



Comparing numerical methods for solving the competitive storage model

Christophe Gouel

► **To cite this version:**

Christophe Gouel. Comparing numerical methods for solving the competitive storage model. Computational Economics, Springer Verlag, 2013, 41 (2), pp.267-295. .

HAL Id: hal-01136976

<https://hal.archives-ouvertes.fr/hal-01136976>

Submitted on 28 Jul 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Comparing Numerical Methods for Solving the Competitive Storage Model*

Christophe Gouel[†]

Abstract

This paper compares numerical methods for solving the competitive storage model. Because storage implies a nonnegativity constraint on stocks, the solution methods must be considered carefully. The model is solved using value function iteration and several projection approaches, including parameterised expectations and decision rules approximation. In considering a storage model with convenience yield, in which the inequality constraint is smoothed, perturbation methods are also applied. Parameterised expectations approximation proves to be the most accurate method, whereas perturbation techniques are shown inadequate for solving this highly nonlinear model. The endogenous grid method allows rapid solution if supply is assumed to be inelastic.

Keywords: binding constraint, nonlinear rational expectations models, numerical methods.

JEL classification: C63, D84, E37.

*The author is grateful to Jean-Marc Bourgeon, Benjamin Carton, Sébastien Jean, Michel Juillard, and an anonymous reviewer for helpful comments and suggestions. The research was supported by the AgFoodTrade project funded under the 7th Framework Programme for Research and Development, DG-Research, European Commission.

This is a pre-copy-editing, author-produced PDF of an article accepted for publication in *Computational Economics* following peer review. The final publication (Gouel C., 2013, 41(2), 267–295, *Computational Economics*) is available at Springer via <http://dx.doi.org/10.1007/s10614-012-9318-y>.

[†]The World Bank – Development Research Group, Agriculture and Rural Development, 1818 H street NW, Washington, DC 20433, USA, Tel.: +1 202 473 7044, email: cgouel@worldbank.org, and CEPII

1 Introduction

The competitive storage model is the workhorse of neoclassical studies on the price behaviour of storable commodities (Williams and Wright, 1991, Wright, 2001). Although its empirical properties continue to be debated (Cafiero and Wright, 2006), this model provides justification for the main properties of commodity prices: positive serial correlation, skewness, successions of long periods of doldrums and short periods of high prices. It also serves as a normative benchmark for analysing public intervention in commodity markets (Miranda and Helmlinger, 1988). Like most dynamic stochastic problems, it cannot be solved analytically. For this reason, understanding the properties of available numerical solutions is important for securing precise econometric estimates or reliable policy conclusions. One example of the importance of good numerical solutions is provided by Cafiero et al. (2011), who show that Deaton and Laroque's (1995, 1996) econometric estimates are not reliable because of the imprecise methods used to approximate the model. Another aspect that requires good knowledge of numerical solutions is its extension to higher dimensional problems. Many important questions on commodity prices (e.g., the comovement of commodity prices, the relationship between trade and prices, the consequences of price stabilisation programmes, and the effect of monetary policies on prices) imply problems with several state and decision variables that become very challenging to solve numerically.

This paper compares different approaches to solving the competitive storage model. It implements three methods not tested previously on this model: perturbation, endogenous grid methods (EGM) and decision rules approximation. As perturbation methods only work with smooth problems, two alternative models are considered: one with a nonnegativity constraint on storage and the other with a precautionary motive for storage that rules out zero stock situations and smoothes the model. Comparing solution methods for this smooth model is the object of Miranda's (1997) study. Our approach differs from Miranda's in considering the model with a nonnegativity constraint—the more traditional representation of storage problems—and in its different focus. Miranda (1997) studies mainly the effect of varying approximation methods for a given projection algorithm, which parameterises the price function. The present paper tests various algorithms but limits variations between approaches by holding the set of tools constant: functions are approximated by cubic splines; convergence to the rational expectations equilibrium is achieved by function iteration; value function iteration (VFI) and projection problems are resolved by collocation; and the same degree of approximation is used for each method. The precision of the solutions is assessed using a measure of the Euler equation error that is derived to account for the switching-regime behaviour.

There is a large body of work on the comparison of numerical methods in economics. Most concerns the stochastic growth model (see, e.g., Aruoba et al. 2006 and Heer and Maußner 2008 for two recent studies). This literature covers only some of the methods used for the storage model. The specificities of the storage model come from the nonnegativity constraint on stocks. For these kinds of constraints, Christiano and Fisher (2000) study numerical methods in a stochastic growth model with irreversible investment. This paper uses some of the same methods as Christiano and Fisher but proposes also additional ones.

The study of numerical methods for the storage model presents interests beyond the study of commodity behaviour. The storage model, in effect, is formally very close to the consumption/saving problem under

income uncertainty (Deaton, 1991, Carroll, 2001), which is central for heterogeneous-agent modelling in macroeconomics (Krusell and Smith, 2006). Their problems are identical: how much an agent should consume and save today when future resources are volatile. In the optimal consumption problem, a consumer's borrowing capacity is limited, or he applies a limit as a precaution. At each period, the consumer chooses how much to save to protect his future consumption against adverse income shocks. In the storage model, there are three agents: a consumer, a storer acting competitively, and a producer using a stochastic production function. Without distortion, this problem can be stated as a planner problem. The planner must choose how much grain to carry over to the next period. There is one main difference between the two models: storage is costly and grain may get spoiled during storage, whereas in a consumption/saving problem, savings are remunerated. But the impatience of consumers makes saving costly, which renders the two model almost identical. This similarity was first noted by Deaton, who worked on both fields (Deaton, 1991, Deaton and Laroque, 1992). The longstanding ignorance about this resemblance had some consequences: a numerical method for the storage model was proposed by Gustafson (1958), whereas the consumption/saving problem under income uncertainty was solved using the misleading certainty equivalence or log-linearisation until the work by Zeldes (Barsky et al., 1986, Zeldes, 1989).

This paper provides several notable results. First, the parameterised expectations approach of Wright and Williams (1982) stands out as the most precise method. It has well-known good properties for solving models with occasionally binding constraints (Christiano and Fisher, 2000), but in our case, its good performances extend also to the smooth version of the model. Indeed, unlike other methods that approximate highly nonlinear functions, this method approximates a function that is almost linear. Second, perturbation methods are shown to be inadequate for this model. Given the nonlinearity of storage behaviour, low-order perturbations are imprecise and lead to negative storage, whereas high-order perturbations present diverging behaviour. Third, when supply is assumed to be inelastic, as is the case in most econometric studies, the endogenous grid method is by far the fastest, which makes it especially valuable for computationally intensive work.

2 The competitive storage model

2.1 Model equations

The storage model analysed in Wright and Williams (1982) is used throughout the paper. It features a market for a storable commodity with a competitive storer, a producer whose output is subject to multiplicative shocks and a final demand.

The activity of the competitive risk-neutral storer is to transfer the commodity from one period to the next. Storing the quantity S_t from period t to period $t + 1$ entails a physical cost $\Phi(S_t)$, a purchase cost $P_t S_t$, with P_t the market price, and an opportunity cost. A share δ of the commodity deteriorates during storage. The benefits valued in period t are $P_{t+1} S_t \cdot (1 - \delta) / (1 + r)$, with r as the interest rate. The profit expected by the storer is

$$E_t \left(\Pi_{t+1}^S \right) = \left[\frac{1 - \delta}{1 + r} E_t (P_{t+1}) - P_t \right] S_t - \Phi(S_t), \quad (1)$$

where E_t is the expectation operator conditional on period t information. Taking into account the possibility of a corner solution (i.e., the nonnegativity constraint of storage), expected profit maximisation yields the following complementary condition:¹

$$S_t \geq 0 \quad \perp \quad \frac{1 - \delta}{1 + r} E_t (P_{t+1}) - P_t - \Phi' (S_t) \leq 0, \quad (2)$$

which means that inventories are null when the marginal cost of storage, including purchase cost, is not covered by the expected marginal benefits; for positive inventories, the arbitrage equation holds with equality.

A representative producer makes his productive choice one period before bringing the output to market. He plans in period t a production level H_t for period $t + 1$, but a disturbance affects final production (e.g., weather disturbances). The expected profit can be written as

$$E_t \left(\Pi_{t+1}^H \right) = \frac{E_t (P_{t+1} H_t \epsilon_{t+1})}{1 + r} - \Psi (H_t), \quad (3)$$

where $\Psi (H_t)$ is the cost of planning the production H_t , and $H_t \epsilon_{t+1}$ is the realised production. ϵ_{t+1} is the realisation of a stochastic process, supposed i.i.d. and distributed following a normal centred on 1 with standard error σ . The planned production derived from expected profit maximisation satisfies

$$E_t (P_{t+1} \epsilon_{t+1}) = (1 + r) \Psi' (H_t). \quad (4)$$

Final demand is given by the inverse demand function $P (D)$, where D is the quantity consumed. The shocks being i.i.d., the state of the model is defined by total availability

$$A_t = (1 - \delta) S_{t-1} + H_{t-1} \epsilon_t. \quad (5)$$

Accounting for carry over storage, market equilibrium can be written as

$$A_t = D_t + S_t. \quad (6)$$

Using Equation (6), the current price can be substituted from the storage arbitrage condition (2) by a function of availability and storage:

$$S_t \geq 0 \quad \perp \quad \frac{1 - \delta}{1 + r} E_t (P_{t+1}) - P (A_t - S_t) - \Phi' (S_t) \leq 0. \quad (7)$$

Thanks to this substitution, the storage model is defined by two equilibrium equations, the first-order conditions (4) and (7), by the transition equation (5), and by the rational expectations hypothesis. A restriction of this model to inelastic supply is the standard tool for econometric works. In this case, the equilibrium reduces to Equation (7), with H_t taken as constant. Although this simplified model is not the object of this work, we will discuss some of its numerical specificities.

¹Complementarity conditions in what follows are written using the ‘‘perp’’ notation (\perp). This means that both inequalities must hold, and at least one must hold with equality.

2.2 Model parameterisation

Physical storage costs are usually assumed to be proportional to the amount stored (Wright and Williams, 1982). In our first parameterisation, we take $\Phi'(S) = 0.01$ (i.e., 1% of the deterministic steady-state price assumed to be 1). One alternative is to include in the storage cost a convenience yield, which would tend to take high negative values for low stocks, preventing the occurrence of stockouts. Miranda (1997) uses a logarithmic function to achieve this. We follow his approach by taking $\Phi'(S) = 0.3 + 0.1 \log(S)$, which conveniently removes the complementarity condition from Equation (7). As the logarithm tends to minus infinity for low stocks, stocks always remain positive. The two models show similar, but distinctive, behaviours (see Figure 1), the model with convenience yield being smoother than the model with the nonnegativity constraint. Negative marginal storage costs can be justified on the grounds that they account for both physical storage and convenience yield. To reconcile the fact that low stocks exist for apparently negative returns to storage, some authors, such as Kaldor (1939) and Working (1948), conclude that storers can expect a convenience yield from holding stock despite an apparent negative yield. This convenience yield comes from the possibility of being able to use stock at any moment. With a convenience yield, storage can be seen as responding to different motives: a speculative motive that allows stockouts and a precautionary motive that precludes them (Carter and Revoredo Giha, 2007).

The inverse demand function is taken to be isoelastic: $P(D) = D^{1/e}$. The production cost function also follows an isoelastic form: $\Psi(H) = H^{\alpha+1}/[(1+r)(\alpha+1)]$. The production cost is normalised by $1+r$ to ensure that the deterministic steady-state production is at 1. Table 1 presents the parameters used to calibrate the model. They are assumed to be such that the deterministic steady state of the model with nonnegativity constraint is 1 for availability, price, and production. Convenience yield implies that there is always some stock, even when price is constant, and there is no speculative opportunity, so the deterministic steady state differs slightly with stock, availability, price, and production, respectively, at 0.0338, 1.0336, 0.9998, and 1.0001. Supply elasticity is 0.2, and demand elasticity is -0.3 .

Table 1. Calibrated parameters

Parameter	δ	r	e	α	σ
Value	0.01	0.03	-0.3	5	0.10

For this calibration, the storage model reproduces some stylised facts about agricultural commodities markets (Table 2). Prices are positively serially correlated. Their distribution is asymmetric with a higher frequency of high than low prices, because storage alleviates low prices by stockpiling but cannot alleviate all episodes of high prices as stocks cannot be negative. The occurrence of zero stock periods should be small to preserve the serial correlation (it is storage that creates autocorrelation in prices) and to limit episodes of high prices. Because market availability is smoothed by the storage and supply reaction, final demand oscillates little (5% and 6% of coefficient of variation). In opposition to prices, final consumption is positively skewed, because high consumption is prevented by storage, and periods of low consumption occur because of stockouts.

Figure 1 illustrates the behaviour of the storage model. It shows the numerical difficulties we encounter when

Table 2. Storage model behaviour

	With nonnegativity constraint	With convenience yield
Occurrence of stockouts (%)	16	0
Price		
Autocorrelation	0.27	0.23
Coefficient of variation	0.20	0.24
Skewness	3.72	2.71
Consumption demand		
Coefficient of variation	0.05	0.06
Skewness	-2.04	-0.79

deciding about the different approximation schemes. The difficulties arise from the nonnegativity constraint of storage. Below a threshold availability, the expected marginal profit from stockpiling grains is negative, as the expected price does not cover storage costs and the high current price. For a negative expected marginal profit, there is no stock carried over to the next period. So the storage rule encompasses two regimes: a no-stock regime below the threshold and a stock increasing with availability above the threshold. This behaviour affects all the variables in the model. Below the threshold, the market price is set by consumption demand, but above it, the demand for storage adds to consumption demand. So the market demand function presents a kink at the availability threshold. The behaviour of producers depends on the expected price in the next period. This expected price is constant for low availability because there is no stock to connect successive periods. Therefore, planned production is constant until the threshold availability is reached and then decreases with the increase in stocks. For a more complete description of the properties of the storage model, see [Wright \(2001\)](#).

The model parameterised with convenience yield has no threshold and is smooth. However, its behaviour is similar to the model with inequality constraint: it displays highly nonlinear curves that feature a transition between a regime with stocks driven not by speculation but by convenience yield to a speculative regime. Planned production or expected price curves are strongly nonlinear when represented with respect to availability, but much more linear when stock is used in the x -axis (figures [1\(c\)](#) and [1\(d\)](#)).

3 Solution methods

3.1 General framework

Before detailing the methods used to solve the storage model, we present an overview of how rational expectations, stochastic, discrete time, continuous state models can be approximated. We adopt the framework used in [Fackler \(2005\)](#), and [Winschel and Krätzig \(2010\)](#). For brevity, we ignore models defined by complementarity equations, but the extension is straightforward. A rational expectations model is composed of four kinds of variables: s the state variables, x the response variables, z the expectations, and e the shocks.

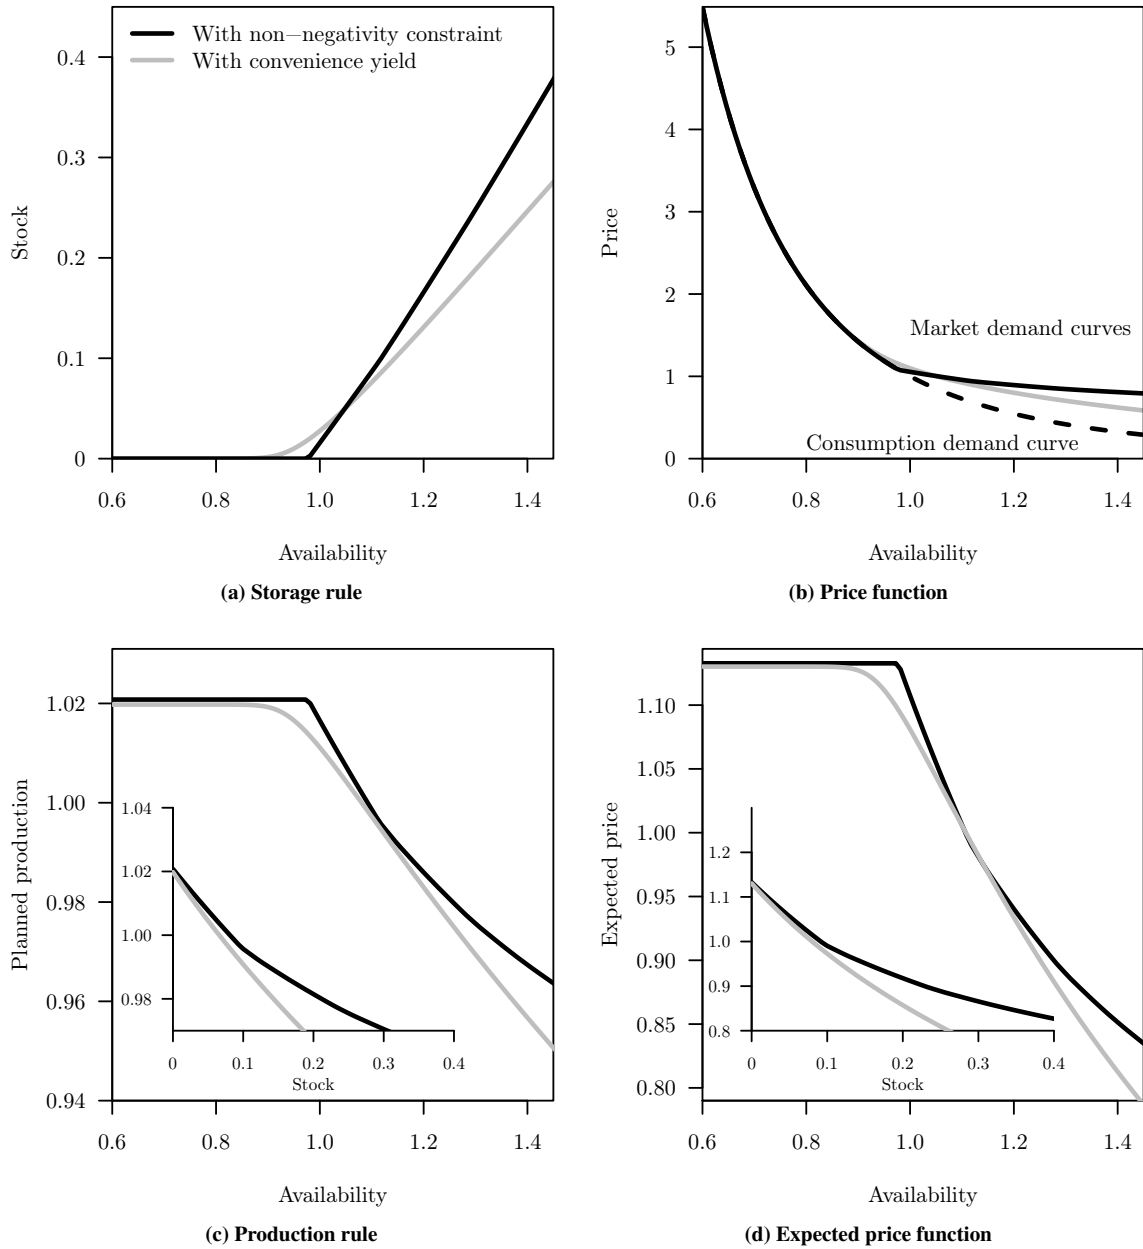


Figure 1. Characterisation of the storage model behaviour. Black curves correspond to the model with nonnegativity constraint, grey curves to the smoother model with convenience yield. For the price curves (b), the dashed line corresponds to demand for final consumption; total market demand (solid lines) includes demand for storage. For the production rule (c) and the expected price function (d), the plot includes another version of the same functions but drawn against the decision variable, stock, rather than the state variable, availability. These additional curves show the important differences in nonlinearity depending on the explanatory variable.

The model is defined by the following three groups of equations:

$$s_{t+1} = g(s_t, x_t, e_{t+1}) \quad \text{state transition,} \quad (8)$$

$$0 = f(s_t, x_t, z_t) \quad \text{equilibrium,} \quad (9)$$

$$z_t = E_t [h(s_t, x_t, e_{t+1}, s_{t+1}, x_{t+1})] \quad \text{expectations.} \quad (10)$$

For the storage model, the state transition equation is (5), and the equilibrium equations are (4) and (7). The expectations equations are already included in the equilibrium equations, but the expectations variable can be defined as the vector $z_t = [E_t(P_{t+1}\epsilon_{t+1}) \ E_t(P_{t+1})]$.

The difficulty of solving this type of models arises from the consistency between expectations and outcomes. Without expectations in Equation (9) or with exogenous expectations, for example with a backward-looking structure, the method would be to solve Equation (9) using a nonlinear solver and to iterate for next period with Equation (8). In a stochastic, infinite-horizon model, it is not possible to solve for all possible future situations in order to build consistent expectations. They are several general approaches to tackle this problem.

Perturbation methods A perturbation method considers equations (8)–(10) as a system and derives a Taylor approximation of the system around its steady state.

Projection methods Projection methods define (8)–(10) as a functional equation problem and find on a deterministic grid an approximation of a function allowing to define next-period conditions. Various functions can be used to approximate future conditions. They, however, differ in their degree of nonlinearity and their algorithmic performance, as we show below with the storage model.

For projection, four functions are approximated in the literature (we note below an approximating function as $\lambda(\cdot)$). (i) The response variables with respect to the state variables: $x_t = \lambda(s_t)$ (implemented, e.g., in Fackler, 2005, Winschel and Krätzig, 2010). (ii) The expectations with respect to the state: $z_t = \lambda(s_t)$ (den Haan and Marcet, 1990). (iii) The expectations with respect to the response: $z_t = \lambda(x_t)$ (Wright and Williams, 1984). (iv) When the expectations function, h , does not depend on time t variables and on shocks, it can be approximated as a function of next-period state: $h(\cdot, \cdot, \cdot, s_{t+1}, x_{t+1}) = \lambda(s_{t+1})$ (Miranda and Glauber, 1995). We compare below for the storage model approximations (i), (iii) and (iv).

Stochastic simulation As with projection methods, stochastic simulation methods, which are not studied in this paper, approximate a functional unknown characterising the expectations. The difference is that, contrary to projection methods that use a deterministic grid, with stochastic simulation, the solution domain is endogenous and generated by simulations. Between each simulation, the approximation is updated based on the discrepancy with respect to rational expectations (see, e.g., den Haan and Marcet, 1990, Judd et al., 2011).

Dynamic programming When the recursive equilibrium problem can be derived from an optimisation problem, such as a central planner problem, it can be solved using dynamic programming. In this case, equations (8)–(10) can be seen as the first-order conditions of

$$V(s_t) = \max_{x_t} \{b(s_t, x_t) + \beta E_t [V(s_{t+1})]\} \quad (11)$$

subject to Equation (8).

3.2 Existing approaches to solving the storage model

The seminal paper by [Gustafson \(1958\)](#) proposes several solution methods: iterating the value function, iterating the marginal value function, and approximating the storage rule by a piecewise linear function, but he carries his work with value function iteration alone. [Gardner \(1979\)](#) extends [Gustafson's](#) value function iteration method to a case with elastic supply (for an early review of the dynamic programming approach to the storage problem, see [Plato and Gordon, 1983](#)). Value function iteration is known to be a reliable method. It has convergence properties based on dynamic programming theory, but it is slow, and, since it solves a central planner problem, it does not enable finding a suboptimal competitive equilibrium.

Two more recent methods are generally used to solve the storage model. [Wright and Williams \(1982\)](#) were the first to propose a solution based on Euler equations. They use low-order polynomials to fit expected price to current stock (i.e., a parameterised expectations approach). They believe that low-order polynomials provide sufficient precision because the expected price function is smooth and not too nonlinear (the small plot in [Figure 1\(d\)](#)). They apply their method to various settings, such as a storage-trade model, a model featuring the interaction of public buffer-stock and private stock, and a storage model with news arrival ([Williams and Wright, 1991](#)). This numerical strategy is applied in several works on the storage model, such as [Miranda and Helmerger \(1988\)](#), [Lence and Hayes \(2002\)](#), and [Park \(2006\)](#).

Another method, which has two alternative forms, was proposed by [Deaton and Laroque \(1992\)](#) (a method similar to that applied to optimal consumption problems under liquidity constraints by [Deaton, 1991](#)) and [Miranda and Glauber \(1995\)](#). Both forms approximate the price function. The first solves the system using a fixed-point algorithm; the second applies a collocation method. [Miranda \(1997\)](#) compares various approximation schemes for this latter method.

The methods described above do not exhaust the possibilities. For a model with convenience yield, [Judd \(1998\)](#) suggests direct approximation of the storage rule but does not apply it. Approximating decision rules is a usual method for solving the stochastic growth model ([Aruoba et al., 2006](#)) and other dynamic models. The method is often used because it allows direct simulations once decision rules are identified. It is generally not used for the storage model because the storage rule is kinked, which makes it difficult to approximate. However, its limited accuracy could be compensated by higher speed; therefore, in the present paper, we implement decision rules approximation.

[Newbery and Stiglitz \(1982\)](#) propose to expand the storage rule as a Taylor series at the kink point. We adopt a similar method in this paper. The model is approximated at low orders around its steady state. The use of a Taylor expansion is restricted to smooth functions, so we apply it only to the case with convenience yield. This method, the perturbation method, is commonly used with the stochastic growth model.

With the exception of the perturbation methods, all the methods applied in this paper share the same set of numerical issues: how to pass from an integral over ϵ to a finite dimensional problem (i.e., how to discretise the shocks); how to solve the first-order conditions (4) and (7), which present specific problems caused by the complementarity condition; and which method to use to approximate the unknown functions. Following the treatment of these common issues, we present the different methods.

3.3 Shock discretisation

Productivity shocks ϵ are discretised, and the integrals over the shocks are calculated using a 7-node Gauss-Hermite quadrature, allowing the exact integration of degree 15 polynomials weighted by a normal distribution. The Gaussian formula transforms an expectation term into a sum. For example, next-period expected price expressed as a function of next-period availability is approximated as

$$E[P(A)] = E[P((1 - \delta)S + H\epsilon)] \approx \sum_{l=1}^L w_l P((1 - \delta)S + H\epsilon_l), \quad (12)$$

with ϵ_l and w_l the nodes and weights of the quadrature.

3.4 Solving the equations

The equilibrium equations of competitive storage model reduce to two equations: the storage arbitrage condition (7) and the producer incentive equation (4). The storage arbitrage condition takes a complementarity form when marginal storage costs are constant. This is not a traditional equation that can be solved with a regular nonlinear solver. For positive stocks levels, it behaves as does a traditional equation, with stock level adjusting to ensure the nullity of marginal profit. But when stock hits the zero lower bound, stock level does not adjust anymore, and marginal profit is allowed to become negative. We solve this complementarity problem by using the `ncpsolve` solver (Miranda and Fackler, 2002).² When the inequality constraint is smoothed by the convenience yield hypothesis, the problem can be solved using any nonlinear equation solvers.³ In our case, we use also the `ncpsolve` solver, which can also solve simple nonlinear problems. Equations are solved with a precision of 10^{-9} .

Given the simplicity of the problem, we could avoid having to use a complementarity solver. The two equations can be solved by any nonlinear solver, and, once a solution is found, the positivity of storage can be checked. For negative storage, we would force it at zero and solve for the planned production. This is the approach taken by Williams and Wright (1991, p. 83). However, this is not advisable for a problem with more complementarity equations, such as a multi-commodity model or a storage-trade model. A complementarity solver introduces only limited overheads compared with a traditional nonlinear solver, so we prefer this more direct solution.

With inelastic supply, it is necessary to solve only the storage arbitrage condition. The problem can be written as a fixed-point problem and solved through simple operations without the need for solvers (Deaton and Laroque, 1992, Judd, 1998). Deaton and Laroque (1992) propose this solution, and it has been applied in most econometric work on the storage model because of its speed (Chambers and Bailey, 1996, Michaelides and Ng, 2000, Osborne, 2004).

²For more complex complementarity problems, the use of more robust solvers can be contemplated. PATH (Dirkse and Ferris, 1995) and LMMCP (Kanzow and Petra, 2004, 2007) are both good candidates for this task.

³Even if smoother, this problem is not simple to solve. Indeed, the logarithm does not tolerate negative storage level, and the solver must be prevented from exploring this region.

3.5 Approximation methods

The first numerical methods for the storage model involve approximating the value function by space discretisation (Gustafson, 1958, Gardner, 1979). Later, Wright and Williams (1984) use low-order polynomials to approximate conditional expectations. The polynomials are fitted to grid values by ordinary least-squares. Miranda and Glauber (1995) solve the storage model using Chebychev collocation methods.

The virtues of the various approximation schemes used to study dynamic economic problems have been discussed by various authors (see, e.g., Judd, 1992, 1998, Christiano and Fisher, 2000). The choice simply reduces to splines versus Chebychev polynomials. In this context, the storage model presents a specific difficulty. Because of the nonnegativity constraint on stocks, most functions of the model present a kink. They are continuous but not differentiable (see Figure 1), which creates problems for both interpolation schemes and requires a high number of free parameters to achieve a satisfactory level of precision. The difficulty is greater with polynomials, because the discontinuity affects the entire interpolation interval (Miranda and Fackler, 2002), which means that, in this case, splines should behave better. For the storage model with convenience yield, which has no kink, Miranda (1997) compares various approximation schemes. He finds that approximation by cubic splines is preferable to the other approaches. We follow his result and use only cubic splines interpolation.

The interpolation is done with the Matlab CompEcon toolbox, developed to accompany Miranda and Fackler (2002). Functions defined over storage are approximated over the interval $[0, 0.5]$ ($[2.2 \times 10^{-16}, 0.5]$ with convenience yield). Functions defined over availability are approximated over the interval $[\min(\epsilon_l), 1.7]$. These intervals were chosen to include the values we are most likely to find during a simulation. They are determined by tâtonnement until the smallest intervals including most points inside the empirical distribution were found. The bounds play a crucial role in the precision achieved by an algorithm. Since most of the distribution is included in these intervals, enlarging them can only decrease precision.

Values outside intervals may be encountered during the resolution of the functional equations problem, and, in these cases, there is no extrapolation: values are taken as approximations of the last points of the interval. For all methods except perturbation where we rely on a separate software, convergence to the true approximating function is assured by a successive approximation through a time iteration approach.

Algorithm performance and convergence depend greatly on the initial guess about the approximated function. Our initial guess is based on the corresponding functions at steady-state values. When price or expected price is approximated, we correct the first guess for the expected effects of storage. Storage will occur in periods of low prices and thus will prevent situations of extremely low prices. We, therefore, restrict the first guess to values above 70% of the steady state.

Approximation methods make it possible to transform an infinite-dimension problem (finding a function satisfying some conditions over a continuum) into a finite-dimension problem (finding the values satisfying the conditions at some nodes and approximate between the nodes). The function $f(\cdot)$ —to be defined later for each method—is approximated by $f_n(\cdot)$: $f(x) \approx f_n(x) = \sum_{j=1}^J \theta_j^n B_j(x)$. θ_j^n are basis coefficients to be determined, and $B_j(\cdot)$ are the basis functions of a cubic spline (Judd, 1998). At each iteration of n , the

basis coefficients are updated in a fitting step. Except for the endogenous grid method where the coefficients are updated through a least-squares regression, the fitting step consists of solving the linear system of J equations in the unknown θ_j^n :

$$\sum_{j=1}^J \theta_j^n B_j(x_i) = f(x_i) \quad \text{for } i = 1, \dots, J, \quad (13)$$

with $\{x_1, x_2, \dots, x_J\}$ as the collocation nodes.

3.6 Value function iteration

The model can be recast as a planner problem (Scheinkman and Schechtman, 1983). Its recursive dynamic formulation is

$$V(A_t) = \max_{S_t \geq 0, H_t} \int_{\bar{D}}^{A_t - S_t} P(D) dD - \Phi(S_t) - \Psi(H_t) + \frac{1}{1+r} E_t V((1-\delta)S_t + H_t \epsilon_{t+1}), \quad (14)$$

where $V(\cdot)$ is the value function of the problem. When \bar{D} is a positive number, $\int_{\bar{D}}^{A_t - S_t} P(D) dD$ is gross consumer benefit. Without any loss of generality, we make \bar{D} sufficiently high such that $\int_{\bar{D}}^{A_t - S_t} P(D) dD = (A_t - S_t)^{1+1/e} / (1 + 1/e)$.

To solve by value function iteration, we define a grid on availability, $\{A_1, A_2, \dots, A_J\}$. For a given value function, it is necessary to determine optimal storage and planned production by solving the two arbitrage conditions. The value function approximation can be updated using these optimal decisions. We stop the iterations when changes in the value function decrease below a threshold. The algorithm runs as follows:

1. Initialise by taking the first guess: $V_0(A) = A^{(1+1/e)/(1+r)} - \Psi(H^{SS})$,⁴ with $V_0(\cdot)$ as a J -breakpoint spline.
2. For each A_i in $\{A_1, A_2, \dots, A_J\}$, find S_i and H_i that solve the following first-order conditions:

$$S_i \geq 0 \quad \perp \quad -P(A_i - S_i) - \Phi'(S_i) + \frac{1-\delta}{1+r} \sum_{l=1}^L w_l V_n'((1-\delta)S_i + H_i \epsilon_l) \leq 0, \quad (15)$$

$$-(1+r)\Psi'(H_i) + \sum_{l=1}^L w_l \epsilon_l V_n'((1-\delta)S_i + H_i \epsilon_l) = 0. \quad (16)$$

3. Update the J -breakpoint spline $V_{n+1}(\cdot)$ using the system of J equations:

$$V_{n+1}(A_i) = \frac{(A_i - S_i)^{1+1/e}}{1 + 1/e} - \Phi(S_i) - \Psi(H_i) + \frac{1}{1+r} \sum_{l=1}^L w_l V_n((1-\delta)S_i + H_i \epsilon_l) \quad \text{for } i = 1, \dots, J. \quad (17)$$

⁴Variables with the superscript SS are the deterministic steady-state values.

$V_n((1 - \delta) S_i + H_i \epsilon_i)$ and its derivative are computed by interpolation.

4. If $\|V_{n+1}(A) - V_n(A)\|_2 \geq 10^{-7}$, then increment n to $n + 1$ and go to step 2.

Comparing this crude dynamic programming algorithm with the projection methods that follow is unfair, because there are several procedures that could accelerate a value function iteration algorithm (Rust, 1996). But given VFI is limited to models amenable to a planner problem, a deeper analysis of this algorithm does not seem relevant.

3.7 Projection methods

Projection methods are the most standard methods applied to the storage model (Wright and Williams, 1982, Miranda and Helmberger, 1988). Their use was formalised in economics by Judd (1992) and consists of defining an approximation scheme and using it to minimise a residual function that defines the rational expectations equilibrium. Storage model is a perfect test-case for projections methods, because they can be applied in various ways depending on the functions being approximated. The present paper demonstrates that choosing the smoothest function for approximation can result in significant precision gains.

Projection methods can differ in terms of which functions they approximate, but the residual is always defined by the two Euler equations (4) and (7). As it is not possible to make the residual function zero for all points, several possibilities exist for defining objectives. They differ in their definition of the norm over the residual. The Galerkin method forces the inner product of the residual with the basis functions to be zero. The collocation method requires the interpolant to make the residual function equal to zero for all nodes. Collocation only obliges that the model be solved for the same number of points as the number of free parameters in the interpolant and is simpler to explain; thus, in this paper, we use collocation for all projection methods, except for endogenous grid.

We follow the usual practice in the literature on storage model in implementing the simplest method to converge to the true rational expectations equilibrium (i.e., time iteration). At each iteration step, the new approximation is defined by the application of the current approximation to the next-period problem. The algorithm stops when the function used for two successive periods are almost identical. The collocation conditions could also be approached as a nonlinear problem and solved using a Newton solver.

We present three versions of the projection method, which differ only in the functions they approximate: price approximation (expectations function approximation); parameterisation of the expected price (parameterised expectations algorithm, PEA); storage rule approximation (decision rules approximation).

3.7.1 Expectations function approximation

Approximating the relation between price and availability is equivalent to approximating the expectations function (the projection method (iv) in Section 3.1), since the expectations only depend on price and productive shock. The literature proposes two approaches to solving the storage model by parameterising

the price function. **Deaton and Laroque (1992)** propose a fixed-point algorithm, which is very fast because it does not require any root-finding step but is applicable only to cases without supply reaction. So it cannot be used here. The other approach is suggested by **Miranda and Glauber (1995)**, who use a projection algorithm with Chebychev collocation.

In this paper, we implement two algorithms to solve the model with a parameterised price function. First, **Miranda and Glauber's** approach, which we call time iteration approach. Second, a fixed-point algorithm, based on the endogenous grid method proposed by **Carroll (2006)**, which is faster than **Deaton and Laroque's** algorithm and is applicable to all the situations of interest in this paper.

The time iteration approach **Miranda and Glauber's** algorithm and the value function iteration are almost identical. The former iterates the marginal of the value function—i.e., the price function—rather than the value function. More precise results can be expected because VFI requires a good approximation of both the value function and its derivative.

1. Initialise $g_0(A) = \max [P(A), 0.7P^{SS}]$ with $g_0(\cdot)$ as a J -breakpoint spline.
2. For each A_i in $\{A_1, A_2, \dots, A_J\}$, find S_i and H_i that solve the first-order conditions

$$S_i \geq 0 \quad \perp \quad -P(A_i - S_i) - \Phi'(S_i) + \frac{1 - \delta}{1 + r} \sum_{l=1}^L w_l g_n((1 - \delta)S_i + H_i \epsilon_l) \leq 0, \quad (18)$$

$$-(1 + r)\Psi'(H_i) + \sum_{l=1}^L w_l \epsilon_l g_n((1 - \delta)S_i + H_i \epsilon_l) = 0. \quad (19)$$

3. Determine the J -breakpoint spline $g_{n+1}(\cdot)$ that solves the following system:

$$g_{n+1}(A_i) = P(A_i - S_i) \quad \text{for } i = 1, \dots, J. \quad (20)$$

4. If $\|g_{n+1}(A) - g_n(A)\|_2 \geq 10^{-7}$, then increment n to $n + 1$ and go to step 2.

The endogenous grid method **Carroll (2006)** proposes a very efficient method for solving small-scale dynamic stochastic optimisation problems. The idea is to solve the Euler equation by iterations of the end-of-period decision variable rather than the state variable. This suppresses the need to solve a nonlinear equation and replaces the root-finding step by arithmetic operations. Applied to the model with convenience yield and inelastic supply, it would simply consist of iterating the function $k_n(A)$ on a grid of storage points $\{S_1, S_2, \dots, S_M\}$. $k_n(A)$ approximates the price function, with

$$A_i = S_i + P^{-1} \left(\frac{1 - \delta}{1 + r} E [k_n((1 - \delta)S_i + \bar{H}\epsilon_l)] - \Phi'(S_i) \right), \quad (21)$$

and $k_{n+1}(A_i) = P(A_i - S_i)$. The grid $\{A_1, S_2, \dots, A_M\}$ on which the function is approximated changes at each iteration, hence the name endogenous grid method. Since there is no root-finding step, we do not need

the derivative of the function k_n . It can thus be approximated by linear interpolation instead of spline, which helps further speed the method.

Carroll (2006) explains also how the method can be applied to problems with inequality constraints. One needs only to change the iteration step in k_n and make sure that the grid on S includes zero:

$$k_{n+1}(A) \text{ approximates } \begin{cases} P(A_i - S_i) & \text{for all grid points,} \\ P(A) & \text{for } A \leq \min(A_i). \end{cases}$$

As soon as there is more than one decision variable, the methods becomes trickier. Barillas and Fernández-Villaverde (2007) extend the method to a case with two decision variables. As these decision variables must be mutually consistent, it is not possible to define a grid on both variables. The idea is to define the grid on one decision variable and to find the other by solving the corresponding first-order equation. In our case, we define a grid on storage, solve Equation (4) for planned production, and compute availability and price with Equation (21). We lose one of the interests of EGM (i.e., avoiding root-finding step), but in this case, root-finding is quite simple, and we can accelerate the algorithm by applying it every two or three iterations. The algorithm runs as follows:

1. Initialise $k_0(A) = \max [P(A), 0.7P^{SS}]$ with $k_0(\cdot)$ as a J -breakpoint spline.
2. For each storage level S_i in $\{S_1, S_2, \dots, S_M\}$ find H_i that solve the Euler equation for production

$$-(1+r)\Psi'(H_i) + \sum_{l=1}^L w_l \epsilon_l k_n((1-\delta)S_i + H_i \epsilon_l) = 0. \quad (22)$$

To accelerate the algorithm, this step can be applied every two or three iterations.

3. For each combination $\{S_i, H_i\}$, calculate A_i with

$$A_i = S_i + P^{-1} \left(\frac{1-\delta}{1+r} \sum_{l=1}^L w_l k_n((1-\delta)S_i + H_i \epsilon_l) - \Phi'(S_i) \right). \quad (23)$$

4. Determine the J -breakpoint spline $k_{n+1}(\cdot)$ that best fits (in a least-squares meaning) the points $\{A_i, P(A_i - S_i)\}$ and the function $P(A)$ for $A < \min(A_i)$.
5. If $\|k_{n+1}(A) - k_n(A)\|_2 \geq 10^{-7}$, then increment n to $n + 1$ and go to step 2.

Unlike the other methods, we take more grid points for the EGM than the number of breakpoints in the approximating function. Consequently, the approximation is updated by a least-squares step. This is required by the fact that the definition of the spline requires at least one observation between two breakpoints, which cannot be guaranteed if we take as many observations as breakpoints because breakpoints and grid points are chosen independently.

The time iteration approach and the endogenous grid method being two solution algorithms to approximate the relation between price and availability should lead to close precision performances, except for the small differences introduced by the fitting step (collocation or least squares). In the following, thus, they are only distinguished with respect to their computing time.

3.7.2 Parameterised expectations algorithm

The parameterised expectations algorithm approximates the expectations that appear in the Euler equations. Popularised in the macroeconomics literature by [den Haan and Marcet \(1990\)](#), this algorithm was originally developed by [Wright and Williams \(1982\)](#) for the storage model.⁵ [Christiano and Fisher \(2000\)](#) note that the two methods are slightly different. [Den Haan and Marcet](#) approximate conditional expectations with respect to the state variables, whereas [Wright and Williams](#) approximate conditional expectations with respect to the current-period decision.⁶

In the case of the storage model, the [den Haan and Marcet](#) PEA would approximate $E(P_{t+1}|A_t)$, whereas [Williams and Wright](#) parameterise $E(P_{t+1}|S_t)$. The latter function should be smoother than the former, because expected price with respect to availability shows two regimes with a kink and is quite nonlinear. For low availabilities, the expected price is constant. It is the stockout case. After a threshold, expected price decreases with current availability. The relation between expected price and storage is smoother and more linear, as it does not include the first stockout regime (compare the two plots in [Figure 1\(d\)](#)). This difference was already appreciated in the case of the stochastic growth model with irreversible investment in [Christiano and Fisher \(2000\)](#). In what follows, we implement only [Wright and Williams's](#) approach.

The algorithm below is adapted from [Wright and Williams \(1984\)](#). This method requires approximation of two functions: expected price with respect to storage, $f_n^S(S) \approx E(P|S)$, and producer incentive price with respect to storage, $f_n^H(S) \approx E(P\epsilon|S)$.⁷

1. Initialise $f_0^H(S) = E\{\max[P((1-\delta)S + H^{SS}\epsilon - S^{SS}), 0.7P^{SS}]\epsilon\}$ and $f_0^S(S) = E\{\max[P((1-\delta)S + H^{SS}\epsilon - S^{SS}), 0.7P^{SS}]\}$; both are J -breakpoint splines.
2. For each storage level S_i in $\{S_1, S_2, \dots, S_J\}$ and shocks ϵ_l define the following availability:

$$A_{i,l} = (1 - \delta) S_i + (\Psi')^{-1} \left(f_n^H(S_i) / (1 + r) \right) \epsilon_l. \quad (24)$$

3. For each element $A_{i,l}$, solve the storage Euler equation to find $S_{i,l}$

$$S_{i,l} \geq 0 \quad \perp \quad \frac{1 - \delta}{1 + r} f_n^S(S_{i,l}) - P(A_{i,l} - S_{i,l}) - \Phi'(S_{i,l}) \leq 0. \quad (25)$$

⁵They describe the algorithm in [Wright and Williams \(1984\)](#).

⁶The two versions differ also in the method used to compute the approximation. [Wright and Williams](#) use a projection method, and [den Haan and Marcet](#) implement a stochastic simulation approach, dubbed by [Judd et al. \(2011\)](#) as simulation-based PEA. Stochastic PEA may be more intuitive than its projection counterpart, but it is less accurate and its convergence is not warranted, so it is not implemented here.

⁷[Wright and Williams \(1984\)](#) approximate directly the value of planned production instead of the producer incentive price. The two schemes are equivalent, but to follow the logic of PEA, we parameterise the expectation term.

4. Determine the J -breakpoint splines $f_{n+1}^S(\cdot)$ and $f_{n+1}^H(\cdot)$ that solve the systems

$$f_{n+1}^S(S_i) = \sum_{l=1}^L w_l P(A_{i,l} - S_{i,l}) \quad \text{for } i = 1, \dots, J, \quad (26)$$

$$f_{n+1}^H(S_i) = \sum_{l=1}^L w_l \epsilon_l P(A_{i,l} - S_{i,l}) \quad \text{for } i = 1, \dots, J. \quad (27)$$

5. If $\|f_{n+1}^S(S) - f_n^S(S)\|_2 \geq 10^{-7}$ or $\|f_{n+1}^H(S) - f_n^H(S)\|_2 \geq 10^{-7}$, then increment n to $n + 1$ and go to step 2.

We expect this method to be more precise than the others, since the conditional expectations are smoother than most of the other functions to be approximated in the model. However, this method requires two conditional expectations to be approximated, each corresponding to a different decision variable. The parameterised price approach or VFI, on the other hand, are more parsimonious, because they require only one function approximation.

3.7.3 Decision rules approximation

Although widely used with other dynamic models, decision rules approximation has not been applied to solving the storage model. The method was suggested in Judd (1998) for the case with convenience yield, but here we apply it to both cases: convenience yield and nonnegativity constraint. In a time iteration approach, the approximated storage rule is used inside the expectation terms to define next-period storage. Storage undertaken at the current period, found by solving the first-order conditions over a grid of availability points, allows the approximation to be updated. The production rule is not necessary to define the rational expectations equilibrium.

1. Initialise $s_0(A) = 0$ with $s_0(\cdot)$ as a J -breakpoint spline.
2. For each A_i in $\{A_1, A_2, \dots, A_J\}$, find S_i and H_i that solve the first-order conditions

$$S_i \geq 0 \quad \perp \quad -P(A_i - S_i) - \Phi'(S_i) + \frac{1 - \delta}{1 + r} \sum_{l=1}^L w_l P((1 - \delta)S_i + H_i \epsilon_l - s_n((1 - \delta)S_i + H_i \epsilon_l)) \leq 0, \quad (28)$$

$$-(1 + r) \Psi'(H_i) + \sum_{l=1}^L w_l \epsilon_l P((1 - \delta)S_i + H_i \epsilon_l - s_n((1 - \delta)S_i + H_i \epsilon_l)) = 0. \quad (29)$$

3. Update the approximation by finding the J -breakpoint spline $s_{n+1}(\cdot)$ that solves the system

$$s_{n+1}(A_i) = S_i \quad \text{for } i = 1, \dots, J. \quad (30)$$

4. If $\|s_{n+1}(A) - s_n(A)\|_2 \geq 10^{-7}$, then increment n to $n + 1$ and go to step 2.

The decision rules approximation method can also be solved by fixed-point iteration, as suggested in Judd (1998) or by EGM.

3.8 Perturbation methods

For the storage model, perturbation methods have been barely considered. Miranda (1997) tries linearisation of the decision rules, but not higher-order approximations. In other settings, particularly dynamic stochastic general equilibrium modelling, perturbation methods are widely used because they allow one to work with a large number of state variables and are easy to implement through user-friendly packages such as Dynare. The method consists of approximating the dynamic problem around the steady state by its Taylor development, which requires a smooth problem.

The storage problem is normally not amenable to such methods because of the nonnegativity constraint. The introduction of convenience yield smoothes the model and allows the use of perturbation. Convenience yield has some theoretical underpinnings in the storage literature (see Section 2.2), but, more generally, this transformation from a problem with inequality constraint to a smooth problem is not uncommon in numerical methods, where it is called the penalty function or barrier method (Judd, 1998, p. 123–125). It is used in the macroeconomic context by Preston and Roca (2007) and Kim et al. (2010) to smooth the borrowing constraint in a consumption/saving decision.

The storage model with convenience yield is approximated around its steady state from the first to the fifth order by using the software Dynare++, included in Dynare v. 4.1.3.⁸

4 Numerical results

A benchmark solution is used to generate the true distribution of the state variable. The benchmark solution is the PEA with 5,000 nodes. The benchmark simulation is defined over 10,000 periods. The same shocks are used for both versions of the model.

4.1 Euler equation error

Because the storage model has no analytical solution, we cannot compare the approximate solutions obtained from the various methods with an exact benchmark. For this case, Judd (1992) designed a bounded rationality measure, the Euler equation error, to measure by how much solutions violate the optimising conditions. This measure defines the resources wasted by using approximated decision rules instead of true ones. In what follows, we derive the definitions of the two measures of the Euler equation error for the storage model. This

⁸<http://www.dynare.org>

differs from the literature on numerical methods in that the measure has to account for the existence of two regimes in the Euler equation.

Expressed in terms of the decision rules, $S(A)$ and $H(A)$, the storage Euler equation is

$$S(A_t) \geq 0 \quad \perp \quad \frac{1-\delta}{1+r} E_t \left[P((1-\delta)S(A_t) + H(A_t)\epsilon_{t+1}) - S((1-\delta)S(A_t) + H(A_t)\epsilon_{t+1}) \right] - P(A_t - S(A_t)) - \Phi'(S(A_t)) \leq 0. \quad (31)$$

To simplify the notations in the following equations, we use $A_{t+1} = (1-\delta)S(A_t) + H(A_t)\epsilon_{t+1}$. We need to distinguish two situations. First, when storage costs include a convenience yield, the Euler equation holds exactly without a complementarity condition. This leads to the following Euler equation:

$$P(A_t - S(A_t)) = \frac{1-\delta}{1+r} E_t \left[P(A_{t+1} - S(A_{t+1})) \right] - \Phi'(S(A_t)). \quad (32)$$

As the decision rules are approximations, this equation will not hold exactly for all availabilities. We can define the Euler equation error function with convenience yield as

$$EE_S(A_t) = 1 - \frac{P^{-1} \left(\frac{1-\delta}{1+r} E_t \left[P(A_{t+1} - S(A_{t+1})) \right] - \Phi'(S(A_t)) \right)}{A_t - S(A_t)}. \quad (33)$$

Second, for the case without convenience yield, we can rewrite the Euler equation as

$$P(A_t - S(A_t)) = \max \left\{ P(A_t), \frac{1-\delta}{1+r} E_t \left[P(A_{t+1} - S(A_{t+1})) \right] - \Phi'(S(A_t)) \right\}. \quad (34)$$

This formulation displays the two regimes. For low next-period expected price, there is no storage, and the current price is given simply by the inverse demand function applied to current availabilities. For positive storage, current price is linked to next-period price. As the decision rules are approximations, this equation does not hold exactly for all availabilities. We can define the Euler equation error as

$$EE_S(A_t) = 1 - \frac{P^{-1} \left(\max \left\{ P(A_t), \frac{1-\delta}{1+r} E_t \left[P(A_{t+1} - S(A_{t+1})) \right] - \Phi'(S(A_t)) \right\} \right)}{A_t - S(A_t)}. \quad (35)$$

As $D_t = A_t - S(A_t)$, this error can be interpreted as a unit-free representation of the amount of commodity consumption that should adjust to make the Euler equation hold exactly.

The error for the production Euler equation is derived in the same way:

$$EE_H(A_t) = 1 - \frac{(\Psi')^{-1} \left(\frac{1}{1+r} E_t \left[P(A_{t+1} - S(A_{t+1}))\epsilon_{t+1} \right] \right)}{H(A_t)}. \quad (36)$$

It is equivalent to the level of production that should be adjusted to make the Euler equation hold.

We consider two measures of the Euler equation error: its maximum and its average. Its maximum is taken

over an interval of availability defined by the first and the last percentiles of the benchmark simulation. The maximum error is useful to identify the algorithm limits. Average error is the error arising from the use of decision rules for a sequence of decisions. We calculate it by averaging absolute Euler equation errors over the benchmark simulation. To facilitate their reading, the errors are presented in a base-10 logarithm. A result of -3 for $\log_{10} |EE_S|$ should be read as a \$1 error in every \$1,000 of commodity consumption; -6 , \$1 error for \$1,000,000 of consumption.

4.2 Limitations of perturbation methods

Because perturbation methods perform significantly worse than other methods, we give them special treatment in order to understand why they are of limited interest in this case, whereas it is the method of choice in several other settings.⁹ Figure 2 presents the storage rules obtained by perturbation methods, at various orders, and compares them with the true one, obtained using the benchmark method. A figure for the production rule would be similar, so it is not displayed here.

Given the nonlinearity of decision rules, a first-order perturbation could not be expected to yield an appropriate approximation. For low availability, the first-order storage rule proposes negative storage and presents a marginal propensity to store current availability much lower than the true rule. These are the classical properties of the certainty equivalent solution to the consumption/saving model under uncertainty (Zeldes, 1989). The certainty equivalent means that the variance of future shocks does not matter. The first-order storage rule passes by the deterministic steady state, which is not part of the true storage rule, because, under certainty, there is no speculative motive for storing, only a convenience yield. Beyond first-order perturbation, the approximated decision rules follow the true one quite closely for availability between 0.94 and 1.16, which corresponds to 69% of the distribution. Below or beyond these bounds, approximations are very poor. Odd-order approximations present negative storage for high availability. And despite the fact that Blanchard-Kahn conditions are satisfied, 4th and 5th order solutions are unstable. They are locally stable around the steady state but lose stability farther away from it. The best precision is achieved with second-order approximation. This degradation of precision with perturbation order was identified by Feigenbaum (2005) for the consumption problem and by Aruoba et al. (2006) for one parameterisation of the stochastic growth model.

There are two problems related to perturbation methods applied to the competitive storage model. Even if smoothed by the convenience yield, the storage model remains highly nonlinear, which calls for high-order perturbation. As perturbation methods are in essence local methods, they should behave correctly around the steady state, which is what we observe. But the typical shocks in agricultural production drive the behaviour in regions distant from the steady state. This work confirms some recent evidence that higher-order perturbation may not be well-behaved. Feigenbaum (2005) and den Haan and de Wind (2009) note that a drawback of perturbation methods is the lack of control over the radius of convergence. In this case, we see that the radius of convergence covers only a small part of the availability distribution. The use of

⁹They are also treated apart because the storage rules found by perturbation are negative for some availability. The Euler equation error, used to assess the precision of the other methods, is not defined for negative storage, because of the logarithm in $\Phi'(\cdot)$.

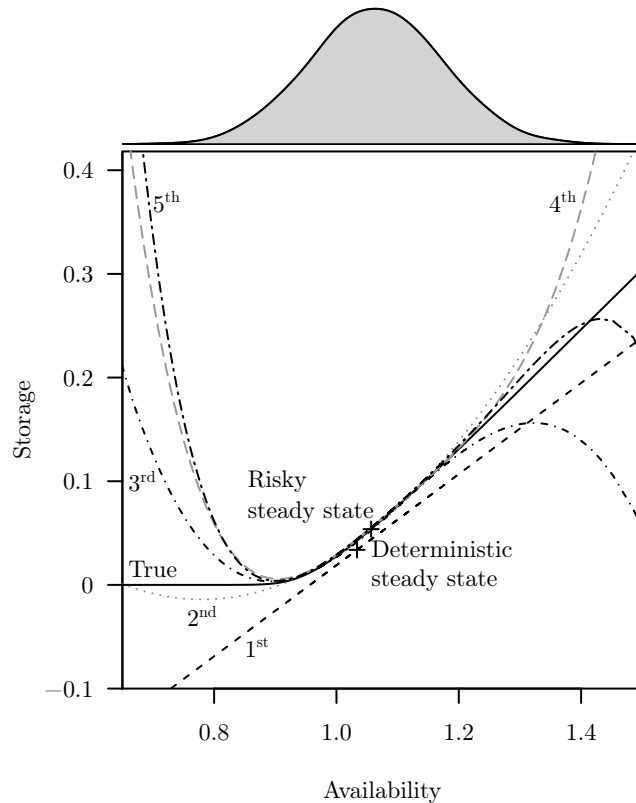


Figure 2. Storage rules obtained with perturbation methods up to the 5th order. Density of availability from the benchmark model above the plot. The “True” storage rule is generated by the benchmark method. Deterministic steady state is the state reached in the absence of shocks and ignoring future shocks. Risky steady state is the fixed point of the stochastic problem in the absence of shocks (i.e., the state reached in the absence of shocks but accounting for the likelihood of future shocks; [Coerdacier et al., 2011](#)).

a logarithm as the barrier function may have increased the lack of convergence of the Taylor development, because its radius of convergence is limited. But [den Haan and de Wind \(2009\)](#) show that this problem arises also with a well behaved barrier, such as an exponential.

4.3 Comparison of methods precision

Before comparing methods, it is useful to assess the relation between storage and production Euler equation errors, and the precision achievable from the most precise method. [Table 3](#) presents detailed results for the PEA—proved below to be the most accurate algorithm—when the number of breakpoints in the spline approximation varies.

The production Euler equation error is systematically lower than the storage Euler equation error. However, they are close and vary in the same way. Augmenting the number of nodes improves the precision slowly for the model with nonnegativity constraint. The general shape is well approximated at a low order (for information, the papers of [Williams and Wright](#) that have derived most of theoretical properties of the storage

Table 3. The effect of approximation order on precision ($\log_{10} |EE|$) for the parameterised expectations algorithm

Order	Storage Euler equation		Production Euler equation	
	Max error	Average error	Max error	Average error
With nonnegativity constraint				
5	-2.88	-3.20	-3.10	-3.43
10	-2.80	-3.57	-3.03	-3.81
20	-3.20	-4.20	-3.43	-4.43
50	-3.91	-5.14	-4.14	-5.37
100	-4.25	-5.76	-4.48	-5.99
200	-4.23	-6.27	-4.45	-6.50
1,000	-5.26	-7.79	-5.48	-8.02
With convenience yield				
5	-3.54	-3.94	-3.79	-4.15
10	-4.68	-5.12	-4.95	-5.37
20	-6.26	-6.87	-6.48	-7.10
50	-8.17	-8.74	-8.38	-8.96
100	-9.45	-9.99	-9.65	-10.22
200	-10.28	-10.87	-10.31	-11.00
1,000	-12.30	-12.56	-12.68	-12.87

model used an order-5 polynomial). With convenience yield, the expected price approximation improves faster with the number of breakpoints. Based on these first results, we decide to restrict presentation of the subsequent results to the storage Euler equation error and to a 20-breakpoint approximation.

Table 4 presents the precision achieved by the various methods. The PEA is consistently more precise than other methods. Expectations function approximation and decision rules approximation achieve very close precision. This is confirmed at various orders in Figure 3. Value function iteration performs similarly to expectations function or decision rules approximation. For all methods, the logarithm of the Euler equation error decreases linearly with the logarithm of the approximation order. The mean slope is -2 , which means a 10-fold increase in the number of breakpoints leads to an error decrease of 2. If at 20 breakpoints, the average error is \$1 in every \$1,000, at 200 breakpoints, it will be \$1 in every \$100,000. Whatever the approximation level, PEA is the most precise method, with always one order of precision below the other methods.

Table 4. Storage Euler equation error ($\log_{10} |EE_S|$) for various methods with a 20-breakpoint spline approximation

Method	With nonnegativity constraint		With convenience yield	
	Max error	Average error	Max error	Average error
Value function iteration	-3.06	-3.86	-3.90	-4.20
Parameterised expectations	-3.20	-4.20	-6.26	-6.87
Expectations function approximation	-2.66	-3.24	-4.10	-4.33
Decision rules approximation	-2.66	-3.24	-4.23	-4.52

4.4 Simulating the model

We now address an issue often neglected in computational economics, which is how the model should be simulated. There are two principal approaches. The first, and most commonly used, consists in approximating the decision rules and simulating the model by applying them recursively. The second uses the function being approximated (the decisions rules, the conditional expectations, or the expectations function as shown in Section 3.1) to approximate next-period expectations and solve the equilibrium equations to find the current decisions (in a time-iteration approach). For example, if applied to the approximation of the conditional expectations ($z_t = \lambda(x_t)$), this method would solve for x_t the following equation:

$$0 = f(s_t, x_t, \lambda(x_t)). \quad (37)$$

The two approaches differ in speed and in precision. Simulating a model through the recursive application of approximated decision rules is very fast because it does not require any nonlinear solve. The second approach is at least 10 times slower in the case of the storage model because each period requires solving a system of complementarity equations. However, its precision is much higher, because the approximation is used only for the expectations of next-period conditions. The precision gain is all the more important when decision rules have kinks, which makes them difficult to approximate. Wright and Williams (1984), who first proposed the PEA, suggest using the second method for simulating the storage model.

The above results were produced by the second, more precise method. In contrast, Table 5 presents the Euler equation errors achieved when approximating the true decision rules at various orders and using




Figure 3. Average storage Euler equation error for the model with nonnegativity constraint, for all methods, at different approximation orders.

this approximation for simulation. The decision rules approximated by a 20-breakpoint spline generate maximum and average storage Euler equation errors of -2.04 and -2.96 , a level of precision lower than what was achieved previously by explicitly solving the first-order conditions (see Table 4 for comparison). This applies especially to the maximum error, which decreases very slowly with the number of breakpoints because of the difficulty involved in approximating the kinks.

Table 5. The effect of approximation order on precision ($\log_{10} |EE|$) for a spline approximation of the true decision rules, for the model with nonnegativity constraint.

Order	Storage Euler equation		Production Euler equation	
	Max error	Average error	Max error	Average error
5	-1.35	-1.82	-1.89	-2.07
10	-1.77	-2.50	-2.87	-3.02
20	-2.04	-2.96	-3.04	-3.39
50	-2.50	-3.88	-3.82	-4.59
100	-3.08	-4.47	-4.11	-5.14
200	-3.14	-5.00	-4.26	-5.64
1,000	-4.05	-6.47	-4.98	-7.04

4.5 Implementation and computing time

This subsection focuses on the time needed to solve the model. As all the existing econometric work uses the model without supply reaction, we include a comparison of computing time over this more restricted model. Comparison of its Euler equation errors would be less interesting because they are very close to those for the model with supply reaction.

The results reported are for a 2.6 GHz Intel PC running Windows Vista and Matlab R2010b. We would expect a faster solve for a model with inelastic supply, because the model is simpler. This is not the case for all methods (see Table 6). While the model has one less equation, the function being approximated is farther away from the first guess and so requires more iterations. The kink point of the storage rule appears at a lower availability level when supply elasticity decreases. For a low availability level, it is better to increase future supply by increasing production rather than by increasing stock, because building stocks are expensive. And the lower the kink, the higher will be the number of iterations needed. This higher number of iterations is not compensated for by the fact that iterations are simpler except for VFI and EGM. In the case of EGM, solving the model without supply reaction requires only simple calculations, and there is no nonlinear problem to solve. It is very fast. It is the case where this method shows all its interest.

The value function iteration, as expected, is the most time-consuming method. It converges very slowly (several hundred iterations) to rational expectations. EGM is the fastest method. The most precise method, parameterised expectations algorithm, is the slowest for a given approximation order (20 breakpoints). However, this is not to imply that it is slower than other methods for a given precision.

In terms of speed, EGM outperforms other methods by a factor of at least 20 in the case of inelastic supply. The speed of this method makes it a particularly suitable alternative to Deaton and Laroque's fixed-point

Table 6. Time (seconds) to solve the rational expectations equilibrium with a 20-breakpoint spline approximation

Method	With nonnegativity constraint		With convenience yield
	Inelastic supply ^a	Elastic supply	
Value function iteration	3.45	3.62	3.47
Parameterised expectations	0.78	0.44	0.48
Expectations function approximation			
Time iteration	0.42	0.36	0.22
Endogenous grid method	0.02	0.26	0.14
Decision rules approximation	0.55	0.47	0.24

^a For inelastic supply, Equation (4) is excluded of the model, and planned production is fixed at 1, its deterministic steady state.

algorithm, for econometric work. When extended to a situation with elastic supply, EGM loses part of its interest. It requires a nonlinear solve and so loses its very high speed but still outperforms other methods.

The model with convenience yield is smooth. Its functions are simpler to approximate than those of the model with nonnegativity constraint. Hence, it takes less time to solve it; convergence is achieved with fewer iterations.

The time needed to find the rational expectations equilibrium and to simulate the model is only a small part of the time devoted to the model. The time needed to program the solution method is likely the largest part of the whole. In this regard, EGM is very simple to implement in the model with inelastic supply, because it boils down to iterations with arithmetic operations. With elastic supply, it requires a nonlinear solve, so it is comparable with other methods. The various projection methods present the same difficulty of implementation, PEA being less intuitive because of its two nested loops, the first on storage and the second on availability.

5 Conclusion

This paper compared several solution methods for the competitive storage model. Three methods are popular in the related literature: value function iteration, parameterised expectations algorithm, and price function approximation. We introduced three other methods, borrowed from the literature on numerical analysis: perturbation, endogenous grid, and decision rules approximation. Perturbation methods proved to be of limited use in this context, because the storage model, smoothed by the convenience yield hypothesis, is too nonlinear, and the disturbances drive approximation out of its radius of convergence. The endogenous grid method is a good alternative to Deaton and Laroque's (1992) algorithm for a fast solve of the model for estimation purpose. One method stands out: the parameterised expectations algorithm, which is very precise, even at small approximation orders, because it approximates smooth and close to linear functions. These results could be generalised to the consumption/saving problem under uncertainty, where the equations are close to the storage model.

Because of modern computational power, any method, except perturbation, would provide a satisfying solution to the competitive storage model and within reasonable time. Numerical challenges arise when

the model is extended to include several state variables, such as in storage-trade interaction or the multiple-commodity model. In these settings, the advantages of the parameterised expectations algorithm over other methods may be even greater, because it is more precise, even for a small number of nodes.

References

- Aruoba, S. B., Fernández-Villaverde, J. and Rubio-Ramírez, J. F. (2006). Comparing solution methods for dynamic equilibrium economies. *Journal of Economic Dynamics and Control*, 30(12), 2477–2508.
- Barillas, F. and Fernández-Villaverde, J. (2007). A generalization of the endogenous grid method. *Journal of Economic Dynamics and Control*, 31(8), 2698–2712.
- Barsky, R. B., Mankiw, N. G. and Zeldes, S. P. (1986). Ricardian consumers with Keynesian propensities. *The American Economic Review*, 76(4), 676–691.
- Cafiero, C., Bobenrieth, E. S. A., Bobenrieth, J. R. A. and Wright, B. D. (2011). The empirical relevance of the competitive storage model. *Journal of Econometrics*, 162(1), 44–54.
- Cafiero, C. and Wright, B. D. (2006). Is the storage model a 'closed' empirical issue? The empirical ability of the storage model to explain price dynamics. In A. Sarris and D. Hallam (eds.) *Agricultural Commodity Markets and Trade. New Approaches to Analyzing Market Structure and Instability*, chapter 4, (pp. 89–114). Northampton: FAO/Edward Elgar.
- Carroll, C. D. (2001). A theory of the consumption function, with and without liquidity constraints. *The Journal of Economic Perspectives*, 15(3), 23–45.
- Carroll, C. D. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters*, 91(3), 312–320.
- Carter, C. A. and Revoredo Giha, C. L. (2007). The Working curve and commodity storage under backwardation. *American Journal of Agricultural Economics*, 89(4), 864–872.
- Chambers, M. J. and Bailey, R. E. (1996). A theory of commodity price fluctuations. *The Journal of Political Economy*, 104(5), 924–957.
- Christiano, L. J. and Fisher, J. D. M. (2000). Algorithms for solving dynamic models with occasionally binding constraints. *Journal of Economic Dynamics and Control*, 24(8), 1179–1232.
- Coourdacier, N., Rey, H. and Winant, P. (2011). The risky steady state. *American Economic Review*, 101(3), 398–401, papers and Proceedings of the One Hundred and Twenty-Third Annual Meeting of the American Economic Association, Denver, CO, January 6–9, 2011.
- Deaton, A. (1991). Saving and liquidity constraints. *Econometrica*, 59(5), 1221–1248.
- Deaton, A. and Laroque, G. (1992). On the behaviour of commodity prices. *Review of Economic Studies*, 59(1), 1–23.
- Deaton, A. and Laroque, G. (1995). Estimating a nonlinear rational expectations commodity price model with unobservable state variables. *Journal of Applied Econometrics*, 10, S9–S40.

- Deaton, A. and Laroque, G. (1996). Competitive storage and commodity price dynamics. *The Journal of Political Economy*, 104(5), 896–923.
- Dirkse, S. P. and Ferris, M. C. (1995). The PATH solver: A non-monotone stabilization scheme for mixed complementarity problems. *Optimization Methods and Software*, 5, 123–156.
- Fackler, P. L. (2005). A MATLAB solver for nonlinear rational expectations models. *Computational Economics*, 26(2), 173–181.
- Feigenbaum, J. (2005). Second-, third-, and higher-order consumption functions: a precautionary tale. *Journal of Economic Dynamics and Control*, 29(8), 1385–1425.
- Gardner, B. L. (1979). *Optimal Stockpiling of Grain*. Lexington, Mass.: Lexington Books.
- Gustafson, R. L. (1958). *Carryover Levels for Grains: A Method for Determining Amounts that are Optimal Under Specified Conditions*. Technical Bulletin 1178, US Dept. of Agriculture.
- den Haan, W. J. and Marcet, A. (1990). Solving the stochastic growth model by parameterizing expectations. *Journal of Business & Economic Statistics*, 8(1), 31–34.
- den Haan, W. J. and de Wind, J. (2009). *How well-behaved are higher-order perturbation solutions?* DNB Working Paper 240, De Nederlandsche Bank.
- Heer, B. and Maußner, A. (2008). Computation of business cycle models: A comparison of numerical methods. *Macroeconomic Dynamics*, 12(05), 641–663.
- Judd, K. L. (1992). Projection methods for solving aggregate growth models. *Journal of Economic Theory*, 58(2), 410–452.
- Judd, K. L. (1998). *Numerical Methods in Economics*. Cambridge: MIT Press.
- Judd, K. L., Maliar, L. and Maliar, S. (2011). Numerically stable and accurate stochastic simulation approaches for solving dynamic economic models. *Quantitative Economics*, 2(2), 173–210.
- Kaldor, N. (1939). Speculation and economic stability. *The Review of Economic Studies*, 7(1), 1–27.
- Kanzow, C. and Petra, S. (2004). On a semismooth least squares formulation of complementarity problems with gap reduction. *Optimization Methods and Software*, 19(5), 507–525.
- Kanzow, C. and Petra, S. (2007). Projected filter trust region methods for a semismooth least squares formulation of mixed complementarity problems. *Optimization Methods and Software*, 22, 713–735.
- Kim, S. H., Kollmann, R. and Kim, J. (2010). Solving the incomplete market model with aggregate uncertainty using a perturbation method. *Journal of Economic Dynamics and Control*, 34(1), 50–58.
- Krusell, P. and Smith, A. A., Jr. (2006). Quantitative macroeconomic models with heterogeneous agents. In R. Blundell, W. Newey and T. Persson (eds.) *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress*, Econometric Society Monographs, chapter 8, (pp. 298–340). New York: Cambridge University Press.
- Lence, S. H. and Hayes, D. J. (2002). U.S. farm policy and the volatility of commodity prices and farm revenues. *American Journal of Agricultural Economics*, 84(2), 335–351.
- Michaelides, A. and Ng, S. (2000). Estimating the rational expectations model of speculative storage: A

- Monte Carlo comparison of three simulation estimators. *Journal of Econometrics*, 96(2), 231–266.
- Miranda, M. J. (1997). Numerical strategies for solving the nonlinear rational expectations commodity market model. *Computational Economics*, 11(1–2), 71–87.
- Miranda, M. J. and Fackler, P. L. (2002). *Applied Computational Economics and Finance*. Cambridge: MIT Press.
- Miranda, M. J. and Glauber, J. W. (1995). Solving stochastic models of competitive storage and trade by Chebychev collocation methods. *Agricultural and Resource Economics Review*, 24(1), 70–77.
- Miranda, M. J. and Helmlinger, P. G. (1988). The effects of commodity price stabilization programs. *The American Economic Review*, 78(1), 46–58.
- Newbery, D. M. G. and Stiglitz, J. E. (1982). Optimal commodity stock-piling rules. *Oxford Economic Papers*, 34(3), 403–427.
- Osborne, T. (2004). Market news in commodity price theory: Application to the Ethiopian grain market. *The Review of Economic Studies*, 71(1), 133–164.
- Park, A. (2006). Risk and household grain management in developing countries. *Economic Journal*, 116(514), 1088–1115.
- Plato, G. and Gordon, D. (1983). Dynamic programming and the economics of optimal grain storage. *Journal of Agricultural Economics Research*, 35(1), 10–22.
- Preston, B. and Roca, M. (2007). *Incomplete Markets, Heterogeneity and Macroeconomic Dynamics*. Working Paper 13260, National Bureau of Economic Research.
- Rust, J. (1996). Numerical dynamic programming in economics. In H. M. Amman, D. A. Kendrick and J. Rust (eds.) *Handbook of Computational Economics*, volume 1, chapter 14, (pp. 619–729). Amsterdam: Elsevier.
- Scheinkman, J. A. and Schechtman, J. (1983). A simple competitive model with production and storage. *The Review of Economic Studies*, 50(3), 427–441.
- Williams, J. C. and Wright, B. D. (1991). *Storage and Commodity Markets*. New York: Cambridge University Press.
- Winschel, V. and Krätzig, M. (2010). Solving, estimating, and selecting nonlinear dynamic models without the curse of dimensionality. *Econometrica*, 78(2), 803–821.
- Working, H. (1948). Theory of the inverse carrying charge in futures markets. *Journal of Farm Economics*, 30(1), 1–28.
- Wright, B. D. (2001). Storage and price stabilization. In B. L. Gardner and G. C. Rausser (eds.) *Marketing, Distribution and Consumers*, volume 1B, part 2 of *Handbook of Agricultural Economics*, chapter 14, (pp. 817–861). Amsterdam: Elsevier.
- Wright, B. D. and Williams, J. C. (1982). The economic role of commodity storage. *The Economic Journal*, 92(367), 596–614.
- Wright, B. D. and Williams, J. C. (1984). The welfare effects of the introduction of storage. *The Quarterly Journal of Economics*, 99(1), 169–192.

Zeldes, S. P. (1989). Optimal consumption with stochastic income: Deviations from certainty equivalence. *The Quarterly Journal of Economics*, 104(2), 275–298.