NHH

INSTITUTT FOR FORETAKSØKONOMI

DEPARTMENT OF BUSINESS AND MANAGEMENT SCIENCE

FOR 19 2016

ISSN: 1500-4066

November 2016

Discussion paper

A Bellman approach to periodic optimization problems

BY

Sturla F. Kvamsdal, José M. Maroto, Manuel Morán AND **Leif K. Sandal**

Norges Handelshøyskole

NORWEGIAN SCHOOL OF ECONOMICS

A Bellman approach to periodic optimization problems

Sturla F. Kvamsdal[a], *José M. Maroto*[b,d,]*, *Manuel Morán*[c,d], *Leif K. Sandal*[e]

[a] SNF – Centre for Applied Research, Helleveien 30, N-5045 Bergen, Norway

[b] Department of Estadística e Investigación Operativa II, Universidad Complutense, 28223 Madrid, Spain

[c] Department of Fundamentos del Análisis Económico I, Universidad Complutense, 28223 Madrid, Spain

[d] IMI – Institute of Interdisciplinary Mathematics, Universidad Complutense, 28223 Madrid, Spain

[e] Department of Business and Management Science, NHH Norwegian School of Economics, N-5045 Bergen, Norway

*Corresponding author: maroto@ccee.ucm.es

November 25, 2016

Abstract

We consider an infinite horizon optimization problem with arbitrary but finite periodicity in discrete time. The problem can be formulated as a fix-point problem for a contraction operator, and we provide a solution scheme for this class of problems. Our approach is an extension of the classical Bellman problem to the special case of non-autonomy that periodicity represents. Solving such problems paves the way for consistent and rigorous treatment of, for example, seasonality in discrete dynamic optimization. In an illustrative example, we consider the decision problem in a fishery with seasonal fluctuations. The example demonstrates that rigorous treatment of periodicity has profound influence on the optimal policy dynamics compared to the case where seasonality is abstracted from by considering average effects only.

Key words: Bellman, optimization, periodicity, contraction operator, solution scheme.

1. Introduction

Periodicity is an important characteristic of many systems that are subject to control. Examples include demand systems subject to supply control (that is, demand for newspapers likely have a daily periodicity; demand for winter garments likely has an annual periodicity), transport and logistics systems subject to routing control (that is, transport and logistics systems likely has periodicity in stress), or natural systems subject to management control (renewable resources such as fish stocks may have periodicity in growth or other natural processes as well as periodicity in prices and costs). Periodicity is a special type of non-autonomy, and non-autonomy typically renders many optimal control problems difficult and costly to deal with or even intractable. Thus, periodicity is either treated in some ad-hoc manner or abstracted from altogether, for example by considering aggregate or mean forcings. To our knowledge, periodicity in infinite horizon optimal control problems in discrete time is not treated formally in the theoretical literature. We show that the problem can be formulated as a fixed-point problem for a contraction operator, and further provide a solution scheme. As such, we extend the classical Bellman problem to include the special case of non-autonomy that periodicity represents.

To illustrate our approach to periodic problems, we develop a numerical [discretization] algorithm that we apply to a stylized fisheries management problem with seasonal fluctuations. Our routine exploits the equivalence between the low-dimensional, coupled equations problem formulation and the high-dimensional, single equation problem formulation. The example serves to demonstrate the feasibility of our approach, and also to suggest significant, practical implications of taking periodicity explicitly into account. We contrast the solution of the periodic problem to the solution of a problem that abstracts from the inherent periodicity.

Given the prevalence of periodic characteristic of many systems subject to control, we think our contribution is important and valuable. We show that the classical Bellman problem approach can be extended to periodic problems, and that this extension is – both conceptually and numerically – feasible and practical.

2. Contraction of the periodic problem

A general, infinite horizon, autonomous, dynamic, discounted, discrete-time optimization problem considers the following:

$$\max_{\{u_k\}_{k=0}^{\infty}} \sum_{k=0}^{\infty} \beta^{k+1} \cdot \Pi(x_k, u_k) \tag{1}$$

such that $x_{k+1} = F(x_k, u_k)$, $u_k \in U(x_k)$, $k = 0, 1, 2, \ldots$, and $x_0 \in X$ given. $0 < \beta < 1$ is a discount factor. $X \subset \mathbb{R}^n$ is the system state space, x_k is the dynamic state variable at the beginning of period k; y_k is the state variable at the end of period k and is identical to the state variable at the beginning of period $k + 1$. $\Pi: \{\mathbb{R}^n \times \mathbb{R}^p\} \rightarrow \mathbb{R}$ is bounded and continuous and gives the return at the end of each period. $F: \{\mathbb{R}^n \times \mathbb{R}^p\} \rightarrow \mathbb{R}^n$ is a continuous dynamic operator that governs the state variable such that $x_{k+1} = y_k$ is the state at the beginning of period $k + 1$. $U: \mathbb{R}^n \rightarrow \mathbb{R}^p$ is a nonempty and compact valued correspondence that specifies the feasible controls u_k from the state x_k. That is, u_k is the decision or control variable that has to be decided for each instant of the infinite time sequence $\{t_0, t_1, t_2, \ldots\}$. With these conditions in place, optimal controls $\{u_k^*\}_{k=0}^{\infty}$ and corresponding paths $\{x_k^*\}_{k=0}^{\infty}$ exist, as does the value function of the problem, $V(x) = \sum_{k=0}^{\infty} \beta^{k+1} \cdot \Pi(x_k^*, u_k^*)$ with $x = x_0$. The value function is the fixed point of the Bellman operator, which is defined on the space $C(X)$ or real, bounded, and continuous functions on X and given by $AV = \max_{u \in U(x)} \{\beta \cdot \Pi(x, u) + \beta \cdot V(y)\}$, with $V \in C(X)$ and $y = F(x, u)$. See Bertsekas (2001) for a general treatment of problems of type (1).

Now consider the non-autonomous but periodic problem where $\Pi_k(x, u)$ is the return function and $F_k(x, u)$ is the dynamic operator for period k. Feasible sets for the state and control variables may also vary with period, such that U_k and X_k are the feasible control and state variable sets for period k. The control set may vary with the state such that we have $U_k = U_k(x_k)$, but we typically omit the state argument. The problem is periodic in the sense that for a finite integer $N \geq 2$, we have $\Pi_k(x, u) = \Pi_{k+N}(x, u)$, $\mathrm{F}_k(x, u) = \mathrm{F}_{k+N}(x, u)$, $X_k = X_{k+N}$, and $U_k = U_{k+N}$. We say that the performance or return measure and the dynamic constraint functionally repeats itself. To avoid notational mess, we redefine the state and control spaces from above as follows: $X = \cup_{k=1}^{N} X_k$, $U = \cup_{k=1}^{N} U_k$. Without adding complexity, we can allow for varying period length. Thus, each different period has potentially different discount factor values. We write the length of period k as $T_k = t_k - t_{k-1}$ and its discount factor as β_k. Periodicity implies $T_k = T_{k+N}$ and $\beta_k = \beta_{k+N}$. The length of the cycle of N periods is then $T = \sum_{i=1}^{N} T_i = t_N - t_0$, and the discount factor for the cycle of N periods is $\beta = \prod_{i=1}^{N} \beta_i$. Figure 1 accounts for period index references.

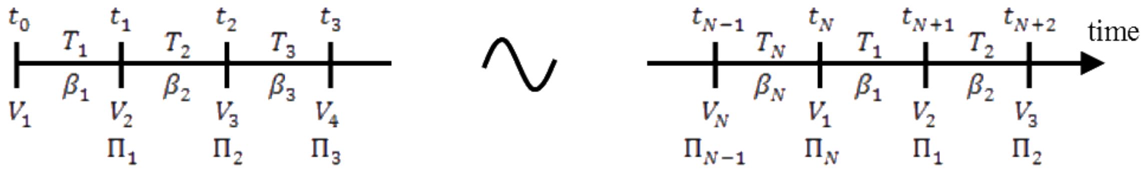

Figure 1: Period index reference for periodic problems. Note that the return (Π_k) is yielded at the end of period k, but that V_k refers to the beginning of period k.

Although a real discounted problem cannot have periodic present value, the running value will be periodic under a constant per time discounting if involved operators (Π_k, F_k) or spaces (U_k, X_k) are periodic. That is, for a problem of type (1) to be periodic, one or more of Π_k, F_k, U_k, and X_k need to have a periodic feature. As suggested above, a periodic feature is such that it repeats itself with some inherent period. If a problem includes several periodic features, the problem period N has to be the least common multiple of the potentially different inherent periods of the different features. (As noted above, we have $N \geq 2$ and thus abstracts from the trivial case with $N = 1$.)

If the optimal control $\{u_k^*\}$ was known, and for a generic state $x_k = x$ at the beginning of the period, we have the current value function V_k:

$$\begin{aligned} V_k(x) &= \beta_k \Pi_k(x, u_k^*) + \beta_k V_{k+1}(F_k(x, u_k^*)), \;\; k = 1, \dots, N-1 \\ V_N(x) &= \beta_N \Pi_N(x, u_N^*) + \beta_N V_1(F_N(x, u_N^*)) \end{aligned} \tag{2}$$

Note that (2) consist of N nested equations that follows from value additivity, that is, that the value at the beginning of a period is equal to the return from that period plus the discounted value at the beginning of the next period. Note also that since the return is yielded at the end of periods (see figure 1), the current period return is also discounted. Because it sometimes is customary to define the control to be the state variable at the beginning of the next period, we rather use $y_k = x_{k+1} = F_k(x_k, u_k)$ as control variable, with $y_k \in Y_k$, $Y_k = Y_{k+N}$ (periodicity), and $Y = \bigcup_{k=1}^{N} Y_k$. For convenience, we simply replace the control argument in the return function and in what follows write $\Pi_k(x, y_k)$. Equation (2) is by definition equivalent to the following set of equations:

$$\begin{aligned} V_k(x) &= \max_{y_k \in Y_k(x)} \{\beta_k \Pi_k(x, y_k) + \beta_k V_{k+1}(y_k)\}, \;\; k = 1, \dots, N-1 \\ V_N(x) &= \max_{y_N \in Y_N(x)} \{\beta_N \Pi_N(x, y_N) + \beta_N V_1(y_N)\} \end{aligned} \tag{3}$$

To see how the optimal control can be derived, we consider the following operator for any bounded continuous function $W \in C(X)$ in the state space:

$$B_k\, W(x) = \max_{y_k \in Y_k(x)} \{\beta_k\, \Pi_k(x, y_k) + \beta_k W(y_k)\} \tag{4}$$

We define $\hat{\beta}_i = \prod_{j=1}^{i} \beta_j$ and further consider the nested operator for k periods, $\hat{B}_k$:

$$\begin{aligned} \hat{B}_k W(x) &= B_k \circ B_{k-1} \circ \dots \circ B_1 W(x) \\ &= \max_{(y_1, y_2, \dots, y_k) \in \Gamma_k(x)} \left\{ \sum_{i=1}^{k} \hat{\beta}_i \Pi_i(y_{i-1}, y_i) + \hat{\beta}_k W(y_k) \right\} \end{aligned} \tag{5}$$

In (5), $y_0 = x$ and $\Gamma_k(x) = \{(y_1, y_2, \dots, y_k) \mid (y_1, y_2, \dots, y_k) \in (Y_1(x), Y_2(y_1), \dots, Y_k(y_{k-1}))\}$. Note that the correspondences $\Gamma_k(x)$ are compact valued if the Y_i are compact. $\Gamma_k(x)$ is the set of k-step feasible paths starting from x.

Let $L(x, y)$ be defined by the argument of the maximum operator in (5). The nested operator (5) is well defined if $L: X \times X^k \to \mathbb{R}$ is continuous and $\Gamma_k: X \to X^k$ is continuous and compact valued. By the Theorem of the Maximum, the operator $\hat{B}_k: C(X) \to C(X)$ is well defined. It follows directly from the classical Bellman problem (Bellman 1957; see also Stokey *et al.* 1989) that this operator is contractive with contraction factor $\hat{\beta}_k$. Alternatively, it is trivial to check that $\hat{B}_k$ satisfies Blackwell's sufficient conditions (Blackwell 1965; see also Stokey *et al.* 1989). In particular, we write $\beta = \hat{\beta}_N$ and consider the N-cycle operator:

$$\hat{B}_N W(x) = \max_{y \in Y(x)} \left\{ \sum_{i=1}^{N} \hat{\beta}_i \Pi_i(y_{i-1}, y_i) + \beta W(y_N) \right\} \tag{6}$$

In (6), we write $y = (y_1, \dots, y_N)$ and $\Gamma_N = Y(x)$. It is trivial to show that the argument of the maximum operator in (5) is continuous and bounded if W and the Π_i are continuous and bounded, and also that $Y(x)$ has compact range if the Y_i are compact. We define

$$\hat{\Pi}_k(x, y) = \beta_k \Pi_k(x, y_k) + \beta_k \beta_{k+1} \Pi_{k+1}(y_k, y_{k+1}) + \dots + \beta \Pi_{k+N}(y_{k+N-1}, y_{k+N}) \tag{7}$$

Equation (6) can now be written

$$\hat{B}_N\, W(x) = \max_{y \in Y(x)} \{\hat{\Pi}_1(x, y) + \beta W(y_N)\} \tag{8}$$

More generally, for the N-cycle operator that starts in period k, written $\hat{B}_{N,k}$ and defined by

$$\hat{B}_{N,k} = B_{k-1} \circ B_{k-2} \circ \dots \circ B_1 \circ B_N \circ B_{N-1} \circ \dots \circ B_k = \hat{B}_{k-1} \circ B_N \circ B_{N-1} \circ \dots \circ B_k \tag{9}$$

That is,

$$\hat{B}_{N,k} W(x) = \max_{y \in Y(x)} \{\hat{\Pi}_k(x, y) + \beta W(y_{k-1})\}, \quad k = 1, \dots, N \tag{10}$$

In (10), $y_0 = y_N$. Equation (10) satisfies Blackwell's sufficient conditions and is thus a contraction with contraction factor β. Let the unique fix points for $\hat{B}_{N,k}$, $k \in \{1, \dots, N\}$, be

$W_k^*(x)$. It follows directly from the definitions that $W_k^*(x) = \hat{B}_{k-1}W_1^*(x)$ is the unique fix point for the N-cycle starting in period k.

Finally, note that (3) implies (10). Thus, our unique set of N fix points must constitute the proper non-autonomous value function and hence solve (3) or vice versa. We have now proved the following proposition:

Proposition: The infinite horizon, N-period optimization problem represented by (3) is well defined provided that $\{\Pi_k, V_k, Y_k\},\ \ k \in \{1, \ldots, N\}$ are continuous and bounded and $\{Y_k\}$ have compact range. Moreover, solving (3) is equivalent with solving the contraction problem (10) and hence has a unique solution for the cycle values, that is, a unique, non-autonomous (periodic) value function $V(k, x) = V_k(x)$.

With regard to boundedness, the proposition can be generalized in the sense of Rincón-Zapatero and Rodriguez-Palmero (2003); see also the related discussion on boundedness in Stokey *et al.* (1989).

Varying period length require suitable adaptions of Π_k, F_k, X, and Y_k, as well as the following specification of β_k. If period k represents a share δ_k of the N-cycle, such that $t_k - t_{k-1} = \delta_k \cdot (t_N - t_0)$, we have $\beta_k = \beta^{\delta_k}$. In many applications, the N-cycle represents a year, and β is then the annual discount factor. The extension to varying period length is an important and useful extension, not least because it allows for reductions in dimensionality. To see this, consider a problem that is formulated on an annual level, but where one month is different such that the problem is periodic. Without the option of varying period length, the model would have $N = 12$. With varying period length, $N = 2$ suffices.

With the above proposition in place, we are equipped to deal with a wide range of optimization problems. That (3) implies (10) and that (10) has a unique solution means that (3) also has a unique solution. Thus, we can work directly with (3). This is useful in numerical schemes. Equation (10) maximizes over an N dimensional (vector) space Y, and numerical solutions can be cumbersome and costly to obtain. In contrast, the equations in (3) are N coupled equations that each maximize over a one-dimensional (vector) space, where established and reliable numerical routines converges fast. A numerical scheme can exploit this coupling and the fast convergence of the equations in (3). Technically, for an arbitrary state x_k at the beginning of period k, we derive the optimal state y_k at the beginning of period $k + 1$. (With periodicity, for $k = N$, we derive y_1.) When the N optimal controls y_k has been found for all k, the optimal decision for any time and initial x_0 in any initial period. Without loss of generality, we may assume that the initial period is the $k = 1$ period. Then, for example, with initial $x_1 = x_0$, the optimal path is obtained as follows:

$$x_0 \to y_1 = x_2 \to y_2 = x_3 \to \cdots \to y_{N-1} = x_N \to y_N = x_1 \to y_1 = x_2 \to \cdots \quad (11)$$

That is, (11) yields the chain of optimal decisions $\{y(t = t_i)\}$, for all i, depending on the initial x_0, and thus also the optimal path $\{x_i\}$.

We have established a numerical routine based on the above proposition – using the set of equations in (3) – and further the inherent logic in (11). Below, we apply this routine to an applied example that suggests that taking account of periodicity may have significant practical implications. The numerical results were obtained from code written in standard FORTRAN.

3. An example

We illustrate the use of our method with a fishery management problem. To make sure our parameter values are grounded in the real world, we consider a discrete time model of Barents Sea capelin. This discrete time model corresponds to the continuous time model for the Barents Sea capelin fishery studied in Agnarsson *et al.* (2008), who established empirical parameter values and suitable functional forms. See Kvamsdal *et al.* (2015) for a detailed discussion of how to properly set up a corresponding discrete time model for a given continuous time model. Because the discrete time functional expressions are complex, and because the specific functions are of minor interest here, we do not delve into details but refer interested readers to the relevant papers. (The details are also available from the authors upon request.) For our purpose, it suffices to say that the return function Π_k and the dynamic operator F_k are both continuous on the compact state space X, and that the control space $Y(x)$ has compact range for all x.

We assume the fishery has seasonal differences in the harvest cost parameter; say, the summer quarter has favorable conditions. The remaining quarters have normal conditions. For simplicity, this variability in costs is the only periodic feature in our example. We thus have $N = 2$, with period 1 being the low cost period with cost parameter $c_{LC} = c_{NC}/2$. Regarding notation, we denote the low cost period with subscript LC and the normal cost period with subscript NC. The length of the 2 period cycle is one year. The two periods hasve different length, with $\delta_{LC} = 1/4$ and $\delta_{NC} = 3/4$. From the set of equations corresponding to (3) above, we derive the optimal escapement biomass levels – the optimal control rules – for the two periods as functions of the stock level at the beginning of a given period: $y_{LC}(x)$ and $y_{NC}(x)$.

To illustrate the potential impact of abstracting from the periodic feature, we also consider the non-periodic (stationary) management problem, that is, of type (1), where the cost parameter is the annual average: $c_{LC} \cdot \delta_{LC} + c_{NC} \cdot \delta_{NC} = c_{NC} \cdot 7/8$. We denote the non-periodic, annual case with subscript A: $c_A = c_{NC} \cdot 7/8$. We thus derive the annual optimal escapement biomass level as a function of the stock level at the beginning of the year: $y_A(x)$.

Figure 2 (top panel) compares the period specific escapement rules $y_{LC}(x)$ and $y_{NC}(x)$ to the annual escapement rule $y_A(x)$. The figure also displays the replacement curve (the 45-degree line), which is helpful because when the escapement rules are below the replacement curve, the stock level is effectively reduced in the relevant period. As seen in the highlighted part, the escapement rule for the low cost period is below the replacement curve for stock levels around 750 and 1,500 (thousand tonnes). Thus, if the initial stock is low, it may get trapped around these levels. An example of such trapping is shown in the bottom panel of Figure 2, where time paths are plotted against the decision rules. The replacement curve is also the identity map and is used to transfer between subsequent periods ($y_k = x_{k+1}$). The bottom panel also shows that in the seasonal model, the long run solution is a sub-annual, cycle with period 2. In comparison, the annual escapement rule is above the replacement curve for all positive stock levels up to the long run equilibrium at around 8,000. These features are further illustrated in Figure 3, where optimal time paths under both the seasonal model and the annual model are displayed. The figure shows that while the path under the seasonal model remains at a level around 1,000, the path under the annual model escapes to a high level near 8,000. Both paths start from an initial stock level of 500. If the initial stock level is higher (here, 1,500), also the path under the seasonal model ends up in a high state. The seasonal and the annual model leads to radically different dynamic system behavior, in other words. Also in Figure 3, we observe the sub-annual cycles of the seasonal model solutions.

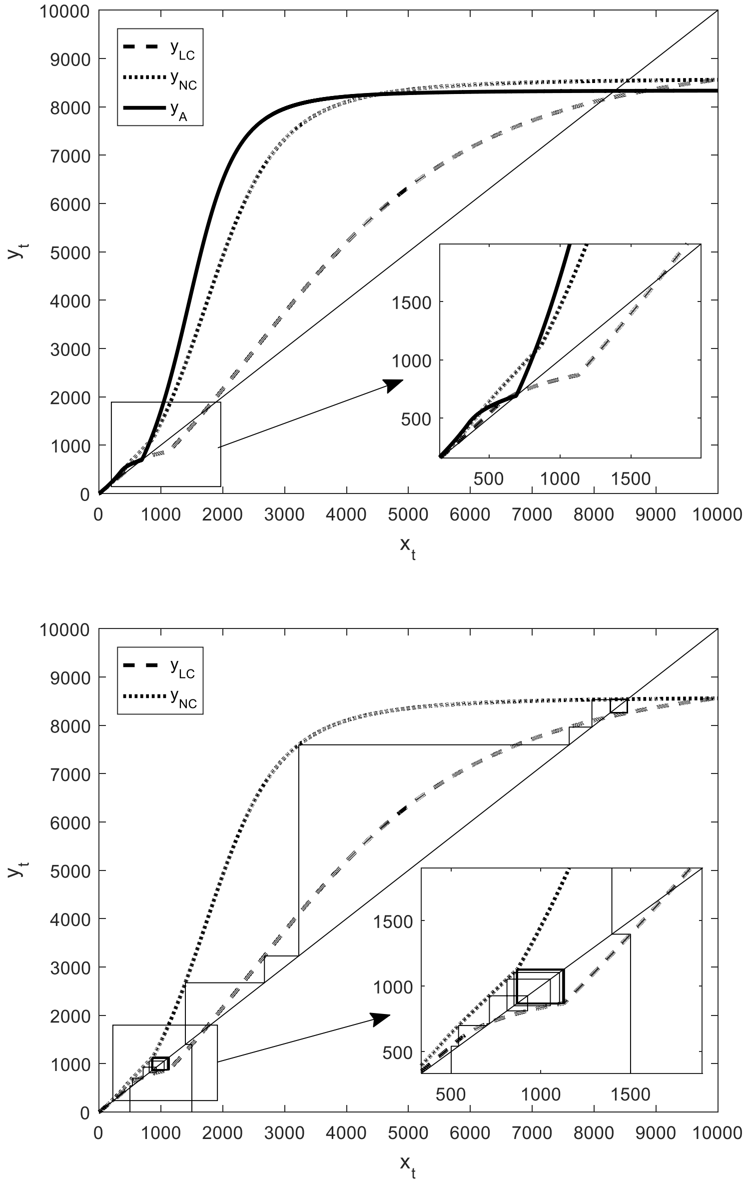

Figure 2: Top panel – Escapement rules for the seasonal (y_{LC}, y_{NC}) and annual (y_A) model. 45-degree line is the replacement curve. Bottom panel – Seasonal escapement rules (y_{LC}, y_{NC}) with dynamic paths for a relatively low initial stock level ($x_0 = 500$) and a higher initial stock level ($x_0 = 1{,}500$) (thin solid lines).

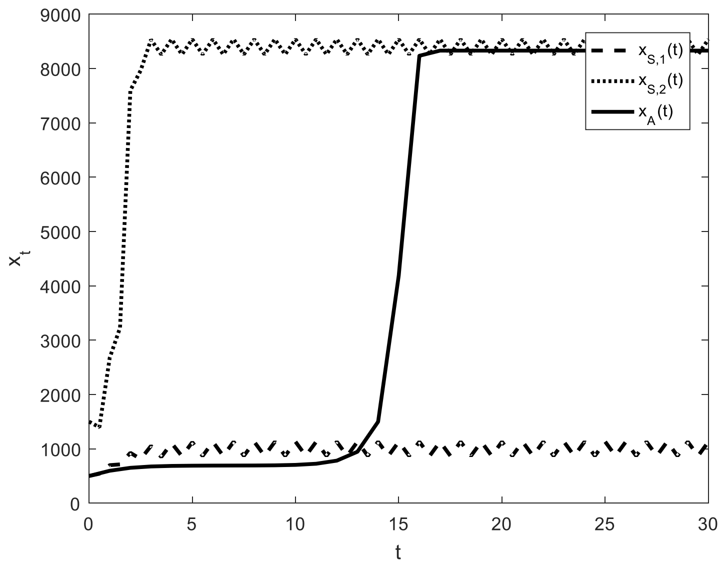

Figure 3: Optimal time paths for the seasonal ($x_{S,1}(t)$) and annual ($x_A(t)$) model for initial stock level $x_0 = 500$. Optimal time path for the seasonal model ($x_{S,2}(t)$) for initial stock level $x_0 = 1{,}500$.

Figures 2 and 3 does not tell the full story of our example, where the seasonal model is the hypothetical true model, while the annual model is a simplified approximation. This approximation is a potentially costly endeavor. To see why, we need to consider two moments. The first is the total, annual harvest in the two models. Figure 4 plots the total harvest as functions of stock level for the low cost and normal cost period in the seasonal model and the annual harvest in the annual model. For sufficiently high stock levels, the harvests are constant. (The constant harvest levels has to do with particular features of the return function; see Agnarsson *et al.* 2008 for discussion.) In the seasonal model, harvest is taken both in the low cost and the normal cost period, and at high stock levels this annual harvest outpaces the (annual) harvest in the annual model by more than 50 %. This feature, which exploits that a model with more periods has more freedom, we call actualization of the growth rates (AGR); see Kvamsdal et al. (2015) for a discussion. At lower stock levels (below 1,500), the case is more unclear as it depends critically, on the initial stock level and the time path, which of the models comes out on top in terms of annual harvest levels.

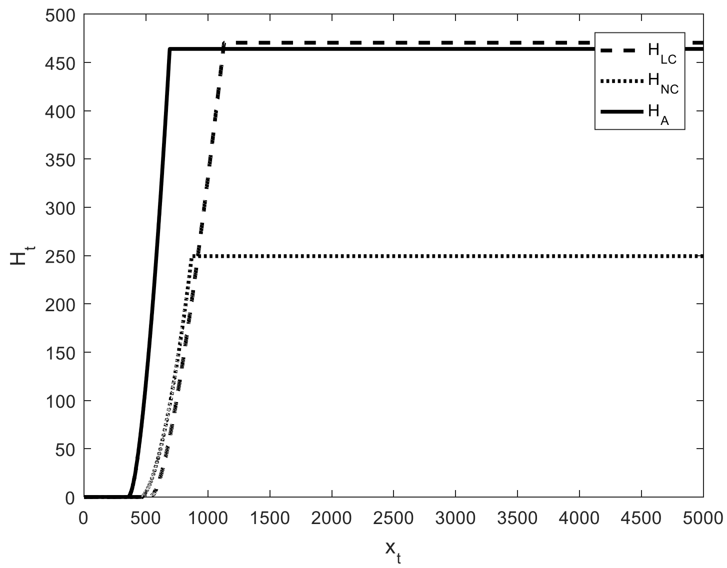

Figure 4: Total harvest (H_t) as functions of period initial stock level (x_t) for the low cost (H_{LC}) and normal cost (H_{NC}) periods of the seasonal model and for the annual model (H_A).

The second moment to consider is how a rational agent (here, a resource user; a fisher) living in the true model adapts to management based on the approximate annual model. In the interest of simplicity, we consider a representative agent who is given a total allowable catch of $h_A(x_t)$ for an initial stock level x_t. The rational agent who observes the different cost levels across the full period of which the catch quota is valid (the N-cycle), will maximize profits via heterogeneous distribution of harvest between periods. If we simplify by ignoring within-year discounting, the structure of the return function is such that the share distributed in the low cost period is independent of x_t and given by $\alpha = 0.85$ (see Agnarsson et al. 2008 for details of the return function). This adaptation of the rational agent subjected to an annual management scheme, but with freedom to exploit within-year differences in underlying conditions, has significant effects upon the resulting system dynamics. Figure 5 displays the effective escapement rules when $\alpha\ H_A$ is harvested in the low cost period and $(1-\alpha)H_A$ is harvested in the normal cost period. The dynamic system has an attractive fix-point around 500, considerably lower than the optimal attractive fix-point for the optimal, seasonal system (around 1,000; see bottom panel of Figure 2).

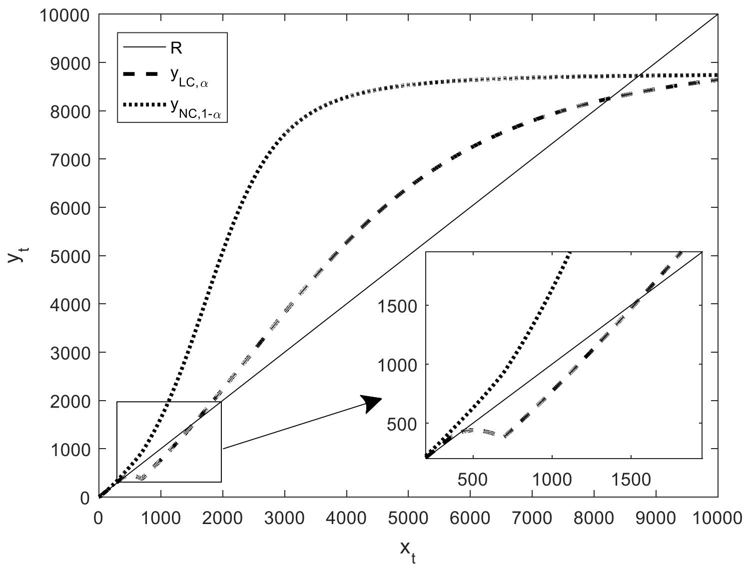

Figure 5: Effective escapement rules for the low cost ($y_{LC,\alpha}$) and normal cost ($y_{NC,1-\alpha}$) periods given that the low cost period harvest is αH_A and the normal cost period harvest is $(1-\alpha)H_A$.

The ability to solve periodic optimization problems is thus a potentially most valuable tool, obviously in fisheries, but likely in a wide range of settings. Our example here contains several valuable lessons about practical implications when periodicity is explicitly taken care of, or rather, the potential pitfalls of abstracting from periodicity. Figure 2 and 3 show that the periodic problem solution has a trap in the sense that for a low initial stock level, the stock level will remain at comparatively low stock levels. In contrast, the annual model, which has a similar initial dynamic behavior (figure 3), has no such trap. Thus, if low stock levels are biologically undesirable in the long run, the simplification inherent in the annual model may prove disastrous.

When we pursue the simplification of the annual model further by letting the resource users adapt rationally to the true, seasonal model while being subject to management based on the annual model, the problems with abstracting from periodicity are exacerbated. As shown in figure 5, management based on the annual model leads to a significantly suppressed stock level. In theory, the effect seen in figure 5 could become so severe that the stock could go extinct. Such inter-annual or within-season inefficiencies have gained some attention in the fisheries economics literature; see review in Smith (2012) and also Huang and Smith (2014).

4. Final remarks

We have shown that the periodic problem (3) implies the problem (10), and that the operator defined in (9) is a contraction which unique fix-point yields the value function of the problem. Further, the formulation in (3) suggest a numerical scheme that is more efficient than that suggested by (10). Our example shows that explicitly taking periodicity into account may have significant, practical consequences in the short and long run.

References

Agnarsson, S., R. Arnason, K Johannsdottir, L. Ravn-Jonsen, L.K. Sandal, S.I. Steinshamn, N. Vestergaard (2008). Multispecies and stochastic issues: Comparative evaluation of the fisheries policies in Denmark, Iceland, and Norway. *TemNord* 540, Nordic Council of Ministers, Copenhagen.

Bellman, R. (1957). *Dynamic Programming*, Princeton University Press, US.

Bertsekas, D.P. (2001). *Dynamic Programming and Optimal Control*, Athena Scientific, US.

Blackwell, D. (1965). Discounted Dynamic Programming, *Annals of Mathematical Statistics* 36: 226-35.

Huang, L., M.D. Smith (2014). The dynamic efficiency costs of common-pool resource exploitation. *The American Economic Review* 104(12), 4071-4103.

Kvamsdal, S.F., J.M. Maroto Fernández, M. Morán Cabré, L.K. Sandal (2015). *Continuous vs discrete-time bioeconomic models: Seasonal fisheries*, Working paper E-print 34774, Universidad Complutense Madrid.

Rincón-Zapatero, J.P, C. Rodriguez-Palmero (2003). Existence and uniqueness of solutions to the Bellman equation in the unbounded case. *Econometrica* 71(5), 1519-1555.

Smith, M.D. (2012). The New Fisheries Economics: Incentives Across Many Margins. *Annual Review of Resource Economics* 4, 379-429.

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

Acknowledgements

This research has been supported by a grant from Iceland, Liechtenstein and Norway through the EEA Financial Mechanism. Operated by Universidad Complutense de Madrid. Ref: 021-ABEL-CM-2013, and by the Research Council of Norway, project no. 257630. Part of the computations of this research was performed in EOLO, the HPC of Climate Change of the

International Campus of Excellence of Moncloa (Universidad Complutense de Madrid), funded by MECD and MICINN.