School of Economics and Finance

A Generalized Endogenous Grid Method for Non-concave Problems

Giulio Fella

Working Paper No. 677 May 2011 ISSN 1473-0278

Queen Mary
University of London

A generalized endogenous grid method for non-concave problems*

Giulio Fella
Queen Mary University of London†

April 21, 2011

Abstract

This paper extends Carroll's (2006) endogenous grid method and its combination with value function iteration by Barillas and Fernández-Villaverde (2007) to non-concave problems. The method is illustrated using a consumer problem in which consumers choose both durable and non-durable consumption. The durable choice is discrete and subject to non-convex adjustment costs. The algorithm yields substantial gains in accuracy and computational time relative to value function iteration, the standard solution choice for non-concave problems.

JEL classification: C63

Keywords: Endogenous grid method, non-concavity

*The original motivation for the algorithm in this paper was the need to reduce computational time in Fella and Gallipoli (2009). I am grateful to Giovanni Gallipoli for his extensive feedback and to Jésus Fernández-Villaverde and Josep Pijoan-Mas for useful discussions. I am indebted to the Department of Economics at the University of Pennsylvania for its generous hospitality while working on this paper.

†School of Economics and Finance, Queen Mary University of London, Mile End Road, London E1 4NS, UK. E-mail: `g.fella@qmul.ac.uk`

1 Introduction

Many interesting dynamic economic problems are non-concave. This is the case, for example, when choice sets are non-convex either because choices are discrete or because they entail fixed costs. Discrete choices arise naturally in the literature on retirement (e.g. Rust 1989), labour supply (Gomes, Greenwood and Rebelo 2001), education (Gallipoli, Meghir and Violante 2009). Fixed adjustment cost are found in the literature on investment (Khan and Thomas 2008) and consumer-durables (Bajari, Chan, Krueger and Miller 2009).

In general, equilibria in these models need to be computed numerically. The non-concavity of the problem implies that the optimal policy correspondence may not be continuous, and the value function not differentiable, even with respect to continuous variables such as saving or the stock of capital. As a consequence, the Bellman maximand for the dynamic programming problem may not be differentiable even on the interior of the choice set. In the absence of differentiability, numerical optimization cannot exploit more efficient methods relying on first order conditions and has to resort to, notoriously slow, global comparison methods.

Discretized value function iteration is perhaps the most common approach to such problems. Yet, it severely suffers from the curse of dimensionality.

This paper develops a much more efficient and accurate algorithm to solve a class of problems that encompasses discrete-choice and fixed-adjustment cost problems. Problems in this class are differentiable in the endogenous continuous state variables at an internal maximum, though not necessarily everywhere in the interior of the choice sets. It follows that first order conditions are still necessary for an internal local maximum for the continous state variables.

The algorithm exploits this property of the class of problems considered to generalize the endogenous grid method (EGM hereafter) first proposed by Carroll (2006), and its extension to value function iteration (VFI hereafter) by Barillas and Fernández-Villaverde (2007), to non-concave, and possibly non-differentiable, problems. The idea behind Carroll's EGM is the following. Consider an optimal saving problem. In the standard approach, one fixes values for the endogenous state variable - wealth - at the beginning of the period and solves the Euler equation forward for the associated values of end-of-period wealth. EGM instead fixes values for end-of-period wealth and solves the Euler equation backward for the associated values of beginning-of-period wealth. The second approach is much faster as the Euler equation is often linear in beginning of period assets, but non-linear in end of period ones.

Since the Euler equation holds at an internal maximum, the algorithm uses EGM to locate an exact solution to the Euler equation - a local extremum. If the solution falls in the region where the problem is non-concave, it then uses standard VFI to verify whether the local extremum is a global maximum. The solution is very accurate because the algorithm determines the value of initial assets for which a given value of future assets solves the Euler equation *exactly.* The imprecise VFI global method is used only to confirm that the candidate global maximum is indeed so or to discard it. The algorithm is very efficient because thanks to EGM it eschews root finding. The algorithm can further be refined by exploiting the monotonicity of the saving/investment function.

The algorithm is illustrated for a consumer problem with discrete durable choice and fixed durable adjustment cost. The model is the same as in Bajari et al. (2009), but with a discrete rather than continuous durable choice. The assumption that the durable choice is discrete has two purposes. It facilitates the illustration of the algorithm and increases the computational challenge by increasing the range of assets over which the problem is non-concave.

Hintermaier and Koeniger (2010) also extend EGM to a consumer problem with durables and borrowing constraints, but their algorithm requires concavity. Clausen and Strub (2010) identify general restrictions on problem primitives under which the value function is differentiable at an internal optimum for the endogenous, continuous state variables and first order conditions hold. These restrictions define the class of problems to which the algorithm applies.[1]

The paper is structured as follows. Section 2 introduces the problem and the properties of the solution - local differentiability and monotonicity - that underpin the solution algorithm. Section 3 describes the basic algorithm while Section 4.2 discusses how to modify it to exploit monotonicity and deal with problems with occasionaly-binding borrowing constraints. Section 5 reports the numerical results, while Section 6 concludes.

2 The problem

2.1 The model

Consider a household with an infinite lifetime who, in each period t, chooses current non-durable consumption c_t, durable consumption[2] d_{t+1} and risk-free financial wealth w_{t+1}. At date 0, the household values alternative durable and non-durable consumption paths according to

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^s u(c_t, d_{t+1}), \tag{1}$$

where $\beta \in (0, 1)$ is a discount factor and the function u is strictly increasing and concave, twice differentiable and satisfies the Inada conditions. Initial financial wealth and durable holdings are (w_0, d_0).

In each period, the household earns a stochastic labour income y_t which follows an $m-$state Markov chain with transition matrix P, and state space $Y = \{y^1, \ldots, y^m\}$, with $y^i > y^{i-1}, i = 2, \ldots, m$. The household also earns capital income rw_t, where r is the risk-free rate of return.

There are two sources of non-convexity. First, the choice set for durables is discrete. More formally, the durable choice satisfies the constraint

$$d_{t+1} \in D \tag{2}$$

with D a discrete, compact subset of $\mathbb{R}_+$, with smallest element $\underline{d} = 0$ and cardinality larger than one. Secondly, the durable stock is subject to non-convex adjustment costs.

[1]An earlier draft of this paper pre-dating Clausen and Strub (2010) defined the class of problems to which the algorithm applies in terms of properties of the first derivative of the value function itself rather than in terms of primitives.

[2]I adopt the notational convention of indexing durable consumption at *time t* by $t+1$ to simplify the notation in the recursive problem.

Each unit of durable purchased involves a cost $(1+\phi)$ but the cost is zero if the stock of durables is not adjusted.

It follows that the household dynamic budget identity can be written as

$$c_t + w_{t+1} + (1 + \mathbb{I}_d\phi)d_{t+1} = y_t + (1+r)w_t + d_t \tag{3}$$

where $\mathbb{I}_d$ is an indicator function equal to zero if $d_{t+1} = d_t$ and one otherwise.

The non-durable consumption choice is bounded below by a non-negativity constraint[3]

$$c_t \geq 0 \tag{4}$$

and above by a borrowing constraint

$$w_{t+1} \geq -\gamma y_1 - \xi d_{t+1}, \tag{5}$$

where $\gamma \in [0, r^{-1}]$ and $\xi \in [0, (1+r)^{-1}]$ are respectively the fraction of minimum labour income and durable stock that can be collateralized and satisfy the following assumption.

The restrictions on the two parameters γ and ξ requires that the lowest feasible wealth level is never lower than the natural borrowing limit which obtains when both parameters are at their upper bounds.[4] The restriction implies that the household choice set is always non-empty.

The household maximizes (1) subject to the constraints (2)-(5).

It is useful to write the household problem in such a way that the borrowing constraint (18) does not depend on the choice variable d_{t+1}. To this effect let define the variable

$$a_t = w_t + \xi d_t. \tag{6}$$

with $a_t \in A$, where A is a Borel set in $\mathbb{R}$. The dynamic budget constraint (3) becomes

$$c_t + a_{t+1} + \lambda d_{t+1} = z_t(a_t, d_t, y_t) \tag{7}$$

where

$$z_t(a_t, d_t, y_t) = y_t + (1+r)a_t + [1 - (1+r)\xi]\lambda d_t \tag{8}$$

denotes total resources and $\lambda = (1 - \xi + \mathbb{I}_d\phi)$.

The transformed household problem can be written in the canonical form

$$\max_{\{a_{t+1}, d_{t+1}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^s u(z(a_t, y_t, d_t) - a_{t+1} - \lambda d_{t+1}, d_{t+1}), \tag{9}$$

$$\text{s.t. } (a_{t+1}, d_{t+1}) \in \Gamma(a_t, d_t, y_t),\ t = 0, 1, \ldots, \tag{10}$$

$$(a_0, d_0, y_0) \in A \times D \times Y \text{ given}, \tag{11}$$

with the feasibility correspondence $\Gamma : A \times D \times Y \to A \times D$ given by

$$\Gamma(a_t, d_t, y_t) = \{a_{t+1}, d_{t+1} : d_{t+1} \in D, a_{t+1} \in [-\gamma y_1, z(a_t, d_t, y_t) - \lambda d_{t+1}]\}.$$

[3]The Inada condition ensures that c_t is always strictly above its lower bound.

[4]To see this, note that when $\gamma = r^{-1}$ and $\xi = (1+r)^{-1}$ the borrowing limit for a household with current durable stock d_{t+1} equals the present value of human plus durable wealth. The household can repay all her debt along the worst possible income history by downsizing its durable stock to zero and consuming zero forever after.

2.2 The recursive problem

2.2.1 The recursive representation

Before introducing the recursive representation of household problem, it is useful to introduce the following, rather trivial, lemma whose main purpose is to characterise the class of problems to which the proposed solution algorithm applies.[5]

Lemma 1. *The Principle of Optimality holds for the household problem* (9)-(11)*. Furthermore,* $u(z(a_t, d_t, y_t) - a_{t+1} - \lambda d_{t+1}, d_{t+1})$ *is differentiable in* a_t *and* a_{t+1} *on the interior of* A*.*

The relevant point of the Lemma is the differentiability of the felicity functional with respect to present and future assets. Clausen and Strub (2010) show that the property is sufficient for the value function to be differentiable in the continuous state variable a_t at an optimum. The algorithm proposed in this paper applies to all problems that satisfy this property.

Since the problem is stationary, the time index can be dropped in what follows. Let $V(a, d, y)$ be the value function for the household sequence problem starting in state (a, d, y). This function is the unique solution to the Bellman equation

$$V(a, d, y) = \max_{(a', d') \in \Gamma(a,d,y)} u(z(a, d, y) - a' - \lambda d', d') + \beta \mathbb{E} V(a', d', y').$$

It is convenient to denote the expectation of the continuation value by

$$\tilde{V}(a', d', y) = \beta \mathbb{E} V(a', d, y') \tag{12}$$

in what follows and rewrite the Bellman equation as

$$V(a, d, y) = \max_{(a', d') \in \Gamma(a,d,y)} u(z(a, d, y) - a' - \lambda d', d') + \tilde{V}(a', d', y). \tag{13}$$

2.2.2 Solution

Before discussing the solution method, it is useful to introduce some non-standard notation and terminology.

Let

$$\Gamma(a, d, y|d') = \{a' : a' \in [-\gamma y_1, z(a, d, y) - \lambda d']\} \tag{14}$$

denote the set of feasible choices for a' for given d'.

Definition. *Given a triplet* $(d, y, d') \in D \times Y \times D$*, the saving correspondence conditional on* (d, y, d')*, or* conditional saving correspondence*, is the mapping* $a'(a, y, d|d') : A \to A$ *solving*

$$a'(a, d, y|d') = \arg\max_{a' \in \Gamma(a,d,y|d')} u(z(a, d, y) - a' - \lambda d', d') + \tilde{V}(a', d', y) \qquad s.t. \;\; d' \text{ given.} \tag{15}$$

[5] All proofs are in the Appendix.

If the function $\tilde{V}(a, d, y)$ were known, the solution to the recursive problem could be found using the following three steps for each $(d, y) \in D \times Y$.

STEP 1 For each $d' \in D$, solve equation (15) for the conditional saving correspondence $a'(a, \cdot|d')$.

STEP 2 For each $d' \in D$, use $a'(a, \cdot|d')$ to replace in the Bellman equation (13) and obtain the correspondence $V(a, \cdot|d') : A \to \mathbb{R}$ satisfying

$$V(a, \cdot|d') = u(z(a, \cdot) - a'(a, \cdot|d') - d', d') + \tilde{V}(a'(a, \cdot|d'), d', y). \tag{16}$$

STEP 3 Solve for the (unconditional) policy correspondences for saving $a'(a, d, y)$ and durables $d'(a, d, y)$ and the value function $V(a, d, y)$ satisfying

$$V(a, d, y) = \max_{d' \in D} V(a, d, y|d'), \tag{17}$$

$$d'(a, d, y) = \arg\max_{d' \in D} V(a, d, y|d'), \tag{18}$$

$$a'(a, d, y) = a'(a, d, y|d'(a, d, y)). \tag{19}$$

Equation (12) closes the system.

The system underpins the VFI solution method. Given an initial guess $\tilde{V}^0(a, y)$ for the function $\tilde{V}(a, y)$ one can iterate on the above system, together with equation (12), until convergence.

In standard, concave, problems one does not necessarily need to solve for the value function. One could use the first order conditions to iterate on the policy functions.[6] In non-concave problems, though, solving for the value function is essential as first order conditions are not sufficient for a global maximum.

2.2.3 Some analytic properties of the solution

This section derives some analytic properties of the solution on which the algorithm is based.

For the class of problems satisfying Lemma 1, an internal local maximum always satisfies the first order condition for assets, as stated in the following corollary.

Lemma 2. *The first order condition*

$$-u_c(z(a, y) - a' - d', d') + \partial\tilde{V}(a', d', y)/\partial a' = 0 \tag{20}$$

is necessary for an internal local maximum.

The lemma is an application of Theorem 2 in Clausen and Strub (2010) for the problem at hand. They show that the upper envelope of differentiable functions is differentiable at an internal optimum for a continuous variable even though it may not be differentiable everywhere.

The intuition behind the result is the following. If Lemma 1 holds, the only source of non-differentiability in a' of the continuation value $\tilde{V}(a', d', y)$ is changes in the future

[6] See, for example, Hintermaier and Koeniger (2010).

discrete choices as a' changes. Clausen and Strub (2010) show that the derivative of the value function at kinks can only jump up as a' increases. Since at an internal local maximum for a', the expression on the left hand side of equation (20) changes sign from positive to negative as a' increases, the maximum cannot be located at a kink. Therefore if a turning point is located at a non-differentiability it can only be a local minimum. It follows that the Euler equation (20) always holds at a internal local maximum and, therefore, at a candidate internal global maximum.

The above discussion implies that Lemma 2 still applies as long as the value function has only upward kinks. That is, it also applies if the differentiability condition in Lemma 1 is replaced by the weaker condition that $u(z(a_t, d_t, y_t) - a_{t+1} - \lambda d_{t+1}, d_{t+1})$ is differentiable almost-everywhere in a_t and a_{t+1} on the interior of A and has only upward kinks at the points were it is non-differentiable.

Finally, the following Proposition establishes that the saving correspondence is strictly increasing, and therefore invertible, in a on the interior of the choice set.

Proposition 1. *The conditional saving correspondence $a'(a, d, y|d')$ is increasing in a. It is strictly increasing if $a'(a, d, y|d') > -\gamma y_1$.*

Proposition 1 holds in both the class of problems considered here and in standard concave problems. The difference is that in non-concave problems the policy correspondence may not be a function. Usefully, though, the correspondence is strictly increasing in a off corners[7] which implies that its inverse with respect to a is a function.

3 The solution algorithm

This section generalizes Carroll's (2006) EGM algorithm, and its extension to VFI by Barillas and Fernández-Villaverde (2007), to the class on non-concave problems satisfying Lemma 1 in the above section.

The contribution of my extension, and of the original EGM in general, lies in solving for the conditional saving correspondence from the present- to the next-period's value of the continuous state variable a for given (d, y, d') in equation (15), namely Step 1 in Section 2.2.2. All other steps are the same as in the standard VFI.

As the conditional policy correspondence $a'(a, d, y|d')$ does not have a closed-form solution, an approximation to it has to be constructed by solving equation (15) on a finite grid for the continuous state variable a, for a given triplet (d, y, d'). Let (d, y, d') be fix in what follows and let $a'(a, \cdot|\cdot)$ denote the conditional saving correspondence.

For standard concave problems, Carroll's (2006) EGM dramatically speeds up the maximization step by exploiting the following three features of the problem.

1. The Euler equation (20) is necessary and sufficient for an internal global maximum.

2. The conditional saving correspondence $a'(a, \cdot|\cdot)$ is invertible for a on the interior of the domain.

3. The Euler equation (20) is much easier to solve for a given a predetermined a' than vice versa.

[7]That is the set of maximizers is increasing in a even though the set is not necessarily a singleton.

The first point trivially implies that the solution to equation (15) can be computed by solving for the unique zero of the Euler equation (20).

The second point implies that in solving for the conditional saving correspondence at a finite set of points one can, interchangeably, proceed in one of two ways.

The usual way is to construct an ordered grid $G_a = \{a_1, a_2, \ldots, a_m\}$ for *initial* assets a and solve the Euler equation (20) for its zero a'_i for every $a_i \in G_a$. The set of pairs $\{a_i, a'_i\}_{i=1}^m$ is the conditional saving correspondence on the set of collocation points G_a.

Alternatively, one could construct an ordered grid $G_{a'} = \{a'_1, a'_2, \ldots, a'_m\}$ for *end-of-period* assets and solve for the value of *total resources* z_i^{end} that satisfies the Euler equation (20) for each $a'_i \in G_{a'}$. The set of pairs $\{z_i^{end}, a'_i\}_{i=1}^m$ is the conditional saving correspondence on the *endogenous* set of collocation points for total resources $G_z^{end} = \{z_i^{end}, z_2^{end}, \ldots, z_m^{end}\}$. Given (d, y), one can use equation (8) to recover the corresponding collocation points for initial wealth a_i^{end} satisfying $z(a_i^{end}, d, y) = z_i^{end}$ thus obtaining the conditional saving correspondence $\{a_i^{end}, a'_i\}_{i=1}^m$ on the set of collocation points $G_a^{end} = \{a_1^{end}, a_2^{end}, \ldots, a_m^{end}\}$ for initial wealth a.

The difference between the two procedures is that in the first one the set of collocation points is pre-determined and the value of the conditional saving correspondence at those points is endogenous. Vice versa, in the second procedure the value of the conditional saving correspondence at the collocation points is predetermined, while the collocation points themselves are endogenously generated. Hence, the name.

The disadvantage of the first procedure is that the Euler equation is non-linear in a'. Solving for a' involves evaluating the Euler equation multiple times for each collocation point a_i. Vice versa, the computational cost of solving the Euler equation for a given a' is very low, as stated in Point 3. above. This can be easily seen in the case in which the felicity function is separable in c and d'; e.g. it satisfies $u(c, d') = \theta \log(c) + (1 - \theta)\log(g(d'))$. In such a case, the Euler equation can be written as $z - a' - \lambda d' = \theta \tilde{V}(a', y)^{-1}$, which is linear in z.[8]

Going from an endogenous collocation point z_i^{end} for total resources to one for initial wealth a_i^{end} involves solving the linear equation (8) in the present problem. In general, though, the relationship between initial wealth and total resources is non-linear[9] Yet, one can recover the saving correspondence at any arbitrary grid point for initial wealth – e.g. a_i – by using the set of pairs $\{z_i^{end}, a'_{end}\}$ to construct an interpolating function. Evaluating such function at the point $z(a_i, y, d)$ returns the value of the conditional saving correspondence at the chosen point a_i.

So, EGM trades off the cost of constructing an interpolating function against the cost of solving a non-linear equation, a very advantageous trade-off.

For general non-concave problems, point 1. does not apply, though, and the Euler equation (20) is neither necessary nor sufficient for a global maximum. One is forced to resort to, notoriously slow, global methods.

For the class of problems satisfying Lemma 1, though, it follows from Corollary 2 and Proposition 1 that EGM is still useful to locate an internal local maximum. Since, given

[8]Even if the $u_c(c, d')$ did not have a closed-form inverse with respect to c, given a' some variant of Newton method converges to the unique solution for z at a quadratic rate. A similar method cannot be applied to solve for a' if, as in the present model, $\partial \tilde{V}(a', d, y)/\partial a'$ is not differentiable in a'.

[9]This is the case, for example, in the neoclassical growth model studied in Carroll (2006) and Barillas and Fernández-Villaverde (2007).

the non-concavity of the problem, a local maximum is not necessary a global one, the algorithm modifies the standard EGM in the following way. First, it partitions the set of grid points for future assets $G_{a'}$ into a *non-concave region* $G_{a'}^{nc}$ in which the Euler equation is not sufficient for a global maximum for a' and its set complement. Secondly, for all a'_i in the non-concave region, the algorithm supplements EGM with a global maximization step. Since the non-concave region is a subset of the choice set $G_{a'}$, the first step restricts the, costly, application of the global maximization step to the non-concave region rather than the whole of $G_{a'}$.

The two steps are illustrated in the next three subsections. To simplify the exposition, we assume the following.

Assumption 1. *The parametric restrictions* $\gamma = r^{-1}$ *and* $\xi = (1+r)^{-1}$ *hold.*

The assumption implies that the borrowing constraint $a' \geq -\gamma y_1$ is the natural borrowing constraint and therefore never holds. We relax this in Section 4.2.

It also follows from Proposition 1 that one can exploit the monotonicity of the conditional saving function $a'(a, d, y|d')$ to accelerate the computation of the solution. It turns out that monotonicity is even more powerful when the policy function is discontinuous. A refined version of the algorithm exploiting monotonicity is described in Section 4.1.

3.1 Identifying the non-concave region

The advantage of identifying the non-concave region in advance is that, outside it, one can use the unmodified EGM algorithm as the Euler equation is both necessary *and* sufficient for an internal global maximum. Since in many problems the non-concave region is a, possibly small, subset of the asset grid, this reduces the set of points at which one has to use a, substantially slower, global method.

Understanding how the algorithm delimits the non-concave region is easier with the help of Figure 1 which draws the marginal utility of present consumption and of future assets as functions of a'. The thick non-monotonic and discontinuous curve plots the marginal utility of future assets $\partial\tilde{V}(a', d', y)/\partial a'$ for given (d', y). The thinner upward sloping curve is the marginal utility of present consumption for a given value of total resources z and durable choice d'. A point where the two curves intersect is a zero of the Euler equation.

In terms of Figure 1, for each abscissa $a'_i \in G_{a'}$ the EGM finds the value of total resources z_i for which an upward sloping curve intersects the thick $\partial\tilde{V}(a', d', y)/\partial a'$ at $a' = a'_i$; namely a'_i is a zero of the Euler equation. The Euler equation is sufficient for a'_i to be a global maximum if a'_i is the unique intersection between the upward sloping curve $u_c(z_i - a' - \lambda d', d')$ through it and the curve $\partial\tilde{V}(a', d', y)/\partial a'$. A sufficient condition for the intersection to be unique is that for all $a'_j \in G_{a'}$ it is $\partial\tilde{V}(a'_j, d', y)/\partial a' > \partial\tilde{V}(a'_i, d', y)/\partial a'$ for all $j < i$ and $\partial\tilde{V}(a'_j, d', y)/\partial a' < \partial\tilde{V}(a'_i, d', y)/\partial a'$ for all $j > i$. In Figure 1, this is the case in the regions where $\partial\tilde{V}(a', d, y)/\partial a'$ is above v_{max} and below v_{min}, or equivalently for any value of assets outside the set $G_{a'}^{nc} = \{a'_2, \ldots, a'_9\}$.[10]

[10] The fact that $\partial\tilde{V}(a', d', y)/\partial a' > v_{max}$ for a' low enough follows from Assumption 1, maintained in this section, that implies that the borrowing constraint is always slack. This not true in general. Section 4.2 discusses how the algorithm needs to be modified when it is not.

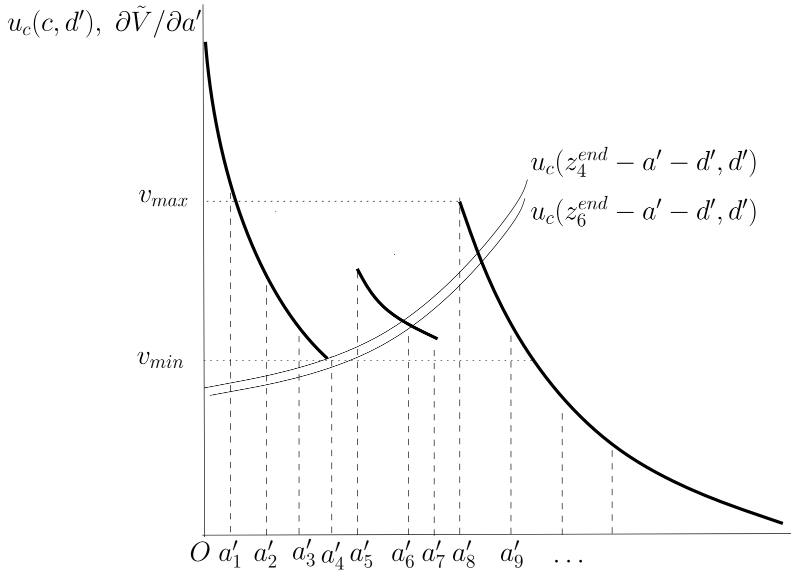

Figure 1: Illustrating the algorithm

Assuming the function $\tilde{V}(a', d', y)$ is known, the bounds v_{min} and v_{max} can be computed, for each given (d', y), as respectively the lowest value of $\tilde{V}(a'_i, d', y)$ and the highest value of $\tilde{V}(a'_{i+1}, d', y)$ for all i such that $\tilde{V}(a'_{i+1}, d', y) > \tilde{V}(a'_i, d', y)$. Given v_{min} and v_{max}, one can compute $\underline{i}$ – the largest i such that $\tilde{V}(a'_i, d', y) > v_{max}$ – and $\bar{i}$ – the smallest i such that $\tilde{V}(a'_i, d', y) < v_{min}$.

By construction, the Euler equation is necessary and sufficient for a maximum for $a' \leq a'_{\underline{i}}$ and $a' \geq a'_{\bar{i}}$. The Euler equation is only necessary though for $a'_i \in G^{nc}_{a'} = \{a_{\underline{i}+1}, \ldots, a_{\bar{i}-1}\}$.

3.2 The basic algorithm

Given (d', y) and the associated non-concave region $G^{nc}_{a'}$ identified in the previous section, the algorithm proceeds in the following way. First for each $a'_i \in G_{a'}$ it applies the standard EGM algorithm and uses equation (20) to solve for z^{end}_i. If a'_i lies outside the non-concave region - e.g. $a'_i = a'_1$ in Figure 1 - the algorithm stores the pair (z^{end}_i, a'_i) and moves to the next point in $G_{a'}$. If instead, as is the case for a'_4 in Figure 1, a'_i belongs to the non-concave region, the algorithm has to locate the *global* maximum associated with z^{end}_i. To do so, for given z^{end}_i the algorithm constructs the discretized Bellman maximand for all a'_i in the non-concave region $G^{nc}_{a'}$ and finds the maximum of the discretized problem

$$a'_g = \arg\max_{a' \in G^{nc}_{a'}} u(z^{end}_i - a' - \lambda d', d') + \tilde{V}_n(a', d', y). \tag{21}$$

If $a'_g = a'_i$, a'_i is both a local and global maximum given z^{end}_i and, again, the pair (z^{end}_i, a'_i) is stored. If instead a'_g is different from a'_i the algorithm does not store any point and just moves onto the next grid point a'_{i+1}.

It should be clear from the above description that the application of the global maximization step to the discretized problem is used only to verify whether a local extremum is a global maximum. It is not used to actually solve for a point on the conditional saving function. If the solution a'_g differs from the original point a'_i, the algorithm does not replace a'_i with the global maximum a'_g for the discretized problem; it just discards a'_i. If $a'_g > a'_i$ the same procedure will be repeated when a'_g is reached and a'_g will be stored only if it is a fixed-point of the procedure. The solution is very accurate because the algorithm determines the value of total resources for which a given grid point a'_i for future assets solves the Euler equation *exactly.* The imprecise global method is used only to confirm that the candidate local extremum is indeed a global maximum or to discard it.

Before presenting the pseudo-code for the algorithm it is useful to tie a few loose ends. First, one has to select an ordered grids $G_{a'}$ for next-period's assets a'. Second, it is useful to store in memory the value of total resource implied by the grid for next period's assets $G_{a'}$ as $G_{z'}(d, y) = z(a, d, y)$ for all $(a, d, y) \in G_{a'} \times D \times Y$. Third, since the function $\tilde{V}(a', d', y)$ is unknown it has to be found by repeated iteration of the system (15)–(17) starting from some initial guess $\tilde{V}^0(a', d', y)$. The initial choice of guess $\tilde{V}^0(a', d', y)$ has to be continuous and increasing and satisfy Lemma 1. It is advisable to choose a differentiable function, to obtain its wealth derivative $\partial\tilde{V}^0(a', d', y)/\partial a'$ by finite differences.

At all subsequent iterations $n > 0$, one solves for $\tilde{V}_n(a', d', y)$ by the usual iterative procedure discussed in Section 2.2.2. The wealth derivative $\partial\tilde{V}_n(a', d', y)/\partial a'$ at the points of the grid $G_{a'}$, can be approximated either by taking finite differences of $\tilde{V}_n(a', d', y)$ or using the envelope condition[11]

$$\frac{\partial\tilde{V}_{n+1}(a', d', y)}{\partial a'} = (1 + r)\mathbb{E}u_c(c_n(a', d', y'), d'_n(a', d', y')) \tag{22}$$

where c_n is given by equation (7) with $a' = a'_n(a, d, y)$.

The corresponding pseudo code is the following.

1. Set $n = 0$. Guess a function $\tilde{V}^0(a', d', y)$ and compute its wealth derivative $\partial\tilde{V}^0(a', d', y)/\partial a'$.

2. For all (d', y) solve for the bounds $\underline{i}$, $\bar{i}$ of the non-concave region $G^{nc}_{a'}$ as derived in Section 3.1. Set $i = 1$, $l = 1$ and do the following.

   2.1. Compute the level of total resources z^{end}_i that solves equation (20) evaluated at point a'_i.

   2.2. If $\underline{i} < i < \bar{i}$, find the solution a'_g to the discretized global problem using (21) with $z = z^{end}_i$. If $a'_g \neq a'_i$, $i = i + 1$ and go to step 2.1.

   2.3. Set $i_l = i$ and store the pair $(z^{end}_{i_l}, a'_{i_l}) = (z^{end}_i, a'_i)$.

   2.4. Replace in equation (16) to obtain the value of the conditional value correspondence
   $$v^{end}_{i_l} = u(z^{end}_{il} - a'_{il} - d', d') + \tilde{V}_n(a'_{il}, d', y) \tag{23}$$
   associated with the level of total resources z^{end}_{il}. If $i < m$, $i = i + 1$, $l = l + 1$ and go to Step 2.1.

[11] In our numerical experiments using the envelope condition results in a slight improvement in accuracy relative to finite differences.

2.5. Since the new guess for the value function has to be defined on the grid $G_{a'}$ interpolate the pairs $\{z^{end}_{i_l}, a'_{i_l}\}$ and $\{z^{end}_{i_l}, v^{end}_{i_l}\}$ on the grid $G_{z'}(d, y)$ for all $d \in D$ to obtain $a'_n(a, d, y|d')$ and $V_{n+1}(a, d, y|d')$ for all $a \in G_{a'}$ and $d \in D$.

3. For all (a, d, y), compute the unconditional policy and values functions $d'_n(a, d, y)$, $a'_n(a, d, y)$ and $V_{n+1}(a, d, y)$ using (18)–(17).

4. Solve equation (12) to obtain $\tilde{V}_{n+1}(a, d, y)$.

5. If $||\tilde{V}_{n+1}(a, d, y) - \tilde{V}_n(a, d, y)||_\infty > 10^{-5}$, with $||\cdot||_\infty$ the sup norm over $G_{a'} \times D \times Y$, use the envelope condition (22) to obtain $\partial\tilde{V}_{n+1}(a_i, d, y)/\partial a'$ and start a new iteration.

4 Generalizations and refinements

4.1 Monotonicity

For points in the non-concave region, the algorithm supplements the standard EGM step with the same global maximization step as in discretized VFI. Proposition 1 implies that for given (d, y, d') the conditional saving correspondence $a'(a, d, y|d')$ is monotonically increasing in initial wealth a. Since $z(a, d, y)$, is strictly increasing in a, as implied by equation (20), the optimal conditional saving choice is increasing in z for given (y, d'). It is easier to frame the discussion that follows in terms of monotonicity between a' and the intermediate variable z.

As in concave problems, monotonicity can be usefully exploited to economize on the number of comparisons at the global maximization step. To understand the implication of monotonicity in the present context, consider again Figure 1. The thick broken line plots the expected marginal utility $\partial\tilde{V}(a', d', y)/\partial a'$ as a function of a' for given (d', y). The upward sloping line plots the marginal utility of current consumption as a function of a' for given d' and total resources z. Keeping d' constant, higher values of z shift the latter curve down. Since, as discussed in Section 3.2 the global maximization step applies only in the non-concave region - namely $G^{nc}_{a'} = \{a'_2, \ldots, a'_9\}$ in the figure - it is only in such region that one needs to exploit monotonicity.

Suppose that one has already evaluated all grid points a'_i up to a'_3 and all of them are (global) maxima given the associated z^{end}_i. The next step is to verify whether a'_4, is a global maximum z^{end}_4 that solves equation (20) for $a' = a'_4$; namely for the value of z for which the upward-sloping curve intersects the thick one at $a' = a'_4$. Given that $\tilde{V}(a'_4, d, y) < \tilde{V}(a'_3, d, y)$, it is $z^{end}_4 > z^{end}_3$. It follows from monotonicity that the maximum associated with z^{end}_4 cannot lie to the left of a'_3. Therefore, in solving the discretized problem in equation (21), one needs to compare values of the maximand on the right hand side of the equation only at grid points in the set $\{a'_3, \cdots, a'_9\}$ rather that at all the points in $G^{nc}_{a'}$. This is the standard way monotonicity is used to speed up the solution of concave problems.

The combination of the monotonicity of the policy correspondence *and* the lack of concavity of the value function (the non-monotonicity of its wealth derivative) can be further exploited, though.

To see this suppose that a'_4 is indeed a global maximum for z^{end}_4 satisfying equation (20). The next step is to verify whether point a'_5 is a global maximum. At point a'_5 it

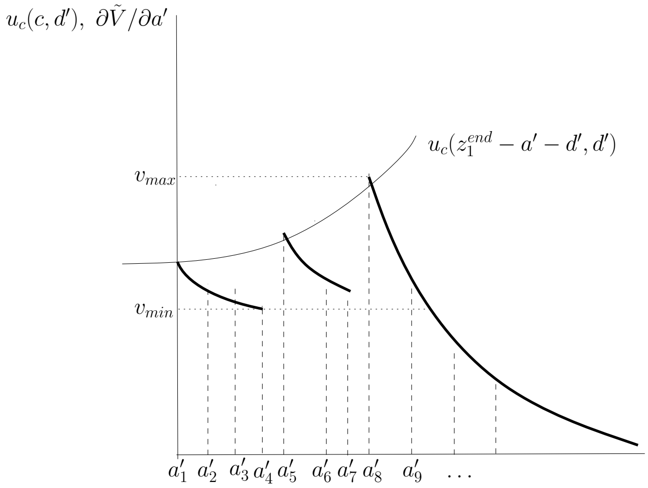

Figure 2: Borrowing constraint

is $\partial\tilde{V}(a'_5,d',y)/\partial a' \geq u_c(z_4^{end} - a'_5 - d', d')$, which implies $z_5^{end} \leq z_4^{end}$. It follows from monotonicity that the global maximum associated with z_5^{end} cannot lie to the right of a'_4. This is true for any point a'_i, such as a'_5 or a'_8 for which $\partial\tilde{V}(a'_i,d',y)/\partial a' \geq u_c(z_4^{end} - a'_i - d', d')$. Evaluating this inequality is all one needs to rule out any such point.

Consider now point a'_6. Suppose that applying the global maximization step (21) for $z = z_6^{end}$ returns a'_j, with $j > 6$, as the associated optimum; e.g $j = 9$.[12] It follows from monotonicity that the optimal saving choice associated with any $z > z_6^{end}$ cannot be lower than the level of a' for which the curve $\partial\tilde{V}(a',d',y)/\partial a'$ and $u_c(z_6^{end} - a' - d', d')$ intersect. This rules out points along the thick curve to the south-west of ditto intersection; namely points a'_i, like a'_7, such that $\partial\tilde{V}(a'_i,d',y)/\partial a' < u_c(z_6^{end} - a'_i - d', d')$ *and* $i < j$.[13]

4.2 Borrowing constraints

Up to now I have maintained, only for expositional reasons, the assumption that the borrowing constraint is never binding. Carroll (2006) though shows that EGM can accommodate occasionally binding borrowing extremely effectively. In what follows I relax the assumption that the borrowing constraint is never binding. Without loss of generality the following assumption normalizes the lower bound on next period's wealth to zero.

Assumption 2. *It is* $\gamma = 0$.

The first grid point in the grid $G_{a'}$ is now $a'_1 = 0$. Figure 2, effectively the counterpart of Figure 1, illustrates how EGM deals with the borrowing constraint.

For given (y, d') EGM calculates the value of total resources z_1^{end} for which the Euler equation is satisfied as an equality. The are two possible cases two distinguish.

In the first case, $a'_1 = 0$ is both a local and global maximum given $z = z_1^{end}$, namely $(z_{i1}^{end}, a'_{i1}) = (z_1^{end}, 0)$ in point 2.3 in the pseudo-code in Section 3.2. Therefore z_1^{end} is the

[12]Since a'_4 has been assumed to be a global maximum for $z = z_4^{end}$, this means that for some $z > z_4^{end}$ the optimal saving choice jumps up discontinuously from some point in the interval $[a'_4, a'_5]$ to some point to the right of the rightmost intersection between $\partial\tilde{V}(a',d',y)/\partial a'$ and $u_c(z_4^{end} - a' - d', d')$.

[13]One does need to consider, though, points like a'_9 for which $\partial\tilde{V}(a'_i,d',y) < u_c(z_6^{end} - a'_i - d', d')$ *but* $i \geq j$.

first interpolating node for the conditional saving and value functions. Let v_{i1}^{end} denote the associated value of the conditional value correspondences from equation (16). Since the Euler equation holds as an equality at (z_{i1}^{end}, a'_{i1}), the borrowing constraint is on the verge of being binding at z_{i1}^{end} and, from monotonicity, the constraint is strictly binding for any $z < z_1^{end}$.

The value of the conditional saving correspondence for all points (a, d, y) for which the borrowing constraint is binding - namely those for which $z(a, d, y) < z_1^{end}$ is just $a'(a, d, y|d') = 0$.

Replacing in (14), the associated value of the conditional value correspondence can be recovered as

$$V(a, d, y|d') = u(z(a, d, y) - d', d') + \tilde{V}(0, d', y).^{14} \tag{24}$$

This case is the only one which applies for concave problems. As first pointed out by Carroll (2006), EGM is extremely efficient in dealing with borrowing constraints in this case.

Consider instead the case in which given z_1^{end} satisfying the Euler equation for $a' = a'_1 = 0$ is not a global maximum given (y, d'). Instead, the solution to equation (21) for $z = z_1^{end}$ is some $a'_g > 0$. Therefore, the borrowing constraint is not necessarily just binding for $z = z_1^{end}$. In fact, the EGM steps 2.1-2.4 in Section 3.2 would return an interpolating function $\{z_{il}^{end}, a'_{il}\}$ whose first value $a'_{i1} > 0$. Therefore, one cannot conclude that the household chooses to be borrowing constraint for z in a left neighborhood of z_{i1}^{end}. The EGM algorithm no longer necessarily determines the lower bound on total resources below which the household is borrowing constraint.

Yet, because the saving function is monotonic such a lower bound exists. An approximation to it can be recovered by finding the value of z_{i0}^{end} that solves the following equation

$$u(z_{i0}^{end} - d', d') + \tilde{V}(0, d', y) = u(z_{i0}^{end} - a'_{i1} - d', d') + \tilde{V}(a'_{i1}, d', y). \tag{25}$$

The solution z_{i0}^{end} is the value of total resources for which the global optimum switches from $a'_1 = 0$ to a'_{i1}. Adding the point $(z_{i0}^{end}, 0)$ as the first point to the vector of interpolating nodes for the unconditional saving correspondence, and the associated value $v_{i0}^{end} = u(z_{i0}^{end} - d', d') + \tilde{V}(0, d', y)$ for the conditional value correspondence, allows to use the same interpolation procedure as in the first case considered.

5 Results

5.1 Parameterization

The parameterization follows Bajari et al. (2009) along a number of dimensions. The chosen felicity function is

$$u(c, d') = \frac{1}{\tau} \log(\theta c^\tau + (1 - \theta)\kappa(0.01 + d')^\tau), \tag{26}$$

[14] Alternatively $V(a, d, y|d') = u(z(a, d, y) - d', d') - u(z(a_1^{end}, y) - d', d') + v_1^{end}$.

Parameter	β	r	κ	ϕ	ξ	σ_y	ρ	σ_η
Value	0.93	0.06	0.075	0.06	0.20	0.063	0.977	0.024

Table 1: Chosen parameters

with the only modification that the marginal utility of durables is bounded. As in Bajari et al. (2009), the durable flow equivalent is $\kappa = 0.075$, the non-durable share $\theta = 0.77$ and the fractions of human and durables wealth that can be collateralized are respectively $\gamma = 0$ and $\xi = 0.2$. The intermediation fee is set to $\phi = 0.06$.

The income process is a discrete approximation to a lognormal process with a persistent and transitory components as in Storesletten, Telmer and Yaron's (2000)[15]. Namely,

$$\begin{aligned} \log y_t &= z_t + \epsilon_t \\ z_t &= \rho z_{t-1} + \eta_t, \end{aligned}$$

with ϵ_t, η_t distributed independently according to $N(0, \sigma_\epsilon)$, $N(0, \sigma_\eta)$.

The Markov chain approximation to the process follows Tauchen (1986). The number of grid points for both the transitory and persistent components is 7 which implies that y can take 49 discrete states.

I choose seven uniformly-spaced points for the durable choice stock and a double exponential grid for assets a. The upper bounds on a and d equal approximately 25 and 10 times unconditional average income. These values are large enough to ensure: (1) that the upper bound of the stationary distribution for a is below the highest grid point, and (2) that the upper bound on d does not constraint the durable choice.

Finally, the interest rate is set to $r = 0.06$, roughly in line with average real mortgage rates, and the discount rate is set to $\beta = 0.93$ to ensure boundedness of the wealth distribution. The chosen values for parameters are collected in Table 1.

The parameter τ governing the elasticity of substitution takes different values in our simulations. In most of our simulations it equals zero, which implies the Cobb-Douglas specification

$$u(c, d') = \theta \log(c) + (1 - \theta) \log(\kappa(0.01 + d'),$$

used in Fernández-Villaverde and Krueger (201). Under this specification the Euler equation is linear in total resources as discussed in Section 3.

In the last part of Section 5.2 I set $\tau = 0.2435$ as estimated in Bajari et al. (2009) to assess how the speed of the algorithm is affected by the non-linearity of the Euler equation in total resources.

5.2 Numerical results

Discretized VFI is the standard method of choice for non-concave and/or non-differentiable problems. It is, therefore, natural to compare the accuracy and speed of my algorithm to those of VFI. Since my algorithm exploits the monotonicity of the

[15]The estimates are from row D. in their Table 1. The permanent, individual-specific random effect is not included as it would play no role in the present set up.

policy function I do the same when solving the model using VFI, so as not to bias the comparison between the two methods.

To compare the accuracy of the two algorithms I compute Euler equation errors following Judd (1992). If $s = (a, d, y)$ denotes the state vector, the Euler equation

$$u_c[c(s), d'(s)] = \beta(1 + r)\mathbb{E}u_c[c(a'(s), d'(s), y'), d''(a'(s), d'(s), y')] \tag{27}$$

should hold exactly for the true policy functions off corners. Given that the computed policy functions are only approximations, equation (27) does not hold exactly when evaluated with the computed policy functions.

Let $c^*(s)$ denote the solution to

$$u_c[c^*(s), \hat{d}'(s)] = \beta(1 + r)\mathbb{E}u_c[\hat{c}(\hat{a}'(s), \hat{d}'(s), y'), \hat{d}'(\hat{a}'(s), \hat{d}'(s), y')], \tag{28}$$

where careted variables denote the approximate policy functions. The (absolute) Euler equation error measured in units of current consumption can then be written as

$$E(s) = \left|1 - \frac{c^*(s)}{\hat{c}(s)}\right| \tag{29}$$

for any point of the state space s.

An Euler error $E(s)$ equal to one per cent means that the agent is making a mistake of one cent for each dollar spent. Following Judd, I report the base 10 logarithm of the Euler error. Therefore, a one per cent error in (29) corresponds to an Euler error of -2.

As standard in the literature, I report both the largest Euler error for any point in the grid $||E(s^i)||_\infty$, and the largest and average Euler error along a simulated path which I denote respectively by $||E(s_t)||_\infty$ and $\bar{E}(s_t)$. To construct the latter two measures, I draw a 50,000-period income history. This together with the policy functions generates a history $\{s_t\}_{t=1}^{50{,}000}$ for the whole state vector. $||E(s_t)||_\infty$ and $\bar{E}(s_t)$ are the largest and average Euler error along such history. Since the Euler equation does not have to be satisfied at the borrowing constraint, I report the Euler errors only at those points in the state space where the borrowing constraint is slack.

The chosen initial conditions are $a_0 = d_0 = 0$ and the unconditional average of the income process. All the computations were carried out on a single core of a Xeon X5570 processor. The programs were written in `Fortran 95`. The code is available at http://webspace.qmul.ac.uk/gfella/research/research.html for download.

To compare my results with comparable papers, such as Barillas and Fernández-Villaverde (2007) and Hintermaier and Koeniger (2010), on concave problems, I first simulate a concave version of the model without a durable choice. For this, I keep d at its lowest value of 0. I solve the model using both EGM and VFI and grids of 400 and 1000 points for the state variable a. Table 2 reports the results.

A remark is in point before discussing the results. The fact that the supremum of the Euler errors on the grid $||E(s^i)||$ is not zero for EGM might appear puzzling. By construction, the Euler equation should be satisfied on the endogenous grid points. As the pseudo-code in Section 3 makes clear the algorithm uses the endogenous grid points and the associated values of the policy and value functions only as interpolating nodes to solve for the those function on the *exogenous* grid for assets $G_{a'}$. While it would straightforward to let the grid for future assets change endogenously at each iteration,

Model	# Grid points	Time(s)	$\|\|E(s^i)\|\|_\infty$	$\|\|E(s_t)\|\|_\infty$	$\bar{E}(s_t)$
VFI	400	9	-0.97	-2.27	-2.91
EGM	400	2	-6.05	-3.88	-6.27
VFI	1000	35	-1.47	-2.67	-3.37
EGM	1000	5	-6.85	-4.39	-7.16

Table 2: Results for the case with no durable choice and $\tau = 0$

there are two reasons why I am not doing it here. First, the endogenous grid for a would depend on the values of the other state variables d and y. This would require storing 7×49 grids here for assets. Second, exactly because $||E(s^i)||$ would be zero by construction it would no longer be an informative statistics to evaluate the relative performance of the two algorithms.

Table 2 makes clear that, in the absence of non-concavities, the accuracy of EGM is dramatically better than that of value function iteration. Furthermore, the computation time is substantially faster. With a wealth grid of 1000 points, EGM solves the problem in one seventh of the time compared to VFI. In terms of accuracy, even with less than half the number of grid points, EGM dramatically outperforms VFI.

I now conduct the same exercise, but allowing the household to choose among seven possible levels for the durable stock. This introduces two sources of non-concavity: the discreteness of the durable choice and the non-convexity of the adjustment cost function. I use grids of 200, 400 and 1000 points for the state variable a. Table 3 reports the results.

Consider first the case in which $\tau = 0$, which directly compares to that reported in Table 2. In terms of computational time, the relative speed of the EGM is similar to the non-concave case. This is very encouraging and not completely expected. The endogenous-grid algorithm has to use a global comparison method over the subset of the state space where the problem is non-concave. Over such a subset the method has no computational advantage compared to VFI. It turns out that, when the policy correspondence is discontinuous, applying monotonicity to the Euler equation, as discussed in Section 4.1, allows to discard a larger number of candidate points than in the standard application of monotonicity in VFI which just restricts the grid over which to search for a solution to equation (21).

Turning to accuracy the results may seem more mixed than in the concave case. Comparing the largest Euler errors either on the set of grid points or along the simulation path does not show a clear superiority of either method. Yet, comparing the average Euler error on the simulated path yields the same picture as in the non-concave case. EGM is roughly twice as accurate as VFI. Furthermore, according to the same criterion EGM significantly outperforms VFI even when the latter methods employs five times the number of grid points – 1000 versus 200 – with a computational time roughly 70 times that of EGM - 1192 versus 17 seconds.

The reason for why EGM is not more accurate according to the first two metrics is apparent once one realizes that the true consumption and saving correspondences are discontinuous and that they are approximated by interpolating using the endogenous grid points as interpolating nodes. As long as the true policy correspondences jump between two interpolating nodes, the Euler equation evaluated at their approximations may be significantly violated at any point in between. This is true independently from

Model	# Grid points	Time(s)	$\|\|E(s^i)\|\|_\infty$	$\|\|E(s_t)\|\|_\infty$	$\bar{E}(s_t)$
$\tau = 0$					
VFI	200	79	-0.60	-1.99	-2.62
EGM	200	17	-1.07	-1.66	-5.43
VFI	400	208	-0.91	-2.15	-2.87
EGM	400	39	-1.22	-1.98	-6.09
VFI	1000	1192	-1.34	-2.40	-3.26
EGM	1000	157	-1.09	-2.35	-6.94
$\tau = 0.2435$					
VFI	1000	2870	-1.33	-2.56	-3.28
EGM	1000	294	-1.38	-4.29	-6.98

Table 3: Result for the case with durable choice

the algorithm used. Therefore, the first two statistics are not particularly meaningful in the presence of discontinuities in the policy functions.

Finally, the last two lines in Table 3 conduct the same analysis for the case in which $\tau = 0.2435$. In such a case, the Euler equation (20) is non-linear in total resources z, as the marginal utility of consumption function is non-invertible in consumption. While non-linear, the Euler equation is twice differentiable with respect to total resources z - but not with respect to a' - and can be solved for z using Newton method.

The results reported in Table 3 make clear that, while the change of utility function nearly doubles computational time for both methods, it leaves their relative performance, both in terms of accuracy and computational time, virtually unaffected. If anything, the advantage of EGM in terms of computational time increases.

6 Conclusion

This paper has presented an extension of Carroll's (2006) EGM, and its combination with VFI by Barillas and Fernández-Villaverde (2007), to non-concave, and possibly non-differentiable problems. The proposed algorithm yields dramatic gains in accuracy and computational time.

I have illustrated the algorithm in the context of a problem with a continuous non-durable and a discrete durable choice and fixed adjustment costs, but one can adapt the techniques in Barillas and Fernández-Villaverde (2007) and Hintermaier and Koeniger (2010) to deal with a continuous Ss durable choice.

The algorithm applies without any modification to policy functions for continuous state variables as long as the objective function is differentiable, and dynamic constraints are differentiable in the continuous variables to which the algorithm is to be applied. An example of a case outside such class is a consumer problem in which the interest rate is a non-differentiable, with a downward-kink, function of wealth as in Bajari et al. (2009). Yet, as long as the location of downward kinks is known, and their number limited, one can still apply the algorithm in this paper piecewise.

References

Bajari, Patrick, Chan, Phoebe, Krueger, Dirk and Miller, Daniel P. (2009), A dynamic structural model of housing demand: Estimation and policy implications. Mimeo.

Barillas, Francisco and Fernández-Villaverde, Jesús (2007), 'A generalization of the endogenous grid method', *Journal of Economic Dynamics and Control* **31**(8), 2698–2712.

Carroll, Christopher D. (2006), 'The method of endogenous gridpoints for solving dynamic stochastic optimization problems', *Economics Letters* **91**(3), 312–320.

Clausen, Andrew and Strub, Carlo (2010), Envelope theorems for non-smooth and nonconcave optimization. Mimeo, University of Pennsylvania.

Edlin, Aaron S. and Shannon, Chris (1998), 'Strict monotonicity in comparative statics', *Journal of Economic Theory* **81**, 201–219.

Fella, Giulio and Gallipoli, Giovanni (2009), Education and crime over the life cycle. Mimeo, Queen Mary, University of London.

Fernández-Villaverde, Jesús and Krueger, Dirk (201), 'Consumption and saving over the life cycle: How important are consumer durables?', *Macroeconomic Dynamics* . forthcoming.

Gallipoli, Giovanni, Meghir, Costas and Violante, Gianluca (2009), Human capital accumulation, education policy and wage dispersion. Mimeo, University College London.

Gomes, Joao, Greenwood, Jeremy and Rebelo, Sergio (2001), 'Equilibrium unemployment', *Journal of Monetary Economics* **48**(1), 109–152.

Hintermaier, Thomas and Koeniger, Winfried (2010), 'The method of endogenous gridpoints with occasionally binding constraints among endogenous variables', *Journal of Economic Dynamics and Control* **34**(10), 2074 – 2088.

Judd, Kenneth L. (1992), 'Projection methods for solving the aggregate growth model', *Journal of Economic Theory* **58**, 410–452.

Khan, Aubhik and Thomas, Julia K. (2008), 'Idiosyncratic shocks and the role of nonconvexities in plant and aggregate investment dynamics', *Econometrica* **76**(2), 395–436.

Rust, John P. (1989), A dynamic programming model of retirement behavior, *in* David A.Wise, ed., 'The Economics of Aging', University of Chicago Press, pp. 359–398.

Stokey, Nancy L., Lucas, Robert E. and Prescott, Edward C. (1989), *Recursive Methods in Economic Dynamics*, Harvard University Press.

Storesletten, Kjetil, Telmer, Chris I. and Yaron, Amir (2000), Consumption and risk sharing over the life cycle. NBER Working Paper 7995.

Tauchen, George (1986), 'Finite state Markov-chain approximations to univariate and vector autoregressions', *Economics Letters* .

A Proofs

Proof of Lemma 1. That the Principle of Optimality holds follows from the assumptions that: (1) the endogenous state vector (a, d) lies in the Borel set $A \times D \subseteq \mathbb{R}^2$, (2) Y is countable, (3) Γ is a non-empty, compact valued and continuous correspondence, (4) $u(\cdot)$ is bounded and continuous and $\beta \in (0, 1)$. The assumptions coincide with assumptions 9.4 to 9.7 in Stokey, Lucas and Prescott (1989).

The differentiability of $u(z(a_t, d_t, y_t) - a_{t+1} - \lambda_{t+1}, d_{t+1})$ follows trivially from (8) and the assumptions on $u(\cdot)$. $\square$

Proof of Lemma 2. The proof follows closely that of Theorem 2 in Clausen and Strub (2010). Let $y^t = (y_0, \ldots, y_t)$ denote a partial history of income shocks from period 0 to t and let $p_t(y^t)$ denote the unconditional probability of history y^t. Given the transition probabilities $p(y'|y)$ associated with the Markov chain, it is $p_t(y^t) = p(y_t|y_{t-1})p_{t-1}(y^{t-1})$ with $p_0(y^0) = 1$.

Let $s_0 = (a_0, d_0, y_0)$ denote the state at $t = 0$ and

$$W(s_0; a_1, d_1) = u(z(s_0) - a_1 - \lambda d_1, d_1) + \tilde{V}(a_1, d_1, y_0) \tag{30}$$

the maximum of lifetime, given the current state s_0 and choices (a_1, d_1).

Finally, let $\Pi(s_1)$ denote the set of feasible plans from $t = 1$ onwards for given s_1, that is

$$\Pi(s_1) = \left\{ \{a_{t+2}, d_{t+2}\}_{t=0}^{\infty} : (a_{t+2}, d_{t+2}) \in \Gamma(s_{t+1}), \text{ for all } t \geq 0; s_1 \text{ given} \right\}, \tag{31}$$

and $f(s_0; a_1, d_1, \pi)$ denote the value

$$f(s_0; a_1, d_1, \pi) = \sum_{t=0}^{\infty} \beta^t \sum_{y^t} p_t(y^t) u\big(z(a_t, y_t, d_t) - a_{t+1} - \lambda d_{t+1}, d_{t+1}\big) \tag{32}$$

of the maximand in (9) for given $s_0; a_1, d_1)$, and an arbitrary $\pi \in \Pi(s_1)$.

Lemma 1 implies that $f(s_0; a_1, d_1, \pi)$ is differentiable in a_1 and that the function $W(s_0; a_1, d_1)$ is the upper envelope of $f(s_0; a_1, d_1, \pi)$ with respect to π. It follows from Theorem 1 in Clausen and Strub (2010) that $W(s_0; a_1, d_1)$ and, by (30), $\tilde{V}(a_1, d_1, y_0)$ is differentiable at an internal maximum for a_1. Since (30) coincides with the right hand side of (16), it follows that the first order condition (20) holds at an internal maximum for a'. $\square$

Proof of Proposition 1. Theorem 1 in Edlin and Shannon (1998) implies that an interior maximiser $x^*(t) \in \arg\max_x g(x, t)$ of a function $g(x, t)$ is strictly increasing in t if $\partial g/\partial x$ is increasing in t at $x^*(t)$. It follows from Lemma 2 that, for given (d, y, d') the objective function on the right hand side of (13) is differentiable, and its partial derivative with respect to a' satisfies (20), at an internal optimum for a'. Since the right hand side of (20) is strictly increasing in a, Theorem 1 in Edlin and Shannon (1998). $\square$

Queen Mary
University of London

This working paper has been produced by the School of Economics and Finance at Queen Mary, University of London

School of Economics and Finance
Queen Mary, University of London
Mile End Road
London E1 4NS
Tel: +44 (0)20 7882 5096
Fax: +44 (0)20 8983 3580
Web: www.econ.qmul.ac.uk/papers/wp.htm