leq 0 \quad \forall k \in [0, k^m].$$

The two extreme cases of this class of estimation functions are on one side the linear approximation

$$\hat{f}(k_1;k) = f(k) + f'(k)(k_1 - k),$$

which uses only information about current output and rate of return, and

$$\hat{f}(k_1;k) = f(k_1), \ \forall k, k_1 \in [0, k^m]$$

where the decision maker has complete knowledge about the production function.

The use of (local) linear approximations of unknown non-linear relationships is a widely used practice in many real world planning problems. For example, the standard approach in production planning textbooks (e.g. Sipper and Bulfin (1997)) is to assume constant productivity of capital in long-run planning models although the existence of actual non-linearities is acknowledged. The productivity parameter estimations are regularly revised in the planning process as new information comes in. Our framework is supposed to capture such an approach in a very stylized way.

In our model the decision maker 's estimation of the production function depends only on current output and rate of return, which implies that no information aggregation about the production function over time occurs. This clearly is a restriction we inherit by adopting the framework of Day (2000). If we allow $\hat{f}$ to depend not only on current output and rate of return but also on a fixed or increasing number of past observations, this

[3]To fulfill the differentiability conditions such an estimation function should be approximated by an $\hat{f}$ where the kinks are 'smoothed'.

adds a large amount of complexity to the model and the theoretical analysis carried out below would hardly be feasible anymore. We will argue below that with respect to the main question addressed, namely the comparison of stability of long and short horizon planning, the additional qualitative insights that could be gained by this added complexity are limited. This is confimred in the numerical example discussed at the end of section 5 where we consider the effects of a longer memory as well.

We consider two types of intertemporal decision making of agents with such limited information about the production function: adaptive economizing and repeated infinite horizon optimization. As pointed out in the introduction, adaptive economizing agents use a rather simple heuristic procedure with a two-period planning horizon. They attach a capital stock in the subsequent period a value which is determined by the utility of a future consumption stream which keeps the capital stock constant. Put more formally, if the current consumption decision of an agent with capital stock k yields a capital stock k_1 in the subsequent period then the agent values this capital stock with

$$\Psi u(k_1, c_1),$$

where c_1 is determined by the condition

$$(1-\delta)k_1 + \hat{f}(k_1; k) - c_1 = k_1.$$

The parameter Ψ is called the 'future weight'. This heuristic evaluation of future utility from capital stock k_1 coincides with the actual discounted utility stream if the capital stock stays at the level k_1 for all future and the per period discount factor is $\alpha = \frac{\Psi}{1+\Psi}$. This leaves the agent with a rather simple optimization problem:

$$\max_{0 \le c \le f(k)} u(k, c) + \Psi u((1-\delta)k + f(k) - c, \hat{f}((1-\delta)k + f(k) - c; k) - \delta((1-\delta)k + f(k) - c)) \tag{7}$$

Given our assumptions it is straight forward to check that there exists a unique solution to this problem for every k. The consumption which solves (7) is always positive, it might however be at the upper boundary $f(k)$. We denote by $g(k; \hat{f})$ the solution of (7). Note that this solution of course depends on the estimation function $\hat{f}$ the agent uses. It is easy to see that $g(k; \hat{f})$ is continuous on $[0, k^m]$. The capital accumulation path generated by such a behavior is given by

$$k_{t+1} = \theta(k_t; \hat{f}) := (1-\delta)k_t + f(k_t) - g(k_t; \hat{f}). \tag{8}$$

We call $\theta(\cdot; \hat{f})$ the adaptive economizing policy of the agent.

In what follows we will compare the capital accumulation paths generated by adaptive economizing with the intertemporally optimal paths and with the paths generated if the decision maker chooses current consumption every period after solving the infinite horizon optimization problem using the current estimation function. Given that the agent uses the estimation function $\hat{f}$ and the current capital stock is k_0, the dynamic optimization he faces if he has an infinite planning horizon reads:

$$\max_{\{c_t\}} \sum_{t=0}^{\infty} \alpha^t u(k_t, c_t)$$
$$\text{s.t.} \quad k_{t+1} = (1-\delta)k_t + \hat{f}(k_t; k_0) - c_t$$
$$0 \leq c_t \leq \hat{f}(k_t; k_0) \qquad \forall t \geq 0$$
$$k_0 = k0,\ \alpha \in [0,1),\ \delta \in [0,1)$$

Standard arguments (see Stokey and Lucas (1989)) establish that this problem has a unique optimal path $\{\tilde{c}_t\}_{t=0}^{\infty}$ which is generated by a continuous policy function $\tilde{\tau}(k; \hat{f}, k_0),\ k \in [0, k^m]$. Even if agents successfully solve this dynamic optimization problem, the actual consumption path will in general differ from $\{\tilde{c}_t\}_{t=0}^{\infty}$. The estimation of the production function will change over time and therefore in this setup the intertemporally optimal path determined at time zero does not solve all dynamic optimization problems the decision maker faces in subsequent periods. Accordingly, the decision maker has to solve the dynamic optimization problem with the current estimation of the production function every period. The actual capital accumulation path is given by:

$$k_{t+1} = \hat{\tau}(k_t; \hat{f}) := \tilde{\tau}(k_t; \hat{f}, k_t). \tag{9}$$

We call this procedure rolling infinite horizon planning and $\hat{\tau}$ the rolling planning policy[4]. The solution of the standard dynamic optimization problem, where the exact form of f is known, is included in this class of models as a special case. We denote the optimal policy function of the intertemporal optimization problem with complete information by τ and obviously $\tau(k) = \hat{\tau}(k, f)\ \forall k$. In what follows we denote by $\tilde{\tau}_k$ the derivative of $\tilde{\tau}(\cdot; \hat{f}, k)$ with respect to k. From the definition of $\hat{\tau}$ we get:

$$\hat{\tau}' = \tilde{\tau}' + \tilde{\tau}_k$$

[4]The concept of finite horizon rolling planning has been introduced in different frameworks with complete information. Asymptotic optimality properties of such plans as the planning horizon goes to infinity have been established for example in Goldmann (1968), Kaganovich (1985) or Bala et al. (1991).

Thus, the effect of a change of the current capital stock on the subsequent capital stock can be divided into the effect of the change of the capital stock on the policy function given the current estimation of f and the effect on the estimation function $\hat{f}$. This distinction will turn out to be useful in order to understand the stability properties of $\hat{\tau}$. It should also be noted that the policy function $\tilde{\tau}(.;\hat{f},k)$ for a given k might yield infeasible paths. On the other hand, the adaptive economizing and rolling planning policies θ and $\hat{\tau}$ always generate feasible paths (of course current plans for future periods might be infeasible but they will be revised later on).

In our treatment of stability issues in the presence of wealth effects it will be useful to consider the reduced form rather than the primitive form of the optimization problem (with complete information). Denote by $v(k_t,k_{t+1}) := u(k_t,(1-\delta)k_t+f(k_t)-k_{t+1})$ the reduced from utility function and by $\Omega = \{(k_t,k_{t+1})|k_t \in [0,k^m],(1-\delta)k_t \leq k_{t+1} \leq (1-\delta)k_t + f(k_t)\}$ the set of feasible sequences of capital stocks. Standard arguments establish that the reduced form utility function is (jointly) concave, increasing in the first and decreasing in the second argument and that Ω is convex and compact. The optimization problem in the reduced form model then reads

$$\begin{aligned} &\max_{\{k_t\}} \textstyle\sum_{t=0}^{\infty} \alpha^t v(k_t,k_{t+1}) \\ &\text{s.t.} \quad (k_t,k_{t+1}) \in \Omega \quad \forall t \geq 0. \\ &k_0 = k0 \end{aligned}$$

3 Steady States

It is well known that, in cases where utility depends only on current consumption, the optimal policy of the one sector optimal growth problem has two fixed points, namely one at zero and one at the capital stock where $\alpha(1-\delta+f'(k)) = 1$. However, if current utility of consumption is influenced by the current capital stock, there might be multiple positive fixed points. In the following proposition we show that the maps τ, $\hat{\tau}$ and θ always have the same fixed points. In all the comparisons of optimal and adaptive paths we assume that the future weight used in the adaptive economizing policy is compatible with the discount factor used for the calculation of intertemporally optimal paths, i.e. $\Psi = \frac{\alpha}{1-\alpha}$. Given this, we have

Proposition 1 *The set of fixed points of $\tau(\cdot)$, $\hat{\tau}(\cdot,\hat{f})$ and $\theta(\cdot,\hat{f})$ in $[0,k^m]$ coincides for all $\hat{f} \in \mathcal{F}$.*

Note that this Proposition does not imply that all fixed points of $\tilde{\tau}(\cdot; \hat{f}; k)$ are fixed points of τ. It is easy to see, that $\tilde{\tau}$ may indeed have additional fixed points. This means that, whereas individuals might expect a certain capital stock to be a steady state given their current estimation of the production function $\hat{f}$, they always revise their opinion if they actually reach this capital stock unless this capital stock is a steady state of the optimal policy. Having established now that the set of fixed points of all three maps under consideration are identical we will discuss the local stability properties of these fixed points. We start this discussion with the simpler case without wealth effects.

4 Local Stability without Wealth Effects

Let us first assume that current utility depends only on current consumption, i.e. $u(k, c) = u(c)$. In most studies of optimal growth or renewable resource exploitation problems this assumption is made and the optimal policy in such a framework has been studied extensively (see e.g. Stokey and Lucas (1989)). It is well known that the optimal policy is strictly increasing on $[0, k^m)$ and thus convergence towards the unique positive steady state of the optimal policy is always monotonous. On the other hand, it was demonstrated in Day (2000) that adaptive economizing with a linear estimation of the production function $\hat{f}$ may lead to persistent fluctuations around the steady state and chaotic behavior. Here we will show that such fluctuations are mainly due to the linear approximation with short memory of the production function and may also occur if agents do infinite horizon planning and use the policy $\hat{\tau}(k; \hat{f})$ every period.

To make the point that the fluctuations observed in Day (2000) are caused by limited information about f rather than by the short planning horizon, we first show that if the agents know the exact form of the production function no fluctuations can occur even under adaptive economizing. Like in the optimal solution the capital stock converges monotonously towards the steady state:

Proposition 2 *If current utility does not depend on the current capital stock and the agents have complete information about the production function, then the adaptive economizing policy $\theta(k; f)$ is strictly increasing in k on $[0, k^m]$.*

This proposition shows that if agents estimate the production function correctly there is no qualitative difference between the intertemporally optimal paths and those generated by adaptive economizing. Both paths monotonously approach the positive steady state. Of course the speed of convergence towards the steady state in general differs.

If the estimation function $\hat{f}$ does not coincide with f this monotonicity property of θ is lost. Fluctuations of the adaptive economizing path around the steady state are possible in such a case. On the other hand, under such circumstances it is also questionable whether the rolling planning path created by period–per–period dynamic programming solutions is stable. The following proposition addresses these questions. We characterize the slope of the adaptive economizing function at the steady state and compare it to the slope of the optimal and rolling infinite horizon planning policy functions.

Proposition 3 *Let k^* denote the unique positive steady state of the optimal policy and $c^* = f(k^*) - \delta k^*$ the corresponding consumption. Then*

$$\theta'(k^*; \hat{f}) \leq \hat{\tau}'(k^*; \hat{f}) \leq \tau'(k^*) \quad \forall \hat{f} \in F, \tag{10}$$

where the first inequality is strict if $\hat{f}''(k^*; k^*) < 0$ and the second inequality is strict if $\hat{f}''(k^*; k^*) > f''(k^*)$. The slope of the adaptive economizing policy at k^* is negative if

$$\hat{f}''(k^*; k^*) > f''(k^*) - \frac{(1-\alpha)u''(c^*)}{\alpha^2 u'(c^*)}$$

and the steady state is unstable with respect to the adaptive economizing policy if

$$\hat{f}''(k^*; k^*) > \frac{1}{2} f''(k^*) - \frac{(1-\alpha)u''(c^*)}{\alpha^2 u'(c^*)}.$$

This proposition shows that there are two effects at work if we compare adaptive economizing behavior with the intertemporally optimal solution and that both effects go in the same direction leading to excess sensitivity of consumption with respect to capital stock. First, the use of an estimation of the production function which is 'less curved' than the actual function leads to a policy function which at the steady state is flatter than the optimal policy regardless whether the agent solves an infinite horizon dynamic optimization problem every period or uses the adaptive economizing heuristics. Both the adaptive economizing and the rolling planning policy become flatter the less curved $\hat{f}$ is at the steady state and the slope might become negative if $\hat{f}''$ is sufficiently close to zero. Second, excess sensitivity of consumption is also induced by the use of an adaptive economizing policy rather than the solution of the infinite horizon problem regardless of the beliefs about the production function. If the adaptive economizing policy is compared to the optimal policy these two effects add up implying that adaptive economizing paths approach the steady state faster than optimal and might even overshoot and fluctuate.

Rolling infinite horizon planning stabilizes the capital accumulation path, compared to rolling planning with a two period horizon, and leads to paths which are qualitatively closer to the intertemporally optimal ones than the adaptive economizing paths. Strictly speaking this holds only true if the estimated production function is not linear at k^*. For a linear $\hat{f}$ the adaptive economizing and rolling infinite horizon planning policy have the same slope at k^*.

Corollary 1 *For $\hat{f}'' = 0$ we have*

$$\theta'(k^*; \hat{f}) = \hat{\tau}'(k^*; \hat{f}) = 1 - \frac{\alpha^2 u' f''}{(1-\alpha)u''}.$$

Surprisingly, if a linear estimation of the production function is used, the qualitative properties of the adaptive economizing policy and the rolling planning policy – which of course needs significantly more computational effort – completely coincide at least locally at the steady state. Together, the previous proposition and this corollary imply that with linear estimation functions the steady state might be unstable not only with respect to adaptive economizing paths but also for infinite horizon planning. Simple examples show that in the our formulation we might indeed have limit cycles of paths generated by θ and $\hat{\tau}$. It follows, however, from standard results in growth theory and from the arguments in the proof of proposition 2 that for any fixed estimation of the production function the policies $\hat{\tau}$ and θ have to be monotonously increasing in the absence of wealth effects. Accordingly, limit cycles of order k are only possible if the memory size is below k, and the fact that we might have 2-cycles here is an implication of our assumption of memory of length one[5]. The two main qualitative insights about the de-stabilizing effects of the use of linearized estimation functions and the de-stabilizing effect of adaptive economizing should however be robust with respect to an increase in the memory size (see also the discussion in section 5).

The proposition also shows that fluctuations of the adaptive economizing policy are facilitated by a large discount factor. In light of the well known turnpike results (e.g. McKenzie (1986)) and research on complex optimal paths (e.g. Mitra (1996), Nishimura and Yano (1996)) this might be surprising. Hoewever, it has to be realized that in our framework the incomplete

[5]We are grateful to a referee for pointing this out.

information about f implies that the value function proxies used are incorrect, both with adaptive economizing and infinite horizon planning, and change from period to period. The larger the discount factor the larger is the weight assigned to these proxies in the decision process and the more likely it is that changes in the value function estimations translate into fluctations of the derived decision.

Summarizing, we have shown that in the absence of wealth effects rolling infinite horizon planning yields paths closer to the optimal ones than adaptive economizing with a two period planning horizon, as long as the estimation function $\hat{f}$ is not linear. For linear estimation functions the two policy functions coincide locally around the steady state. These results and also numerous numerical examples suggest that rolling infinite horizon planning generates a larger discounted payoff stream than the adaptive economizing policy if there are no wealth effects. Unfortunately, no general analytical proof could be found for this conjecture.

5 Models with Wealth Effects

We now return to the general case and assume that current utility depends on current consumption and the current capital stock. Even under standard concavity assumptions the optimal policy in such problems may have several positive fixed points and the optimal paths may exhibit persistent fluctuations or chaotic behavior (see Majumdar and Mitra (1994)). It has also been shown that the optimal policy at a positive steady state is downward sloping if and only if the cross derivative of the reduced form utility function $v_{12}(k^*, k^*) = -u_{12}(k^*, c^*) - (1 - \delta + f'(k^*))u_{22}(k^*, c^*)$ is negative (Benhabib and Nishimura (1985)). Here we will concentrate on this case since the case where the optimal policy is increasing at the steady state is qualitatively very similar to the case without wealth effects we have discussed above[6]. In the next proposition we show that for $v_{12} < 0$ the slope of the adaptive economizing policy at the steady state is negative but larger than that of the optimal policy if $\hat{f}''$ is sufficiently close to f''.

Proposition 4 *Assume that $v_{12}(k^*, k^*) < 0$ at a fixed point k^* of the optimal policy. Then $\theta'(k^*; \hat{f})$ is negative for all $\hat{f} \in \mathcal{F}$ and $\theta'(k^*; \hat{f})$ is smaller the larger $\hat{f}''(k^*)$ is. The slope of the adaptive economizing policy is larger*

[6] Note that in the absence of wealth effects the concavity of u implies $v_{12} > 0$.

than that of the optimal policy if

$$\hat{f}''(k^*;k^*) < f''(k^*) - \frac{v_{12}(k^*,k^*)}{u_2(k^*,c^*)}. \tag{11}$$

An implication of this proposition is that, if the agents correctly estimate the second derivative of the production function and $\tau' < 0$, adaptive economizing always leads to a policy function which is less steep than the optimal one. Hence adaptive economizing has a *stabilizing* effect and – as in the case without wealth effects – implies that the steady state is reached faster than optimal. In particular, this means that, in cases where the slope of the optimal policy is only slightly smaller than -1, the steady state is unstable with respect to the optimal policy whereas adaptive economizing with complete knowledge about f leads to dampening fluctuations and (local) convergence towards the steady state. If $v_{12} < u_2 f''$ such a scenario even occurs for the linear estimation functions used in Day (2000). Thus, the assertion that short horizon planning is more likely to lead to fluctuations and less stable behavior than would be optimal does not necessarily hold in the general case with wealth effects.

Concerning the rolling planning policy, we can derive the following result about the slope of the function $\tilde{\tau}(\cdot;\hat{f},k^*)$ at a steady state k^*.

Proposition 5 *Let k^* be a steady state of τ with $v_{12}(k^*,k^*) < 0$. Then we have*

$$\tilde{\tau}'(k^*;\hat{f},k^*) \leq \tau'(k^*) < 0 \ \forall \hat{f} \in \mathcal{F},$$

where the inequality is strict if $\hat{f}''(k^*) \neq f''(k^*)$.

If an agent would stick to his consumption policy determined using the initial estimation of the production function the generated path of capital stocks would also fluctuate and be less stable than the optimal one in a sense that instability of k^* with respect to τ always implies instability with respect to $\tilde{\tau}$ but not vice versa[7]. Under condition (11) this implies also that the slope of $\tilde{\tau}$ is smaller than that of θ and there are situations where $\tilde{\tau}' < -1 < \tau' < \theta'$, which means that given the estimation $\hat{f}$ adaptive economizing leads to a path which is qualitatively similar to the intertemporally optimal path with complete knowledge about the production function – namely convergence towards the steady state in dampening fluctuations – whereas rolling infinite

[7]This statement makes only sense locally around an interior steady state since in general there is no guarantee for feasibility of $\tilde{\tau}$ and external bounds would have to be added to convert it into a feasible strategy. Due to continuity such bounds are not binding locally at the interior steady state.

horizon planning yields expanding fluctuations and no convergence towards the steady state.

The situation is less clear if we assume that agents every period solve the dynamic optimization problem anew using their current estimation of the production function. As pointed out above such behavior leads to a policy function $\hat{\tau}(k;\hat{f})$ and we have $\hat{\tau}' = \tilde{\tau}' + \tilde{\tau}_k$. Concerning $\tilde{\tau}_k$ at a steady state k^* we get from differentiating (12) with respect to k_0:

$$\tilde{\tau}_k = \frac{\alpha u_2(f'' - \hat{f}'')}{-\alpha u_{11} + \alpha(1-\delta+f'-\tilde{\tau}')v_{12} + (\alpha(1-\delta+f')-1)u_{22} - \alpha u_2 \hat{f}''}.$$

Whereas the numerator is always negative it is a priori not clear whether the denominator is positive or negative. Contrary to the case without wealth effects we are not able to give general analytical conditions which determine whether the slope of $\hat{\tau}$ at k^* is larger or smaller than that of τ. However, below we provide an example where adaptive economizing gives a policy function which is qualitatively closer to the optimal one than that created by rolling infinite horizon planning and also yields a larger discounted utility.

Before we discuss the example it is helpful to get a better intuitive understanding of these results. Again, we like to distinguish between the effects of adaptive economizing per se and that of incomplete information. Let us first focus on the effect of adaptive economizing under complete information. Assume that $\hat{f} = f$ and consider a steady state k^* with $v_{12}(k^*) < 0$. Furthermore, consider a current capital stock with $k_t > k^*$ but sufficiently close to k^* such that $(1-\delta)k_t < \tau(k_t) < k^*$. The optimal period $t+1$ capital stock is defined by $v_2(k_t, k_{t+1}) + \alpha W'(k_{t+1}) = 0$, where W is the value function. For adaptive economizing the condition reads $v_2(k_t, k_{t+1}) + \alpha R'(k_{t+1}) = 0$, where $R(k_{t+1}) = \frac{1}{1-\alpha} v(k_{t+1}, k_{t+1})$ is the adaptive economizing proxy of the value function. Obviously, we have $W(k) \geq R(k)$ for all k with equality at k^*, which implies that the proxy of the value function touches the actual value function from below at k^*. Accordingly, at least close to k^* the proxy is more concave than the actual value function. With other words, by looking at a two-period rather than an infinite horizon, adaptive economizing induces an overestimation of the marginal effect on future utility of a decrease of k_{t+1} below k^*. On the other hand, the current marginal effect of an increase in period t consumption is evaluated correctly. Hence, an adaptive economizer under-estimates the marginal effect of an increase of current consumption on total discounted utility at capital stocks below k^* which implies that there is too little consumption and a capital stock in $[\tau(k_t), k^*]$ is chosen. For capital stocks above k^* this marginal effect is overestimated, too much is consumed and therefore $\theta(k_t) \in [k^*, \tau(k_t)]$ if $\tau(k_t) > k^*$. This explains nicely why under perfect information adaptive economizing leads to faster approach towards the steady state than optimal. It also explains why adaptive economizing under full information leads to overshooting if and only if the optimal paths exhibit local overshooting at the steady state. In particular, the arguments given above show that (11) guarantees for $v_{12}(k^*, k^*) > 0$ that the slope of the adaptive policy is larger than that of τ (which of course is positive in such a case).

Now let us consider the effect of incomplete information. For the estimated production function $\hat{f}(\cdot; k_t)$ we have $f'(k; k_t) \leq f'(k)$ for $k < k_t$ which should translate into $\hat{W}'(k; k_t) < W'(k)$ for $k < k_t$ where $\hat{W}'$ denotes the value function of the dynamic optimization problem with production function $\hat{f}$. In other words, the 'less curved' estimation of the production function induces a 'less curved' value function and an underestimation of the marginal effect of a decrease in next period's capital stock for $k < k_t$. By an analogous argument to that given in the previous paragraph this implies that as long as $\tau(k_t) < k_t$ we have $\hat{\tau}(k_t) < \tau(k_t)$. It follows from the discussion in the previous paragraph that under adaptive economizing the underestimation of the marginal effect of a decrease of the capital stock is less pronounced for $\tau(k_t) < k^*$. At least close to k^* the proxy of the value function, $\hat{R}(\cdot; k_t)$, is closer to the actual value function W than $\hat{W}(\cdot; k_t)$ not only with respect to value, but, more importantly, also with respect to slope. Here with infinite horizon planning the estimation errors are indeed magnified (as discussed in the introduction) and the linearization of the production function estimate leads to a stronger lineraitzation of the value function proxy than under adaptive economizing. The simplifying assumption about future actions made in the adaptive economizing heuristic leads to a better proxy of the value function, a less extreme underestimation of the negative future effects of current consumption and to decisions that are closer to the optimal ones than intertemporal optimization does. On the other hand, if $\tau(k_t) > k^*$ the incomplete information about f induces an overestimation of the negative marginal effect of current consumption on discounted future utility. Here application of adaptive economizing leads to a proxy for the value function $\hat{R}$ that is again closer to W than $\hat{W}$ with respect to value but further with respect to slope, which means that the induced policy is further away from τ than $\hat{\tau}$. This discussion is illustrated in figure 1 where we depict $W, \hat{W}, R$ and $\hat{R}$. Note that at k^* the graph of R touches W from below and $\hat{R}$ touches $\hat{W}$ from below at a state close to k^* (the fictitious steady state under the estimated growth function $\hat{f}$). Furthermore, R touches $\hat{R}$ from below at k_t. The slope of $\hat{R}$ lies between that of $\hat{W}$ and W or even above that of W for $k_t < k^*$ whereas it is even smaller than that of $\hat{W}$ for $k \in [k^*, k_t]$. This implies directly that with fluctuating optimal policies ($\tau(k_t) < k^*$) adaptive economizing leads to decisions closer to the optimal ones than $\hat{\tau}$. On the other hand, for monotonous optimal policies ($\tau(k_t) \in [k^*, k_t]$) the outcome of adaptive economizing is further away from $\tau(k_t)$ than $\hat{\tau}(k_t)$.

Insert figure 1 here

The discussion above suggests that the exact specification of the way the production function is estimated is not crucial for our qualitative findings. The basic tradeoff pointed out should always exist as long as the optimal policy exhibits fluctuations and the estimated production function is less concave than the actual one.

Numerical examinations we have have carried out with different specifications for utility and production functions confirm the conclusions from this discussion. Indeed typically at steady states where $v_{12} < 0$ we have $\hat{\tau}' < \tau'$ which means that infinite horizon planning leads to a too slow convergence towards the steady state or even destabilizes the fixed point. This suggests that, in cases where capital and consumption are complementary goods (u_{12} sufficiently large), adaptive economizing might generally lead to long run behavior which is qualitatively closer to the optimal path than that of the rolling planning policy and also yield larger discounted utility. We close the section with the discussion of one such numerical example.

Example: We consider a model with wealth effects where the production function is given by[8]

$$f(k) = 1 - (1 - x)^{\beta}$$

and the concave utility function is of the Cobb-Douglas form

$$u(k, c) = k^{\gamma} c^{\epsilon}, \quad \gamma, \epsilon \in (0, 1), \; \gamma + \epsilon = 1.$$

Further, we choose $\delta = 1$ – which implies $k^m = 1$ and $\beta = 9, \; \gamma = 0.8, \; \epsilon = 0.2$. Simple calculations show that the cross derivative of the resulting reduced form utility v_{12} becomes negative for large capital stocks. We consider the case where the discount factor is $\alpha = 0.175$. The positive steady state then is $k^* = 0.414$ and at this steady state we have $f'' - \frac{v_{12}}{u_2} = 0.0026$. This

[8] Although this production function violates $\lim_{k \to 0} f' = \infty$, this is irrelevant for the results reported here since in our examples maximal consumption is never optimal.

implies by proposition 4 that the adaptive economizing policy is flatter than the optimal one for all estimation functions $\hat{f} \in \mathcal{F}$. To illustrate this point, we first show in figure 2 the optimal policy and the adaptive economizing policy for $\hat{f} = f$.

Insert figure 2 here

The slopes at k^* are $\tau'(k^*) = -0.995$ and $\theta'(k^*) = -0.539$. As demonstrated in figure 3 (initial value $k_0 = 0.43$) both optimal and adaptive paths converge in dampening fluctuations towards the steady state where the speed of convergence is much faster for the adaptive economizing path, as predicted by our theoretical findings.

Insert figure 3 here

In the other extreme case where the decision maker does not know the production function and estimates it with a linear function $\hat{f}$ we get figure 4.

Insert figure 4 here

We have to compare three different policies $\tau, \hat{\tau}$ and θ. The slope of the adaptive economizing policy θ now is steeper than in the case with $\hat{f} = f$ but still flatter than the optimal policy ($\theta'(k^*) = -0.732$). On the other hand, the rolling planning policy is much steeper than the optimal one. Since the slope is given by $\hat{\tau}'(k^*) = -1.486$ we even have to expect divergence of the rolling horizon planning paths from the steady state. Looking at the capital accumulation paths (again for $k_0 = 0.43$) this is indeed confirmed (see figure 5). The optimal and the adaptive economizing paths converge towards the steady state whereas the rolling infinite horizon planning path diverges from the steady state and ends up in a period-four cycle oscillating around the steady state.

Insert figure 5 here

The impression that the adaptive economizing path is qualitatively closer to the optimal one than the rolling infinite horizon planning path is confirmed by looking at the discounted utility values generated by the three paths. Whereas the optimal policy generates utility of $U_\tau = 0.550376$ we get $U_\theta = 0.550373$ and $U_{\hat{\tau}} = 0.550368$. So, the adaptive economizing policy generates a strictly higher discounted utility value where the difference is quite significant in the sense that the difference to the optimal utility value is less than half the gap created by rolling infinite horizon planning.

In figure 6 we depict the slopes of the three policies at the positive steady state for discount factors between 0.1 and 1 (in this model there is a unique positive steady state for all these discount factors) and a linear estimation function $\hat{f}$. It can be clearly seen that infinite horizon planning always leads to a steeper policy than optimal and therefore destabilizes the paths at the steady state. On the other hand, two period planning destabilizes the paths only for very small α and for the largest part of the range has a stabilizing effect. The reason that τ' and $\hat{\tau}'$ get very close for large α is that the steady state wanders to the right and for large k the production function is almost linear which means that the deviation of $\hat{f}$ from f in the neighborhood of k^* becomes very small.

Insert figure 6 here

One might wonder how important the rather strong assumption that individuals only have a memory of length one is for the relatively bad performance of rolling infinite horizon planning in our example. To check this aspect we extend the framework considered so far and assume that individuals have a longer memory. First, we consider the case with memory of length two. The prediction of the production function $\hat{f}$ is now given by the piece-wise linear estimation based on output values and rates of return in the previous two periods. This means that at period t the estimation reads

$$\hat{f}(k;\{k_t,k_{t-1}\}) = \min[f(k_t) + f'(k_t)(k-k_t), f(k_{t-1}) + f'(k_{t-1})(k-k_{t-1})].$$

In such a setting the permanent significant changes in the estimation of the production function, which occur in runs with oscillating trajectories and memory length one, should be avoided. The policy functions under adaptive economizing and rolling infinite horizon planning are mappings from $[0,k^m]^2$ into $[0,k^m]$ and we refrain from presenting graphical representations of these mappings here. However, in figure 7 we show the three different trajectories again for $k_0 = 0.43$ but now for memory length two. Qualitatively the picture is similar to that of figure 5.

Insert figure 7 here

The adaptive economizing path again converges –like the optimal one – in dampening fluctuations towards the steady state whereas the rolling infinite horizon planning path exhibits persistent fluctuations. Here no convergence towards a cycle can be observed although the trajectory is quite close to forming a period-two cycle. The changes of the production function estimation along this trajectory are insignificant. The estimation functions used between periods 20 and 40 are all within a range of $3 * 10^{-4}$ where the distances are measured according to the supremum norm. The discounted utility generated by the rolling infinite horizon planning trajectory increases to $U_{\hat{\tau}} = 0.550372$ but is still slightly below that of the adaptive economizing path.

Extending the memory to four periods has virtually no additional effect. We get a a rolling infinite horizon planning trajectory very similar to that with memory two (see figure 8) and exactly the same discounted utility of $U_{\hat{\tau}} = 0.550372$.

Insert figure 8 here

So, in this example the interplay between a 'linearized' estimation of the production function and the length of the planning horizon can be clearly seen. If the decision maker indeed tries to plan over an infinite horizon this leads to unstable trajectories even under (almost) constant production function estimations. As pointed out in the discussion preceeding this numerical example, the reason for the instability is the strong 'linearization' of the value function proxy induced by the (piece-wise) linear estimation of the production function under long horizon planning. Under the short horizon heuristic this effect is less pronounced thereby generating a trajectory that is qualitatively closer to the optimal one and also yielding higher discounted utility. Hence, our theoretical results and also the economic intuition provided above is nicely confirmed.

6 Conclusions

This paper makes three basic points about the implications of short and long horizon planning in one-sector growth models where individuals have incomplete information about the production function. First, in simple settings without wealth effects where the optimal policy is monotonous, fluctuating paths, which have been observed in experiments, are rather due to the assumption of short memory yielding repeated re-estimation of the production function than to a short planning horizon. Second, in simple environments, where the optimal policy is monotonous, long horizon planning is advantageous even when the decision maker has an incorrect model of the world. Third, in more complex settings, where fluctuating paths are optimal, the profitability of long horizon planning is questionable if the model of the decision maker does not exactly match the real world and he revises it over time. We have argued that these results are due to the fact that, on the one hand, depending on whether the future capital stock is below or above the current one, 'linearized' estimations of the production function lead to an under- or overestimation of the future marginal effects of current consumption whereas, on the other hand, adaptive economizing leads to an over- or underestimation of these future marginal effects depending on whether the future capital stock is below or above the steady state. Accordingly depending on whether the optimal path is monotonous or fluctuating at the steady state the two effects might reinforce or weaken each other. Whereas we can not show that generally adaptive economizing will do better than rolling infinite horizon planning in scenarios with fluctuating optimal policies, we have given an example where this is indeed true not only for a memory size of one but also for longer memories. Taking into account the difference in computational costs for the determination of the two policies, this makes a strong case for short horizon planning in complex environments.

References

Bala, V., Majumdar, M. and Mitra, T. (1991), 'Decentralized Evolutionary Mechanisms for Intertemporal Economies: A Possibility Result', *Journal of Economics*, 53, 1-29.

Benhabib, J. and Nishmura, K. (1985), 'Competitive Equilibrium Cycles', *Journal of Economic Theory*, 35, 284 - 306.

Clark, C.W. (1971) 'Economically optimal policies for the utilization of biologically renewable resources', *Mathematical Biosciences*, 12, 245-260.

Dawid, H and Kopel, M. (1997), 'On the Economically Optimal Exploitation of a Renewable Resource: The Case of a Convex Environment and a Convex Return Function', *Journal of Economic Theory*, 76, 272 - 297.

Day, R. (2000), *Complex Economic Dynamics, Volume II*, MIT Press, Cambridge, MA.

Goldman, S.M. (1968), 'Sequential Planning and Continual Planning Revision', *Journal of Political Economy*, 77, 653 - 664.

Kaganovich, M. (1985), 'Efficiency of Sliding Plans in a Linear Model with Time-Dependent Technology', *Review of Economic Studies*, 52, 691 - 702.

Kydland, F. and Prescott, E.C. (1982), 'Time to Build and Aggregate Fluctuations', *Econometrica*, 50, 1345 - 1370.

Majumdar, M. and Mitra, T. (1994), 'Periodic and Chaotic Programs of Intertemporal Allocation in an Aggregate Model with Wealth Effects', *Economic Theory*, 4, 649 - 679.

McKenzie, L.W. (1986), 'Optimal Economic Growth, Turnpike Theorems and Comparative Dynamics'. In: Arrow, K.J., Intrilligator, M.D. (Eds.), *Handbook of Mathematical Economics*, Vol III, Elsevier Science, Amsterdam.

Mitra, T. (1996), 'An Exact Discount Factor Restriction for Period-Three Cycles in Dynamic Optimization Models', *Journal of Economic Theory*, 69, 281-305.

Intrilligator, M.D. (Eds.), Handbook of Mathematical Economics, Vol. III, Elsevier Science, Amsterdam.

Nishimura, K. and Sorger, G. (1996), 'Optimal Cycles and Chaos: A Survey', *Studies in Nonlinear Dynamics and Econometrics*, 1, 11 - 23.

Nishimura, K. and Yano, M. (1996), 'On the Least Upper Bound of Discount Factors That Are Compatible with Optimal Period-Three Cycles', *Journal of Economic Theory*, 69, 306-333.

Noussair, C. and Matheny, K. (2000), 'An Experimental Study of Decisions in Dynamic Optimization Problems', *Economic Theory*, 15, 389 - 419.

Ramsey, F. (1928), 'A Mathematical Theory of Saving', *Economic Journal*, 38, 543-559.

Rust, J. (1994), "Do People Behave According to Bellman's Principle of Optimality?", unpublished manuscript, University of Wisconsin.

Sipper, D., and Bulfin, R.L. (1997), *Production: Planning, Control, and Integration*, McGraw-Hill.

Stokey, N.L. and R.E. Lucas (1989), *Recursive Methods in Economic Dynamics*, Harvard University Press, Cambridge.

Appendix

Proof of Proposition 1:
We first consider the steady states of $\hat{\tau}$. Note that the concavity of u and $\hat{f}$ implies the concavity of the reduced form utility function $v(k_t, k_{t+1}) = u(k_t, (1-\delta)k_t + \hat{f}(k_t; k_0) - k_{t+1})$ for all k_0.
Together with the positivity of the state space and some standard assumptions which are fulfilled in our setup this implies that the Euler equation is a sufficient and necessary optimality condition for optimal paths in our model (see Stokey and Lucas, 1989). Accordingly, for a given $\hat{f}(\cdot; k_0)$ the equation

$$\begin{aligned} &u_2(k, (1-\delta)k + \hat{f}(k) - \tilde{\tau}(k)) - \alpha u_1(\tilde{\tau}(k), (1-\delta)\tilde{\tau}(k) + \hat{f}(\tilde{\tau}(k)) - \tilde{\tau}(\tilde{\tau}(k))) \\ &-\alpha u_2(\tilde{\tau}(k), (1-\delta)\tilde{\tau}(k) + \hat{f}(\tilde{\tau}(k)) - \tilde{\tau}(\tilde{\tau}(k)))(1-\delta + \hat{f}'(\tilde{\tau}(k)) \\ &= 0 \end{aligned} \tag{12}$$

holds for all k and k_0 where $\tilde{\tau}(k; \hat{f}, k_0) > (1-\delta)k$. A state k^* is a steady state of $\hat{\tau}(\cdot; \hat{f})$ if and only if $k^* = \tilde{\tau}(k^*; \hat{f}, k^*)$. Thus, taking into account (4), we see that k^* is a steady state of $\hat{\tau}(\cdot; \hat{f})$ if and only if

$$u_2(k^*, f(k^*) - \delta k^*) - \alpha u_1(k^*, f(k^*) - \delta k^*) - \alpha u_2(k^*, f(k^*) - \delta k^*)(1-\delta + f'(k^*)) = 0.$$

Obviously this condition thus not depend on $\hat{f}$, which implies that the set of fixed points of all maps $\hat{\tau}(\cdot; \hat{f})$ coincide and in particular coincide with the set of fixed points of $\tau(\cdot) = \hat{\tau}(\cdot; f)$.

On the other hand, if we consider the adaptive economizing process the agents have to maximize

$$\max_{k_1 \in [(1-\delta)k, (1-\delta)k + f(k)]} u(k, (1-\delta)k + f(k) - k_1) + \frac{\alpha}{1-\alpha} u(k_1, (1-\delta)k_1 + \hat{f}(k_1) - k_1).$$

The first order condition for this problem yields (after multiplication with $(1-\alpha)$) that

$$\begin{aligned} &(1-\alpha)u_2(k, (1-\delta)k + f(k) - \theta(k)) - \alpha u_1(\theta(k), \hat{f}(\theta(k)) - \delta\theta(k)) \\ &-\alpha u_2(\theta(k), \hat{f}(\theta(k)) - \delta\theta(k))(\hat{f}'(\theta(k)) - \delta) \\ &= 0 \end{aligned} \tag{13}$$

has to hold if $\theta(k) \in ((1-\delta)k, (1-\delta)k + f(k))$. Straightforward calculations show that the concavity of u guarantees that (13) is indeed a condition for a maximum. Using (4) this yields that a capital stock k^* is a fixed point of $\theta(.)$ if and only if it satisfies

$$(1-\alpha)u_2(k^*, f(k^*) - \delta k^*) - \alpha u_1(k^*, f(k^*) - \delta k^*) - \alpha u_2(k^*, f(k^*) - \delta(k^*))(f'(k^*) - \delta) = 0.$$

Obviously, this is exactly the same condition as the condition for a fixed point of the optimal policy and we have shown the proposition. □

Proof of Proposition 2:
Rewriting the first order maximization condition (13) for a utility function $u(c)$ shows that

$$\begin{aligned} &-(1-\alpha)u'((1-\delta)k+f(k)-\theta(k;\hat{f})) \\ &+\ \alpha u'(\hat{f}(\theta(k;\hat{f});k)-\delta\theta(k;\hat{f}))(\hat{f}'(\theta(k;\hat{f});k)-\delta) \\ &=\ 0 \end{aligned} \tag{14}$$

for all $k \in [0,k^m]$ where $\theta(k;\hat{f}) < (1-\delta)k+f(k)$. Total differentiation of this expression, yields after collecting terms

$$\begin{aligned} &\theta'\Big[(1-\alpha)u''((1-\delta)k+f(k)-\theta)+\alpha\hat{f}''(\theta;k)u'(\hat{f}(\theta;k)-\delta\theta) \\ &+\ (\hat{f}(\theta;k)-\delta)^2u''(\hat{f}(\theta;k)-\delta\theta)\Big]-(1-\alpha)(1-\delta+f'(k))u''((1-\delta)k+f(k)-\theta) \\ &+\ \alpha\hat{f}_k'(\theta;k)u'(\hat{f}(\theta;k)-\delta\theta+\alpha(\hat{f}'(\theta;k)-\delta)\hat{f}_k(\theta;k)u''(\hat{f}(\theta;k)-\delta\theta) \\ &=\ 0. \end{aligned} \tag{15}$$

Under our assumption that $\hat{f}(k_1;k)=f(k_1)\ \forall k_1$ we get

$$\theta'(k;f)=\frac{(1-\alpha)u''(1-\delta+f')}{(1-\alpha)u''+\alpha f''u'+\alpha(f'-\delta)^2u''},$$

where we have omitted all functional arguments to keep the expressions as simple as possible. Given our assumptions about f and u it follows that $\theta'(k;f) > 0\ \forall k \in [0,k^m]$ where $\theta(k) < (1-\delta)k+f(k)$. This shows that θ is increasing on the set of capital stocks where adaptive economizing leads to positive consumption. Since θ is also increasing on all intervals where it coincides with $(1-\delta)k+f(k)$ we have shown the proposition. □.

Proof of Proposition 3: Consider first the adaptive economizing policy. Taking into account that $\theta(k^*)=k^*$, $\hat{f}_k(k^*,k^*)=0$, $\hat{f}_k'(k^*,k^*)=f''(k^*)-\hat{f}''(k^*;k^*)$ and $\alpha(1-\delta+f'(k^*))=1$ we get from (15)

$$\theta'(k^*)=\frac{(1-\alpha)u''(c^*)-\alpha^2u'(c^*)(f''(k^*)-\hat{f}''(k^*;k^*)}{(1-\alpha)u''(c^*)+\alpha^2u'(c^*)\hat{f}''(k^*;k^*)}. \tag{16}$$

The inequalities $\theta' < 0$ and $\theta' < -1$ immediately yield the conditions given in the second part of the proposition.

Next, we characterize $\tilde{\tau}'$ and $\tilde{\tau}_k$. The Euler equation for the dynamic optimization problem given the estimation function $\hat{f}(.;k_0)$ reads

$$\begin{aligned} & -u'((1-\delta)k + \hat{f}(k;k_0) - \tilde{\tau}(k;\hat{f},k_0)) + \alpha(1-\delta+\hat{f}'(\tilde{\tau}(k;\hat{f},k_0);k_0)) \\ \times & \quad u'((1-\delta)\tilde{\tau}(k;\hat{f},k_0) + \hat{f}(\tilde{\tau}(k;\hat{f},k_0);k_0) - \tilde{\tau}(\tilde{\tau}(k;\hat{f},k_0);\hat{f},k_0)) \\ = & \quad 0. \end{aligned} \tag{17}$$

This equality has to hold for all values of k and k_0 such that $\tilde{\tau}(k;\hat{f},k_0) < (1-\delta)k + \hat{f}(k;k_0)$. In particular it has to hold for all values of k and k_0 in the neighborhood of k^*. Total differentiation of this equality with respect to k yields for $k = k^*,\ k_0 = k^*$

$$-u'' + ((1+\alpha)u'' + \alpha^2 u' \hat{f}'')\tilde{\tau}' - \alpha u'' \tilde{\tau}'^2 = 0 \tag{18}$$

The left hand side is positive for $\tilde{\tau} = 0$ and convex. For $\tilde{\tau} = 1$ the left hand side gives $\alpha^2 u' \hat{f}''$, which is non-positive and negative for non-linear estimation functions. Therefore, this quadratic equation has two positive solutions where one of them is in $(0,1]$. Since the slope of the policy function at the unique steady state cannot be larger than one this shows that $\tilde{\tau}'(k^*;\hat{f},k^*) \in (0,1)$. Furthermore, if $\hat{f}''$ increases the constant and the coefficient of $\tilde{\tau}'^2$ in the quadratic equation are unaffected whereas the coefficient of $\tilde{\tau}'$ – which is negative – increases. This implies that the solutions of (18) in $[0,1]$ increases as $\hat{f}''$ increases. Hence,

$$\tilde{\tau}'(k^*;\hat{f},k^*) > \tilde{\tau}'(k^*;f,k^*) = \tau'(k^*)$$

for all $\hat{f}$ with $\hat{f}''(k^*;k^*) \neq f''(k^*)$. Total differentiation of (17) with respect to k_0 yields

$$\tilde{\tau}_k(k^*;\hat{f},k^*) = \frac{\alpha^2 u'(f'' - \hat{f}'')}{u''(\alpha\tilde{\tau}'(k^*;\hat{f},k^*) - 1) - \alpha^2 u' \hat{f}''}.$$

In order to compare θ' with $\hat{\tau}'$ we rewrite (16) as

$$\theta'(k^*) = \frac{(1-\alpha)u''}{(1-\alpha)u'' + \alpha^2 u' \hat{f}''} + \frac{\alpha^2 u'(f'' - \hat{f}'')}{(\alpha-1)u'' - \alpha^2 u' \hat{f}''}.$$

Since $\tilde{\tau}' < 1$ for $\hat{f}'' < 0$ and $\tilde{\tau}' = 1$ for $\hat{f}'' = 0$, the second term of this sum is always smaller or equal than $\tilde{\tau}_k$ where equality holds only for $\hat{f}'' = 0$. To show that the first term is smaller or equal than $\tilde{\tau}'$ we insert it for $\tilde{\tau}'$ into the left hand side of (18). This gives

$$\frac{-\alpha^5 u'^2 u'' \hat{f}''^2}{((1-\alpha)u'' + \alpha^2 u' \hat{f}'')^2} \geq 0,$$

again with equality only for $\hat{f}'' = 0$. This implies

$$\tilde{\tau}' \geq \frac{(1-\alpha)u''}{(1+\alpha-2\alpha\tilde{\tau}')u'' + \alpha^2 u' \hat{f}''}$$

with equality only for $\hat{f}'' = 0$ and we have shown the first of the two inequalities in (10). Finally, we show the second inequality in (10). By the means of implicit differentiation we get from (18)

$$\frac{\partial \tilde{\tau}'}{\partial \hat{f}''} = -\frac{\alpha^2 u' \tilde{\tau}'}{(1-\alpha)u'' + \alpha^2 u' \hat{f}''} > 0$$

Differentiating $\hat{\tau}'$ with respect to $\hat{f}''$ gives after collecting terms

$$\begin{aligned}
&\frac{d\hat{\tau}'}{d\hat{f}''} \\
&= \frac{d(\tilde{\tau}' + \tilde{\tau}_k)}{d\hat{f}''} \\
&= \left(1 + \frac{\partial \tilde{\tau}_k}{\partial \tilde{\tau}'}\right) \frac{\partial \tilde{\tau}'}{\partial \hat{f}''} + \frac{\alpha^2 u'(u''(1-\alpha\tilde{\tau}') + \alpha^2 u' f'') + \alpha u'^2 (f'' - \hat{f}'')}{(u''(1-\alpha\tilde{\tau}') + \alpha u' \hat{f}'')^2} \\
&\leq \frac{\partial \tilde{\tau}'}{\partial \hat{f}''} + \frac{\alpha^2 u'}{u''(1-\alpha\tilde{\tau}') + \alpha u' \hat{f}''} \\
&= \alpha^2 u' \left(\frac{\tilde{\tau}'}{(1+\alpha-2\alpha\tilde{\tau}')u'' + \alpha^2 u' \hat{f}''} + \frac{1}{u''(1-\alpha\tilde{\tau}') + \alpha u' \hat{f}''} \right)
\end{aligned}$$

where the inequality in line 4 is strict if $\hat{f}'' > f''$. It is easy to see that the bracket in the last line is strictly negative for $\tilde{\tau}' \in [0,1)$ and zero for $\tilde{\tau}' = 1$. We know that $\tilde{\tau}' = 1$ holds only for $\hat{f}'' = 0$ and hence we have shown that

$$\frac{d\hat{\tau}'(k^*; \hat{f})}{d\hat{f}''(k^*; k^*)} < 0 \quad \forall \hat{f} \in F.$$

Keeping in mind that $\tau'(k^*) = \hat{\tau}'(k^*; f)$ establishes the second part of (10). □

Proof of Corollary 1:
The expression for θ' can be obtained by inserting $\hat{f}'' = 0$ into (16). From (18) it follows that $\tilde{\tau}' = 1$ for $\hat{f}'' = 0$ and using this value we get $\tilde{\tau} = -\frac{\alpha^2 u' f''}{(1-\alpha)u''}$. □

Proof of Proposition 4: In what follows we again omit all functional arguments with the understanding that all expressions are evaluated at $k^* = \theta(k^*) = \tau(k^*)$. Total differentiation of (13) with respect to k gives

$$\begin{aligned}&-\theta'(\alpha u_{11} + \alpha 2(f' - \delta)u_{12} + (1 - \alpha + \alpha(\hat{f}' - \delta)^2)u_{22} + \alpha u_2 \hat{f}'')\\&+(1 - \alpha)u_{12} + (1 - \alpha)(1 - \delta + f')u_{22} - \alpha u_2 \hat{f}'_k\\&= 0\end{aligned}$$

Using (6) yields

$$\theta' = \frac{-(1 - \alpha)v_{12} - \alpha u_2(f'' - \hat{f}'')}{\alpha u_{11} + 2\alpha(f' - \delta)u_{12} + \alpha(f' - \delta)^2 u_{22} + (1 - \alpha)u_{22} + \alpha u_2 \hat{f}''}.$$

Obviously the numerator is positive if $v_{12} < 0$. Concavity of u implies that the sum of the first three terms in the denominator is negative which implies that the whole denominator is negative. Thus, $\theta' < 0$. It is further easy to see from this expression that θ' decreases with increasing $\hat{f}''$.

Total differentiation of the Euler equation (12) of the dynamic optimization problem with perfect knowledge about f shows that the slope of the τ at the steady state k^* has to be a root of

$$\xi^f_{op}(x) := -\alpha v_{12}x^2 - (\alpha u_{11} - \alpha(1-\delta+f')v_{12} + \alpha(1-\delta+f')u_{12} + u_{22} + \alpha u_2 f'')x - v_{12}$$

Concavity of u implies that the coefficient of x is positive and hence this function has two negative roots for $v_{12} < 0$. The larger of the two converges to 0 as v_{12} goes to zero and since we know that $\tau' > 0$ for $v_{12} > 0$ a continuity argument establishes that τ' is given by the larger of the two negative roots of ξ^f_{op}. Furthermore, we have

$$\begin{aligned}&\xi^f_{op}(\theta')\\&= -[\alpha u_{11} + 2\alpha(f' - \delta)u_{12} + (1 - \alpha + \alpha(f' - \delta)^2)u_{22} + \alpha u_2 \hat{f}'']\theta'\\&\quad +2\alpha v_{12}\theta' - \alpha u_2(f'' - \hat{f}'')\theta' - v_{12} - \alpha v_{12}\theta'^2\\&= -[-(1 - \alpha)v_{12} - \alpha u_2(f'' - \hat{f}'')] + 2\alpha v_{12}\theta' - \alpha u_2(f'' - \hat{f}'')\theta' - v_{12} - \alpha v_{12}\theta'^2\\&= -\alpha v_{12}(1 - \theta')^2 + \alpha u_2(f'' - \hat{f}'')(1 - \theta')\end{aligned}$$

Since we know that $\theta' < 0$ condition (11) guarantees that $\xi_{op}^{f}(\theta') > 0$. Thus, θ' either has to be larger than the larger root of ξ_{op}^{f} or smaller than the smaller root of this function. We know that for $\hat{f}'' = f''$ and $v_{12} = 0$ the slope of θ is zero which implies by continuity that $\theta' > \tau'$ if (11) holds. $\square$

Proof of Proposition 5:
Similar calculations to those in the proof of proposition 4 show that given an estimation function $\hat{f}(\cdot; k^*)$ the slope of the policy function which solves the corresponding dynamic optimization problem has to be the larger root of

$$\xi_{op}^{\hat{f},k}(x) := -\alpha v_{12}x^2 - (\alpha u_{11} - \alpha(1-\delta+f')v_{12} + \alpha(1-\delta+f')u_{12} + u_{22} + \alpha u_2 \hat{f}'')x - v_{12}.$$

Since we know that $\tilde{\tau}' < 0$ for $v_{12} < 0$, we know that this function has at least one negative root and thus the coefficient of x is negative. The coefficient of x decreases with increasing $\hat{f}''$ whereas the constant and the coefficient of x^2 are unaffected. Thus the largest root of $\xi_{op}^{\hat{f},k^*}$ is smaller than that of ξ_{op}^{f} if $\hat{f}''(k^*) > f''(k^*)$. $\square$

Figure Captions

Figure 1: Comparison of the actual value function W (bold line), the proxy for the value function used in adaptive economizing R (sparsely dotted line), the estimated value function $\hat{W}$ (dotted line) and the estimated proxy $\hat{R}$ (solid line) in the neighbourhood of k^*.(linear estimation function $\hat{f}(\cdot;k_t), f(k) = 1-(1-k)^4, u(k,c) = k^{0.8}c^{0.2}, k_t = 0.8, \delta = 1, \alpha = 0.6$).

Figure 2: Comparison of the optimal (solid line) and the adaptive economizing (bold line) policy in a model with wealth effects and a correct estimation function $\hat{f} = f$. The production function is given by the sparsely dotted line.

Figure 3: Comparison of the capital accumulation paths generated by the optimal (solid line) and the adaptive economizing (bold line) policies given in figure 2. Initial capital stock is $k_0 = 0.43$.

Figure 4: Comparison of the optimal (solid line), the rolling planning (dotted line) and the adaptive economizing (bold line) policy in a model with wealth effects and a linear estimation function $\hat{f}$. The production function is given by the sparsely dotted line.

Figure 5: Comparison of the capital accumulation paths generated by the optimal (solid line), the rolling planning (dotted line) and the adaptive economizing (bold line) policies given in figure 4. Initial capital stock is $k_0 = 0.43$.

Figure 6: Comparison of the slope of the optimal (solid line), the rolling planning (dotted line) and the adaptive economizing (bold line) policies at the positive steady state for linear estimation functions $\hat{f}$ and $\alpha \in [0.1, 1]$.

Figure 7: Comparison of the capital accumulation paths generated by the optimal (solid line) policy, rolling planning (dotted line) and adaptive economizing (bold line) policies for a memory of two periods. Initial capital stock is $k_0 = 0.43$.

Figure 8: Comparison of the capital accumulation paths generated by the optimal (solid line) policy, rolling planning (dotted line) and adaptive economizing (bold line) policies for a memory of four periods. Initial capital stock is $k_0 = 0.43$.

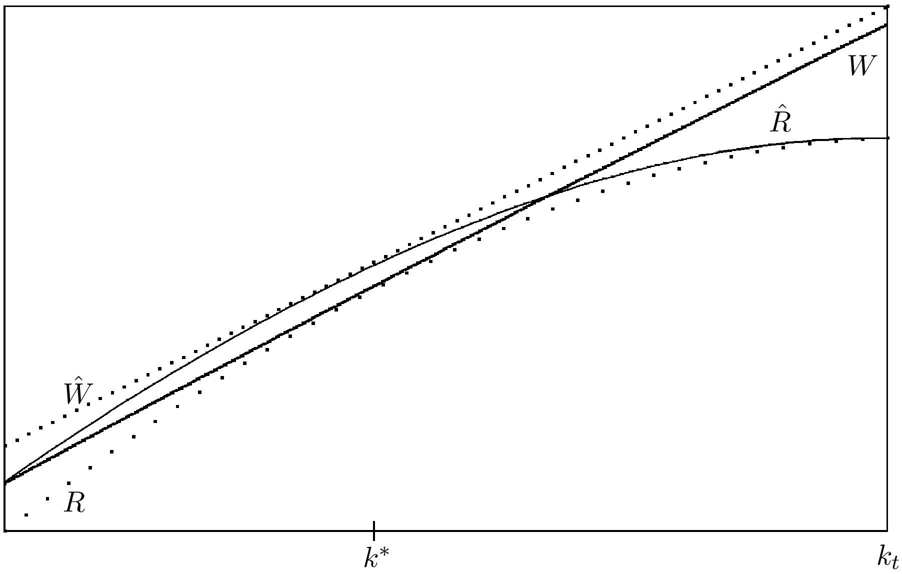

Figure 1

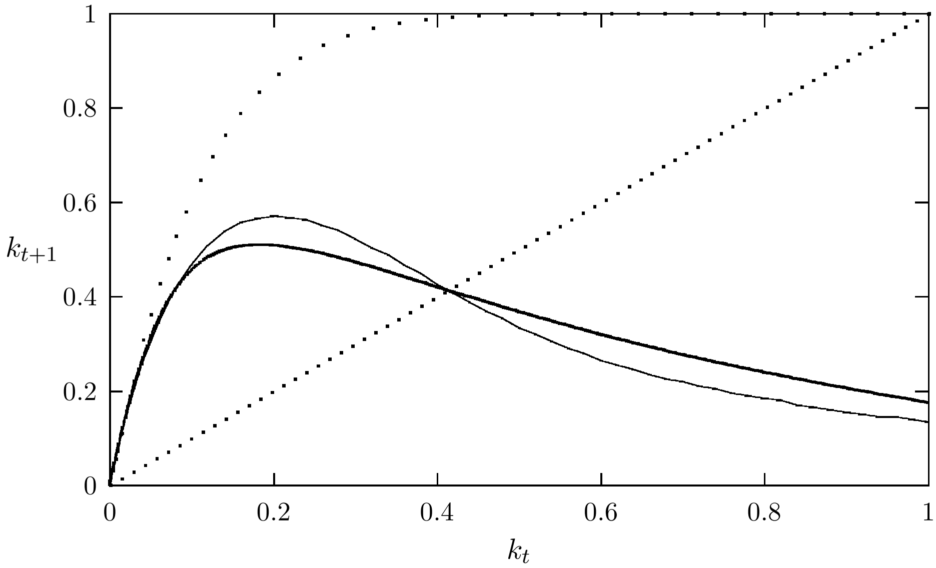

Figure 2

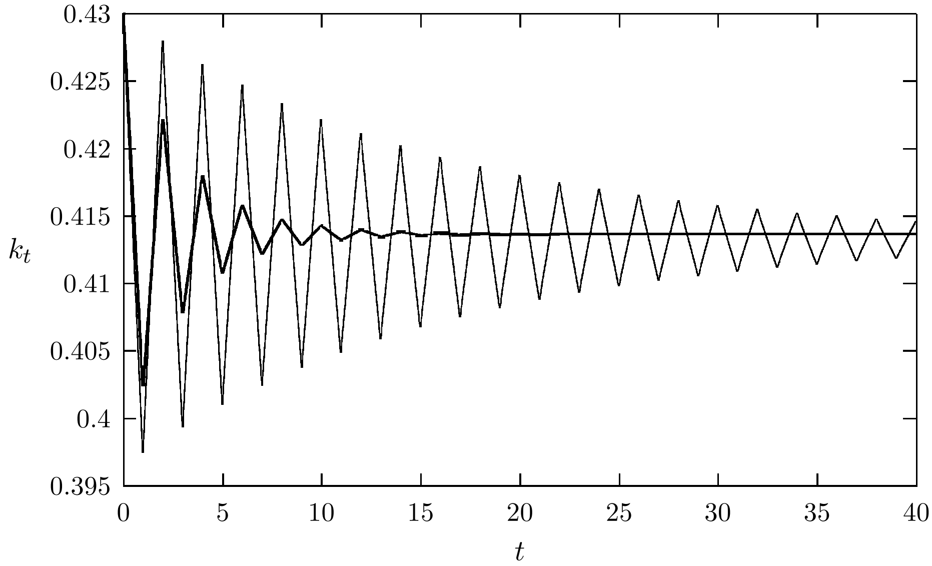

Figure 3

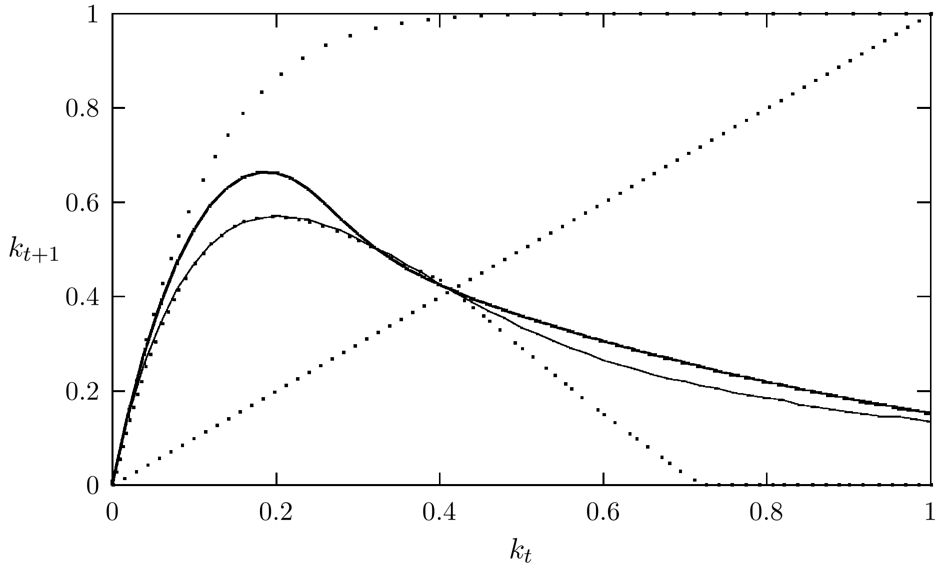

Figure 4

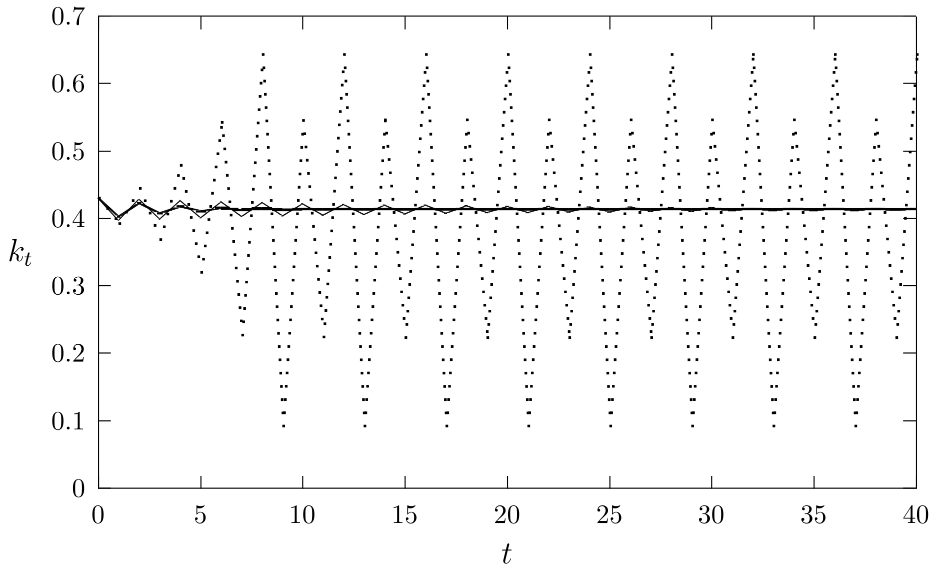

Figure 5

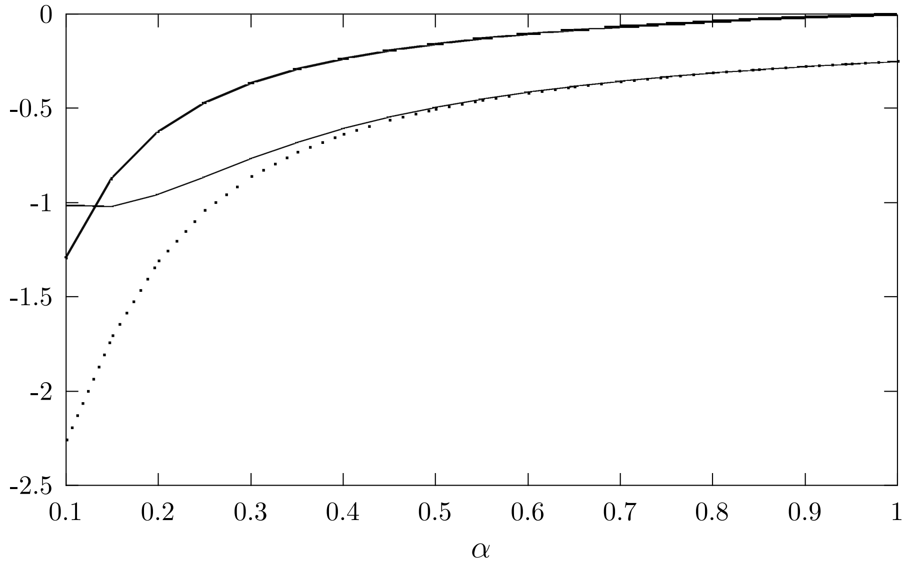

Figure 6

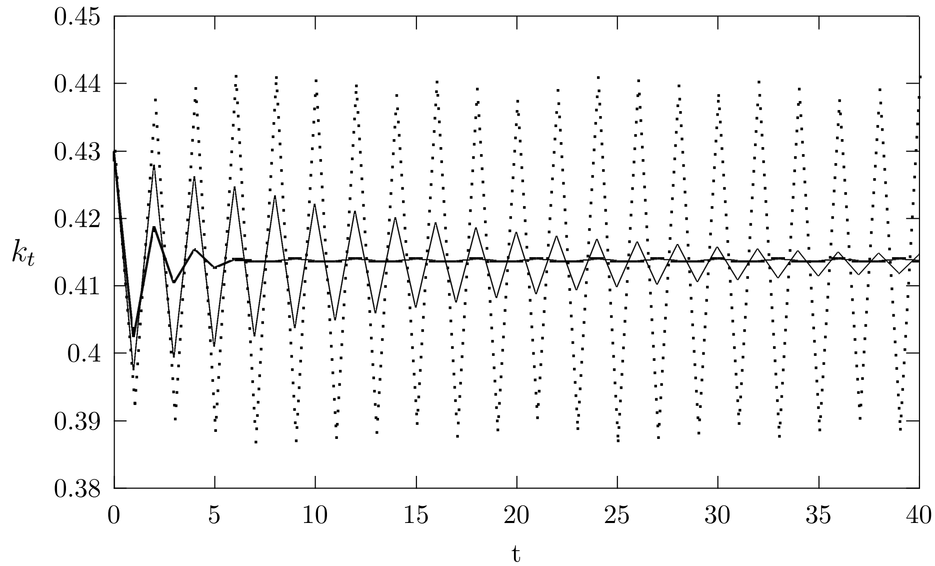

Figure 7

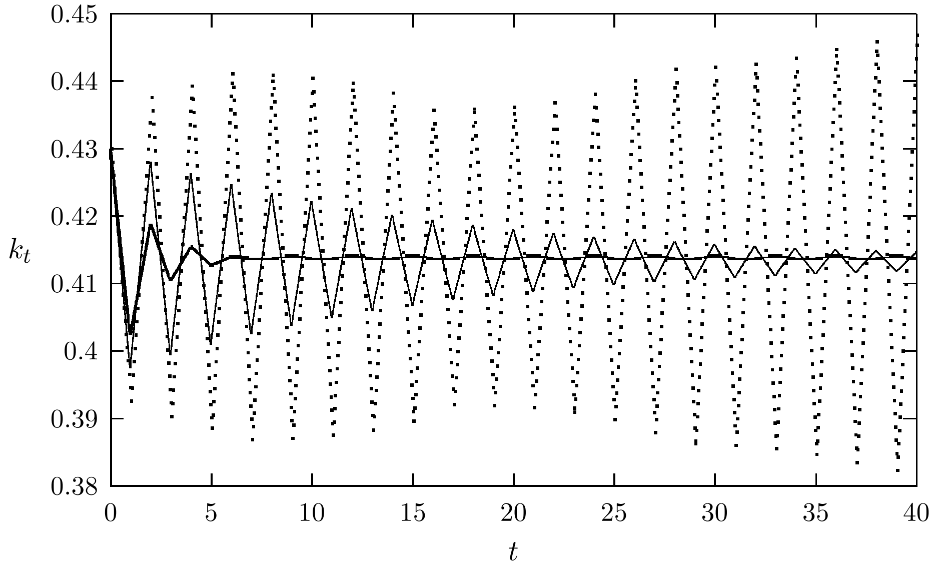

Figure 8