The twin-problem in numerical solutions of continuous-time Bellman equations

Klaus Wälde[1]
University of Mainz, Université catholique de Louvain and CESifo

August 2010

A maximized Bellman equation in continuous time is satisfied by two continuously differentiable value functions. This is illustrated for a deterministic setup for which an analytical solution exists. The second solution is derived numerically. We find that the second solution implies a convex value function, i.e. consumption would fall in wealth. Assuming (or proving) concavity of the value function therefore allows to identify the correct value function.

1 Introduction

We consider the maximization problem of an individual that solves a standard infinite-horizon optimal saving problem. There are various methods to numerically solve for the value function of a maximization problem of this type (see e.g. Judd, 1998). In continuous-time setups, a maximized Bellman equation has the structure of an implicit differential equation. The maximized Bellman equation is an implicit differential equation as the shadow price, i.e. the derivative of the value function, appears twice in the maximized Bellman equation and one can not solve explicitly for this derivative. One can use existing routines (e.g. ode15i in matlab) to solve for the value function. This is very convenient as the solution is very fast.

There is one shortcoming of implicit differential equations, however. Looking at the implication of this implicit structure shows that there are two *derivatives* of the value function for any *level* of the value function. Both of these derivatives satisfy the Bellman equation. This implies that there are two continuously-differentiable value functions which satisfy the Bellman equation.

Luckily, we also find that one of these value functions is convex, implying that consumption falls in wealth. As the latter can be ruled out on economic grounds (or can in some setups even be proven to contradict certain fundamentals), we are able to identify a unique solution.

The motivation for this analysis comes from the numerical solution of the matching and saving model of Bayer and Wälde (2010a,b). The numerical solution method - which was

[1]Gutenberg School of Management and Economics, University of Mainz, Jakob-Welder-Weg 4, 55131 Mainz, Germany, klaus@waelde.com, http://www.waelde.com, Phone + 49.6131.39-20143, Fax + 49.6131.39-23827. I would like to thank Andrey Launov, Jeremy Lise and Jakob Schwab for discussions. All matlab codes are available upon request.

not based on solving an implicit differential equation - also led to two solutions. In an attempt to understand the principles behind this “twin-problem”, the nature of the maximized Bellman equation implying an implicit differential equation with two solutions turned out to be at the heart of the problem. This paper illustrates this problem for the simpler standard deterministic consumption-saving problem. The conclusion for Bayer and Wälde (2010a,b) is that in the case of multiple solutions it must be checked wether one of these solutions implies convex value functions. If it does, a unique solution can be identified for these setups as well.

The next section presents a consumption-saving problem in continuous time with infinite horizon. Section 3 presents the analytical solution of the optimal consumption path and of the value function. Section 4 computes the numerical solution of the consumption path. It first shows why there are two solutions for the value function. It then discusses the identification of the unique consumption path. The final section concludes.

2 The model

We consider a classic deterministic continuous-time maximization problem. An individuals maximizes her utility function

$$U(t) = \int_t^\infty e^{-\rho[\tau - t]} u(c(\tau))\, d\tau, \tag{1}$$

where the felicity function is standard CRRA,

$$u(c(\tau)) = \frac{c(\tau)^{1-\sigma} - 1}{1-\sigma}, \quad \sigma > 0, \tag{2}$$

subject to the budget constraint

$$\dot{a}(t) = ra(t) + w - c(t). \tag{3}$$

The current point in time is t and future points in time are denoted by $\tau \geq t$.

3 The analytical solution to this model

3.1 Optimal consumption and wealth paths

The Keynes-Ramsey Rule has the usual structure

$$\frac{\dot{c}}{c} = (r - \rho)/\sigma \equiv g. \tag{4}$$

This, together with the BC (3) is an ODE system that can be solved, taking a No-Ponzi-Game-condition into account, to yield an explicit closed-form solution for consumption, expressing consumption $c(\tau)$ as a function of the only state variable $a(\tau)$,

$$c(\tau) = (r - g)\left[a(\tau) + \frac{w}{r}\right]. \tag{5}$$

Inserting this into the BC (3) yields the following corresponding optimal time path for wealth,

$$a(\tau) = e^{g[\tau - t]}\left[a(t) + \frac{w}{r}\right] - \frac{w}{r}.$$

where $a(t)$is initial wealth at time t.

When we are interested in consumption as a function of time τ only, we can derive this by inserting this into (5). We then get

$$c(\tau) = (r - g)\, e^{g[\tau - t]}\left[a(t) + \frac{w}{r}\right]$$

3.2 The value function and shadow prices

Inserting this into the instantaneous utility function and the latter then into the lifetime utility function yields the value function

$$V(a(\tau)) = \frac{v(\tau)}{\rho - (1-\sigma)g} - \frac{1}{(1-\sigma)\rho}$$

where

$$v(\tau) \equiv \frac{\left((r-g)\left[a(\tau) + \frac{w}{r}\right]\right)^{1-\sigma}}{1-\sigma}$$

The value function in complete form thus reads

$$V(a(\tau)) = \frac{\left((r-g)\left[a(\tau) + \frac{w}{r}\right]\right)^{1-\sigma}}{1-\sigma} * \frac{1}{\rho - (1-\sigma)g} - \frac{1}{(1-\sigma)\rho} \tag{6}$$

The shadow price, i.e. its derivative can be computed directly as

$$V'(a(\tau)) = \frac{dV(a(\tau))}{da(\tau)} = \frac{dv(\tau)/da(\tau)}{\rho - (1-\sigma)g} = \left((r-g)\left[a(\tau) + \frac{w}{r}\right]\right)^{-\sigma} \tag{7}$$

(which is the same as taking FOC's from the Bellman equations we will see and inserting the analytical consumption level).

4 The numerical solution to this model

4.1 Multiple solutions to the Bellman equation

Let us now assume we want to solve this model numerically. Let us further assume we would like to do this in a way which is widely used, i.e. by solving the Bellman equation taking the first-order condition into account. Doing so, we start from the Bellman equation

$$\rho V\left(a\left(t\right)\right)=\max_{c(t)}\left\{u\left(c\left(t\right)\right)+V_{a}\left(a\left(t\right)\right)\left[ra\left(t\right)+w-c\left(t\right)\right]\right\}. \tag{8}$$

The first-order condition reads

$$u'\left(c\left(a\right)\right)=V_{a}\left(a\right)\Leftrightarrow c\left(a\right)^{-\sigma}=V_{a}\left(a\right)\Leftrightarrow c\left(a\right)=V_{a}\left(a\right)^{-1/\sigma}, \tag{9}$$

where we suppress time arguments for simplicity. The second derivative is given by

$$\left(\frac{d}{dc\left(t\right)}\right)^{2}\{.\}=u''\left(c\left(a\right)\right)<0$$

making sure we found a maximum.

Combining the first-order condition (9) with the Bellman equation (8), we obtain the maximized Bellman equation

$$\begin{aligned}\rho V\left(a\right)&=\frac{c\left(a\right)^{1-\sigma}-1}{1-\sigma}+V_{a}\left(a\right)\left(ra+w-c\left(a\right)\right)\\&=\frac{V_{a}\left(a\right)^{-(1-\sigma)/\sigma}-1}{1-\sigma}+V_{a}\left(a\right)\left[ra+w-V_{a}\left(a\right)^{-1/\sigma}\right]\\&=\frac{V_{a}\left(a\right)^{-(1-\sigma)/\sigma}-1}{1-\sigma}+V_{a}\left(a\right)\left[ra+w\right]-V_{a}\left(a\right)^{-(1-\sigma)/\sigma}\Leftrightarrow\end{aligned}$$

$$\rho V\left(a\right)=\frac{\sigma V_{a}\left(a\right)^{-(1-\sigma)/\sigma}-1}{1-\sigma}+\left[ra+w\right]V_{a}\left(a\right). \tag{10}$$

This equation nicely shows why there is a risk of obtaining multiple solutions. For a given level $V\left(a\right)$ of the value function, there will generally be two shadow prices that solve this equation. As the shadow price is tightly linked to the consumption level via (9), this means that there are generally two consumption levels which satisfy the Bellman equation. When we express the maximized Bellman equation as

$$\Delta\equiv\frac{\sigma V_{a}\left(a\right)^{-(1-\sigma)/\sigma}-1}{1-\sigma}+\left[ra+w\right]V_{a}\left(a\right)-\rho V\left(a\right)$$

with $\Delta = 0$, then we see easily that there can be two shadow prices for which Δ is zero. Compute the derivative of Δ w.r.t. V_a and obtain

$$\begin{aligned}\frac{d\Delta}{dV_a} &= -\frac{1-\sigma}{\sigma}\frac{\sigma V_a\left(a\right)^{-1/\sigma}}{1-\sigma} + ra + w \\ &= -V_a\left(a\right)^{-1/\sigma} + ra + w\end{aligned}$$

This difference has a minimum $\Delta^{\min}$ defined by

$$V_a^{\min}\left(a\right) = \left(\frac{1}{ra+w}\right)^{\sigma}$$

as the second derivative of Δ w.r.t. the shadow price is positive.

We now plot $\Delta = \Delta\left(V_a\right)$, i.e. we understand Δ as a function of the shadow price. The parameter values we use are the ones used in TwoRootsForDeterministicBE_.m. We then obtain the following figure where the horizontal axis shows V_a and the vertical axis plots Δ.

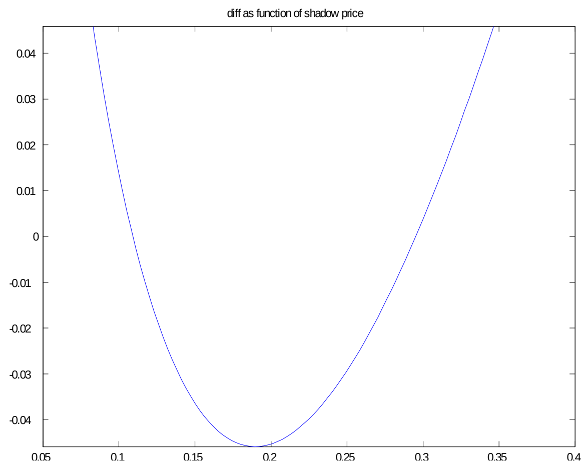

Figure 1 *The Bellman equation is satisfied by two shadow prices*

sigma=1.5, V(a)=179.384, VPrime(a)=0.109216
minimum of BE at -0.0458912
root1: 0.109216, root2: 0.295195, analytical root/shadow price: 0.109216

Figure 2 *Output from matlab*

One root (root1) is the analytical one known from (7), the other one (root2) is the larger one. This does not seem to be a numerical precision error. As the second root is larger, the consumption level implied by the first-order condition will be lower than the analytical consumption level.

4.2 Solving the Bellman equation numerically

Let us now solve the maximized Bellman equation (10) numerically. We look at this equation as an implicit differential equation for $V(a)$ where the exogenous variable is wealth a. It is an implicit differential equation as one can not solve for the differential $V_a(a)$ explicitly (apart from the special case of $\sigma = .5$). Let us assume we have some initial condition $V(a_0)$. For simplicity, we take the one from the analytical solution above in (6) for some given a_0, i.e.

$$V(a_0) = \frac{\left((r-g)\left[a_0 + \frac{w}{r}\right]\right)^{1-\sigma}}{1-\sigma} * \frac{1}{\rho - (1-\sigma)g} - \frac{1}{(1-\sigma)\rho}.$$

We then have, as the discussion in 4.1 has shown, two derivatives, i.e. two shadow prices root1 and root2 which satisfy the maximized Bellman equation.

As there are two derivatives, there are two solutions to this differential equation with initial condition $V(a_0)$, the one with the higher and the one with the lower derivative. Hence, one can therefore expect two continuously differentiable solutions, the one where always the lower derivative is chosen and the one where always the higher derivative is chosen. There is also a continuum of non-differentiable value functions (by arbitrarily jumping back and forth between the continuous solutions). The following figure shows the two continuous value functions in the left panel and the same value functions with a smaller range in the right panel.

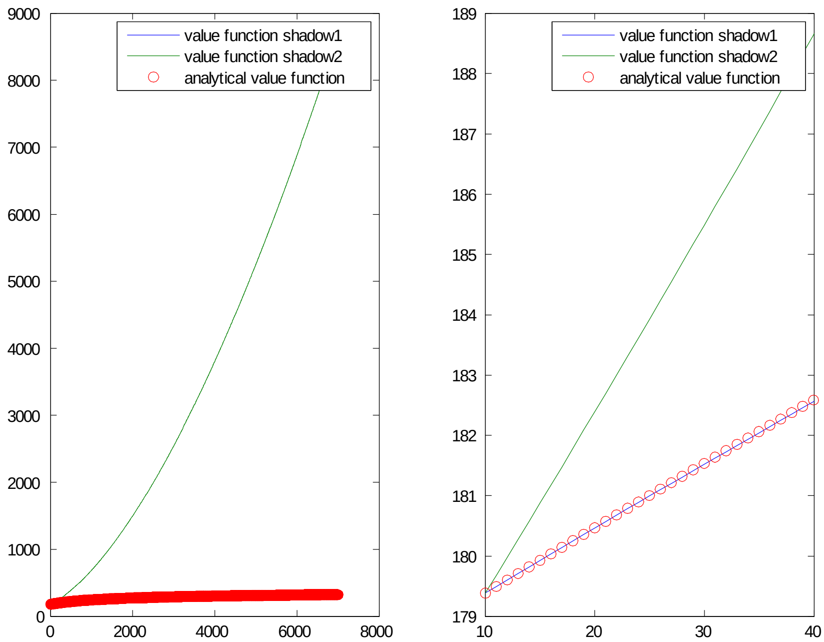

Figure 3 *Two value functions*

The "value function shadow 2" is the one with the higher root2. It is convex! The "value function shadow 1" is the one with the root 1 which corresponds to the analytical root. This one is concave. What is more, it coincides with the graph of the analytical value function. Hence, we can be sure that the numerical solution does what it should do.

4.3 Identifying the correct numerical solution

How can we now make sure which of these two solutions is the correct one? The answer comes from looking at the first-order condition again.

Compute the derivative of the first-order condition (9) with respect to a. This gives

$$u''(c(a))\,c'(a) = V_{aa}(a).$$

With a convex value function, i.e. with $V_{aa}(a) > 0$, this implies that $c'(a) < 0$, i.e. that consumption falls in wealth. This is an implausible property and any numerical solution implying a convex value function should therefore be ruled out. Hence, if we are willing to impose that consumption rises in wealth, we have identified the correct numerical solution. It is the one which implies a concave value function.

Note that it might be possible to prove on analytical grounds that the value function *must* be concave. In fact, Bayer and Wälde (2010b) prove that in the consumption-wealth space they are interested in consumption rises in wealth. This means that a numerical solution of their model which implies a convex value function would be a contradiction. Any valid numerical solution therefore must imply concave value functions.

5 Conclusion

We have shown that two continuously differentiable value functions exist which satisfy the maximized Bellman equation. One value function implies consumption and wealth dynamics which correspond to the analytical solution. The other one implies that consumption falls in wealth. This is a property which allows to rule out this second solution. One unique solution remains.

Concerning stochastic setups, this implies that the solution with a convex value function can also be ruled out. It also suggests why value function iteration always finds one unique solution: if value function iteration imposes concavity of the value function, only one solution is left.

Generally speaking, obtaining numerical value functions by solving maximized Bellman equations as implicit differential equations still seems to be a good idea as the solution tools are well-developed (there are routines e.g. in matlab). One would have to make sure, however, that both solutions are identified and that the relevant one is identified. The criterion suggested here is concavity of the value function.

References

Bayer, C., and K. Wälde (2010a): "Matching and Saving in Continuous Time : Theory," CESifo Working Paper 3026.

——— (2010b): "Matching and Saving in Continuous Time: Proofs," CESifo Working Paper 3026-A.

Judd, K. L. (1998): Numerical Methods in Economics. MIT Press.