Learning and Price Volatility in Duopoly Models of Resource Depletion*

Martin Ellison
University of Oxford

Andrew Scott
London Business School and CEPR

August 2, 2010

Abstract

We introduce learning into a Hotelling model of a non-renewable resource market. By combining learning and scarcity we add significantly to the dynamics implied by learning and substantially enhance the volatility of commodity prices. In our learning model we show how a self confirming equilibrium exists but is not constant over time. As scarcity increases the SCE shifts from a non-cooperative rational expectations equilibrium to a cooperative rational expectations outcome. As a result prices trend at a rate faster that the rate of time preference. We show the existence of escape dynamics which generate substantial volatility in commodity prices despite the fact the model is subject only to i.i.d shocks. The shifting SCE significantly alters escape dynamics with the time to escape shortening and the magnitude of dynamics reducing as scarcity rises. In terms of the Hotelling model, a shifting SCE and variable escape dynamics introduces greater volatility at low frequencies and substantially larger cyclical volatility. These price fluctuations show sharp upward breaks in price and non-linear, non-stationary and asymmetric price fluctuations. We show these results are robust to a range of extensions, including extractions costs, stochastic shifts in demand and learning assumptions closer to rational expectations.

Keywords: Commodity Prices, Escape Dynamics, Hotelling, Learning, Scarcity, Self Confirming Equilibria

JEL classification: D43, D83, Q31

---

*For useful comments we thank Klaus Adam, Chryssi Giannitsarou, Albert Marcet, Tom Sargent, Ulf Soderstrom, Noah Williams and seminar participants at Amsterdam, Bocconi, Cambridge, ESSIM, London Business School, London School of Economics, Mannheim, Oxford, Pavia, the Royal Economic Society Conference and the Swiss National Bank.

1 Introduction

A number of recent papers have explored the potential for learning and associated changes in agents' beliefs to create interesting economic dynamics. Marcet and Sargent (1989) focus on whether learning dynamics converge on a Rational Expectations Equilibria. Evans and Honkapohja (1995) investigate issues of E-stability and how learning can generate stable or unstable fluctuations. Bullard and Mitra (2002) show the relevance of this issue for the evaluation of monetary policy rules. In a series of seminal papers Fudenberg and Levine (1993a and b, 2009) analyse the concept of a Self Confirming Equilibrium and its relationship to Nash (and Rational Expectations) Equilibria. The concept of a Self Confirming Equilibrium and the possibility of escape dynamics that arise are examined by Sargent (1999), Cho, Williams and Sargent (2002) and Sargent, Williams and Zha (2009) who consider cases of inflation and hyperinflation in the Americas. This paper extends this literature by introducing learning into a new setting – non-renewable commodity markets. Non-renewable commodities introduce issues of scarcity into the analysis and we show how this fundamentally affects the nature of learning dynamics. As a result of the interaction of scarcity and learning our model adds substantially to the volatility at both low and high frequencies of non-renewable commodity prices.

The application of learning to commodity markets is a novel one. Most of the papers cited above relate to inflation and monetary policy issues. Lettau and Uhlig (1999) consider the macroeconomic implications of learning for consumption fluctuations, Adam, Marcet and Nicolini (2007) introduce learning into the analysis of stock markets and Piazzesi and Schneider (2007) show how learning can add important and empirically plausible dynamics to yield curve fluctuations. However as noted by Deaton and Laroque (1996) commodity prices show a persistence and volatility that is hard to explain with reference to underlying shocks, making them a natural candidate with which to explore the propagation role of learning in price dynamics. Williams (2001) offers a general duopoly model of a product market and shows how escape dynamics provide additional volatility.[1] This paper uses Williams' model as a starting point but critically adds the assumption that the product being supplied is a finite resource.

The extension of Williams' model to the case of non-renewable and depletable resources is important. Firstly, unlike agricultural commodity markets, nonrenewable commodities such as oil, copper, platinum and tin, are by definition characterised by a finite stock and as a consequence these commodities are often found in the control of a small number of suppliers, an obvious case being OPEC. Reflecting this the seminal work on nonrenewable commodities has always involved firms with market power (e.g. Hotelling (1931), Dasgupta and Heal (1974), Salant (1976), Stiglitz and Dasgupta (1982), Loury (1986)). The fact that firms have market power is critical if learning and escape dynamics are to generate important fluctuations. The second important feature of non-renewable commodities is it introduces a role for scarcity. We show how scarcity fundamentally impacts on market dynamics.

Hotelling (1931) shows how nonrenewable commodities should optimally see their price trend upwards over time at the rate of time preference. We show in the context of our model and no learning that there are two Rational Expectations Equilibria. In the case where there is no scarcity we show how producers set a low price and a high quantity and the market is in a noncooperative equilibrium. In the case where scarcity binds, producers set a high price and a low quantity and the market is in a cooperative equilibrium. In a nonrenewable commodity setting as firms produce their perceptions of scarcity will inevitably rise and the market shifts from a noncooperative to a cooperative equilibrium. We show in the context of our learning model that this occurs because of a shift in the self confirming equilibrium and we characterize the mean dynamics and e-stability of this process. Although agents do not act strategically we show how scarcity acts as a coordinating device. Growing scarcity causes prices to rise for familiar Hotelling reasons. At the same time each producer is cutting back

[1] A more recent version of Williams (2001) was posted on Noah William's website in 2009, but the new version does not discuss the duopoly model of a product market.

their production. This encourages producers to lower their estimate of the elasticity of demand and so they continue to further cut production. Prices continue to rise due to both scarcity and the scaling back of production. Therefore scarcity coordinates these uncoordinated firms so that the market tends towards the cooperative equilibrium. Along the mean dynamics the market shows escape dynamics – these in our model take the form of shifts from the mean dynamics towards the cooperative equilibrium. When scarcity is not apparent these escape dynamics show large abrupt jumps in prices but happen infrequently. As scarcity binds more and more the escape dynamics become more frequent although less dramatic in nature – the mean dynamics become increasingly close to the cooperative equilibrium making escape dynamics less important.

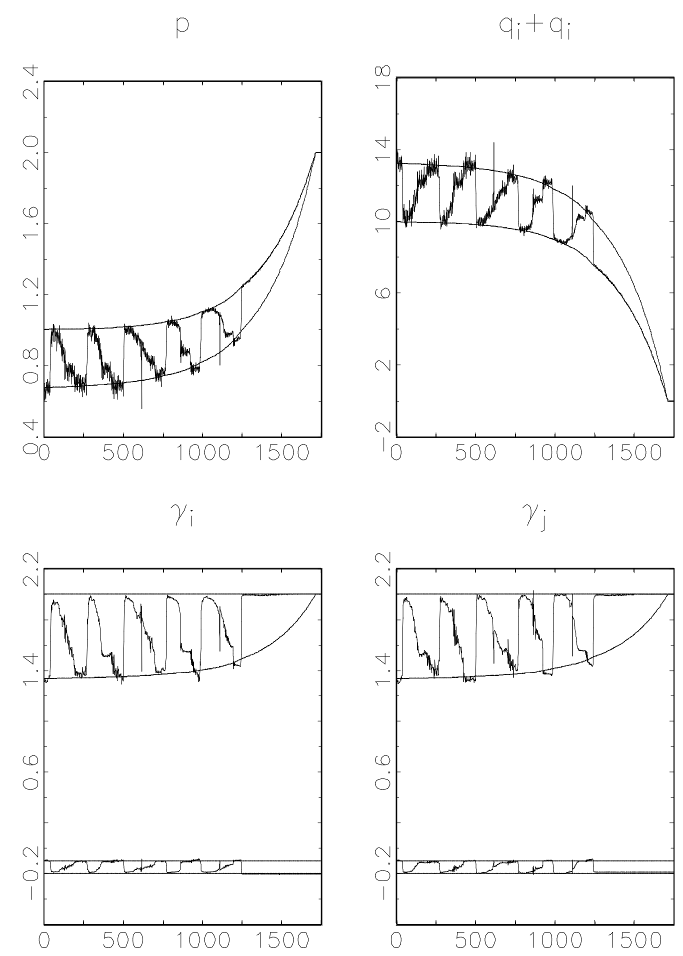

Figure 1: Simulated time paths of price, supply and beliefs

As an indication of the importance of learning see Figure 1 (which shows simulated time paths for prices p, quantities $q_i + q_j$ and beliefs γ_i, γ_j and is explained in detail in Section 5.3). The Hotelling trend in commodity prices is readily apparent but so too is the substantial extra volatility generated by the escape dynamics. Commodity prices show sharp and aggressive upward increases and then later corrections. The upward price surges coincide with agents believing the commodity is more scarce than is the case and the market shifting to the cooperative equilibria.

The notion that learning can produce sharp fluctuations characterized by regime changes has been explored by others. For instance, Evans and Honkapohja (1993, 1994) focus on models of multiple equilibria and how business cycle fluctuations may be caused by shifts between these equilibria. A similar notion is explored by Kasa (2004) who shows how escape dynamics can generate sharp changes in exchange rates and models Latin American currencies using a Markov switching regime. However these are models of multiple equilibria whereas our model always has a unique equilibrium. Although scarcity leads to a shift in the Self Confirming Equilibrium for any given level of scarcity our model implies a unique equilibrium. Brock and Hommes (1997) also show how learning leads to substantial fluctuations by considering a cobweb model where agents decide optimally which of two forecasting rules to utilize. This endogenously leads to cycles as agents are induced to mainly use one forecasting rule and then switch to the alternative. Our model differs from theirs in that all our agents use the same forecasting rule and our fluctuations are driven by underlying changes in scarcity. As scarcity will eventually bind completely the long run properties of our model are well defined and do not show strange dynamics[2]. Referring not to the learning literature but to work in IO the key mechanism underlying our commodity price volatility is the fluctuations between a cooperative and a non-cooperative equilibria. This is exactly similar to the mechanism in

[2] Kandori, Mallaith and Rob (1993) show in the context of an evolutionary game a result with a similar flavour to ours. They show how mutations can restrict the long run equilibria of the model and create shift in outcomes. In our case scarcity means that only the cooperative equilibria can be a long run outcome of the model.

Green and Porter (1984) although their fluctuations are driven by strategic interactions and retaliation strategies and have no role for learning. We show by contrast in a model absent of strategic interactions how learning alone can generate the same market fluctuations.

The plan of the paper is as follows. In Section 2 we outline our basic model but abstract from learning so the key features are only imperfect competition and the non-renewable nature of the resource. We characterise the two Rational Expectations Equilibria of our model : a non-cooperative equilibrium where each producer maximises their profits by setting output assuming their competitor's production is exogenous and a cooperative equilibrium where producers maximise joint profit by coordinating their production plans. In Section 3 we move away from Rational Expectations and introduce learning. We assume producers know the market is characterised by a linear demand curve but do not know the elasticity of demand and so each period update their beliefs before choosing their output. We derive the firm's optimal output decision and show how it depends on the producer's estimates of the demand curve and the perceived level of scarcity. In Section 4 we examine the implications of learning for market dynamics by deriving the mean dynamics and show that a self confirming equilibrium (SCE) exists. We show how in our model, in contrast to the standard result, the SCE evolves over time depending on the level of scarcity. In Section 5 we turn our attention to escape dynamics and show, via simulation, how the combination of an evolving SCE and escape dynamics produces considerable greater volatility at all frequencies in commodity prices relative to the sole i.i.d source of volatility in the model. As well as showing considerable volatility at all frequencies, prices also show substantial non-linear, non-stationary and asymmetric movements. In Sections 6 and 7 we pursue a number of important robustness issues. In Section 6 we review our learning assumption and let each firm assume its competitor's production follows a trend, including the case implied by the Hotelling model. In Section 7 we introduce a general class of extraction costs and again derive our main results, although obviously the existence of extraction costs changes the quantitative predictions of our model. We also show the same conclusion holds for stochastic shifts in demand. A final section concludes.

2 A simple duopoly model

In this section we outline our core model[3] of a depletable resource market characterised by duopoly and focus on two equilibria - a cooperative and non-cooperative outcome. We temporarily abstract from learning considerations in order to focus on the other key features of our model. We intend this as a general model rather than a specific attempt to match any particular commodity market. We therefore abstract from the various specific institutional factors that are critical to explain fluctuations in any one particular market. In later sections we add a number of additional features (extraction costs, stochastic demand, etc.) and show these merely change some of the quantitative features of our model rather than influence our main findings.

The price p_t of the homogenous depletable resource is assumed to depend on total market supply through a linear demand curve:

$$p_t = a - b(q_{it} + q_{jt}), \tag{1}$$

where $a > 0$ and $b > 0$ are structural parameters. The two suppliers i and j are assumed to have imperfect control over the quantity of depletable resource they each supply to the market. Each supplier sets intended supply $\hat{q}_{it}$ but makes a control error W_{it} so market supply q_{it} is given by:

$$q_{it} = \hat{q}_{it} + W_{it}, \tag{2}$$

where W_{it} is i.i.d. with mean zero and standard deviation σ. Control errors are uncorrelated across suppliers and have sufficiently bounded support to ensure that supply is always positive. In the context of the oil industry a natural interpretation of these control errors would be issues concerning pipeline control or variations in delivery times.

We denote the estimate of firm i's total possible future supply at time t by Q_{it} and so firm i faces the additional constraint $\sum_{k=0}^{T_i} q_{it+k} = Q_{it}$ where T_i is the time to exhaustion.

[3] Our core model is based on an example of a standard goods market, without resource depletion, in Williams (2001).

We assume the firm faces no extraction costs (a restriction we remove in a later section) and seeks to maximise the discounted sum of future profits which it discounts at the rate β where $0 < \beta < 1$. We consider two possible objectives for firms (see Hansen, Epple and Roberds (1985) for a fuller treatment of a more generic version of our model and the broad range of equilibrium it supports). The first case is when each firm chooses its own output and time to exhaustion so as to maximise their profits taking the other firms output decisions as exogenous and with full knowledge of the parameters of the model. We shall refer to this as the non-cooperative Rational Expectations equilibrium or NCE. In this case our model is:

$$\max_{\hat{q}_{it},T_i} E_t \sum_{k=0}^{T_i} \beta^k p_{t+k} q_{it+k}$$

$$s.t.$$

$$p_{t+k} = a - b(q_{it+k} + q_{jt+k}),$$

$$q_{it+k} = \hat{q}_{it+k} + W_{it+k},$$

$$E_t \sum_{k=0}^{T_i} q_{it+k} = Q_{it}.$$

The first order condition is a reaction function for firm $i's$ supply conditional on firm j supply:

$$\hat{q}_{it+k} = \frac{a - b\hat{q}_{jt+k} - \lambda_{it}\beta^{-k}}{2b}, \tag{3}$$

where λ_i is the multiplier attached to the finite resource constraint. By analogy, the reaction function for firm $j's$ supply conditional on firm i supply is:

$$\hat{q}_{jt+k} = \frac{a - b\hat{q}_{it+k} - \lambda_{jt}\beta^{-k}}{2b}. \tag{4}$$

Equations (3) and (4) simultaneously determine market supply and combine with the linear demand curve (1) to solve for the market price. Under the assumption that firms are symmetric $(q_i = q_j, \lambda_i = \lambda_j, T_i = T_j)$ then the NCE is given by:

$$\begin{aligned} \hat{q}_{it} &= \hat{q}_{jt} = \frac{a - \lambda_t}{3b} \\ p_t &= \frac{a}{3} + \frac{2\lambda_t}{3} - b(W_{it} + W_{jt}) \end{aligned}$$

Our second equilibrium is a cooperative rational expectations equilibrium (CE) in which both duopolists coordinate and choose their output jointly so as to maximise their combined future profits. That is:

$$\max_{\hat{q}_{it}+\hat{q}_{jt},T} E_t \sum_{k=0}^{T} \beta^k p_{t+k}(q_{it+k} + q_{jt+k})$$

$$s.t.$$

$$p_{t+k} = a - b(q_{it+k} + q_{jt+k}),$$

$$q_{it+k} = \hat{q}_{it+k} + W_{it+k},$$

$$q_{jt+k} = \hat{q}_{jt+k} + W_{jt+k},$$

$$E_t \sum_{k=0}^{T} (q_{it+k} + q_{jt+k}) = Q_{it} + Q_{jt}.$$

The first order condition determines total supply to the market:

$$\hat{q}_{it+k} + \hat{q}_{jt+k} = \frac{a - \lambda_t \beta^{-k}}{2b},$$

so in the simple case where the suppliers are symmetric we obtain:

$$\begin{aligned} \hat{q}_{it} &= \hat{q}_{jt} = \frac{a - \lambda_t}{4b}, \\ p_t &= \frac{a}{2} + \frac{\lambda_t}{2} - b(W_{it} + W_{jt}), \end{aligned}$$

and under the CE firms set their output lower and the market price higher relative to the NCE.

The above focuses on how the level of output and price depends on the degree of cooperation between the two producers. However the intertemporal behaviour of output and prices is also important for our analysis. Consider once again the non-cooperative case. We know that at $t + T$ the output of each producer, assuming symmetry, is zero and so:

$$\beta^T a = \lambda_t. \tag{5}$$

Further summing the first order condition (3) over all periods gives:

$$Ta - 2bQ_{it} - bQ_{jt} = \lambda_t \frac{1 - \beta^{-T}}{1 - \beta^{-1}}. \tag{6}$$

These two equations can be used to jointly solve for T and λ_t. Using a continuous time representation of our system we can derive an implicit expression for the Lagrange Multiplier λ_t:

$$a(1 + \log \lambda_t - \log a) - 3bQ \log \beta = \lambda,$$

where we have invoked complete symmetry so that $Q_{it} = Q_{jt} = Q$. Totally differentiating and using the fact that $\dot{Q}_{it} = -q_{it}$ leads to the standard Hotelling result:

$$\frac{\dot{\lambda}_i}{\lambda_i} = -\log \beta \tag{7}$$

In other words the Lagrange Multiplier rises at the rate of time preference. An exactly similar approach yields an identical expression in the case of the cooperative model. Under both cases prices and output have their trend driven by the Hotelling factor $\log \beta$ but differ in the level of output and prices across the two scenarios. In the case of the non-cooperative model, output is always higher and prices always lower at every moment compared to the cooperative case.

3 Introducing learning

The previous section assumed suppliers know market price is determined by the linear demand curve (1), that they can observe their own output and know the current output of their competitor, and they know with certainty the structural parameters a and b. In this section we drop the last two assumptions and assume agents know only that demand is linear and they can only observe their own output, opening up the possibility of learning. The models outlined in the previous section only have two sources of fluctuations - the white noise control errors W_t and the trend dynamics of the Hotelling component λ_t. By introducing learning we show how commodity price fluctuations can be much more complex. The linear structure of our model and its emphasis on duopoly is the same as that in Williams (2001), enabling us to replicate his approach. The key difference is the fact that ours is a market for a depletable commodity.

3.1 An approximating model

A producer wishing to learn the true value of a and b has to impute the supply of its competitors. In the market, supply arises as the solution to a dynamic optimisation problem constrained by perceptions of market power and scarcity as a result of which it is difficult for a firm to deduce the supply of its competitor. We therefore permit a minor departure from full rationality and allow each supplier to use an approximating model. The approximating model we assume is one where each supplier believes the supply of its competitor has had no systematic variation in the *recent* past. This is a strong assumption[4] which we modify in Section 6 by allowing perceptions of competitor output to follow a trend. However even the assumption that the competitor's output is unchanged is not as strong as it might at first appear. We show, for instance, that this assumption leads to a self-confirming equilibrium so that our approximating model is correctly specified on the equilibrium path and only misspecified out of self-confirming equilibrium.

Using this approximating model supplier i sees recent market data as generated by the process:

$$p_t = (a - bq_j + \eta_{it}) - bq_{it}, \tag{8}$$

in which case recent data can be used by the agent to estimate a regression of the form:

$$p_t = \gamma_{it}^0 + \gamma_{it}^1 q_{it} + \eta_{it}, \tag{9}$$

and thereby obtain estimates of $(a - bq_j)$ and $-b$. The presence of the residual η_{it} represents an acknowledgement by the producer that they only have an approximating model. The perception of no systematic variation in competitor supply means η_{it} is seen as orthogonal to q_j and standard econometric techniques produce estimates $(\hat{\gamma}_{it}^0\ \hat{\gamma}_{it}^1)$ that are unbiased predictors of $(a - bq_j)$ and $-b$. In response to the perception that (9) only holds in recent data, we assume

[4] Although one that is also made *inter alia* in Salant (1976), Porter (1983) and Green and Porter (1984).

each supplier estimates this regression equation using discounted least squares:

$$\hat{\gamma}_{it+1} = \hat{\gamma}_{it} + \varepsilon R_{it}^{-1} X_{it}' \left(p_t - X_{it} \hat{\gamma}_{it} \right),$$
$$R_{it+1} = R_{it} + \varepsilon \left(X_{it}' X_{it} - R_{it} \right),$$

where $X_{it} = (1 \; q_{it})$, $\hat{\gamma}_{it} = (\hat{\gamma}_{it}^0 \; \hat{\gamma}_{it}^1)'$ and ε is the rate at which data is discounted. Allowing for the producer to discount data reflects the agent's belief that it is only in the recent past that their competitor's output has remained unchanged. These standard recursive updating equations can be interpreted as describing the constant gain learning algorithm suppliers use to form perceptions of the economic environment in which they operate. R is a matrix of estimated second moments.

3.2 Supplier decisions

The decision faced by a supplier is how much depletable resource to supply to the market now, given their current knowledge of market conditions and expectations of how market power and scarcity will develop in the future. Expectations of future market power ultimately depend on what suppliers expect their competitors will do and how much confidence they have in their estimates of the structural parameters of the model. We operationalise this aspect by assuming that the supplier maximises *anticipated* rather than expected profits, in the sense suggested by Kreps (1998). Under this decision criterion, the supplier projects forward by assuming that the degree of market power it holds is known with certainty and will remain unchanged in the future. The current perception of the demand curve is summarised by equation (9) and parameter estimates $(\hat{\gamma}_{it}^0 \; \hat{\gamma}_{it}^1)$ so the (approximately) optimal depletion plan solves the problem:

$$\max_{\hat{q}_{it},T_i} E_t \sum_{k=0}^{T_i} \beta^k p_{t+k} q_{it+k}$$

$$s.t.$$

$$p_{t+k} = \hat{\gamma}_{it}^0 + \hat{\gamma}_{it}^1 q_{it+k} + \eta_{it+k},$$

$$q_{it+k} = \hat{q}_{it+k} + W_{it+k},$$

$$E_t \sum_{k=0}^{T_i} q_{it+k} = Q_{it}.$$

We solve this by first forming the Lagrangian:

$$\mathcal{L} = \sum_{k=0}^{T_i} \beta^k p_{t+k} q_{it+k} - \lambda_{it} \left[\sum_{k=0}^{T_i} q_{it+k} - Q_{it} \right].$$

The first order condition with respect to $\hat{q}_{it+k}$ is:

$$\beta^k (\hat{\gamma}_{it}^0 + 2\hat{\gamma}_{it}^1 q_{it+k}) = \lambda_{it}, \tag{10}$$

from which it follows that supply in the first period is given by:

$$\hat{q}_{it} = \frac{\lambda_{it} - \hat{\gamma}_{it}^0}{2\hat{\gamma}_{it}^1}. \tag{11}$$

The value of the Lagrange multiplier λ_{it} can be calculated by defining T_{it} as the time to exhaustion of the depletable resource. At the time of exhaustion, optimality requires supply to be zero so first order condition (10) reduces to:

$$\beta^{T_i} \hat{\gamma}_{it}^0 = \lambda_{it}, \tag{12}$$

which defines a relationship between T_i and λ_{it}[5]. In similar fashion, optimality requires the stock of the depletable resource to be zero at the time of exhaustion. Summing the first order condition (10) between t and $t + T_i$ then gives a second relationship between T_i and λ_{it}:

$$T_i \hat{\gamma}_{it}^0 + 2\hat{\gamma}_{it}^1 Q_{it} = \lambda_{it} \frac{1 - \beta^{-T_i}}{1 - \beta^{-1}} \tag{13}$$

Equations (12) and (13) simultaneously determine the time to exhaustion and the Lagrange multiplier, the latter determining market supply through equation (11).

[5] Given these results the elasticity of perceived demand, η, $= \frac{\lambda + \gamma_0}{\lambda - \gamma_0} = \frac{\beta^T + 1}{\beta^T - 1}$. As we get close to exhaustion $T \to 0$ this elasticity is rising, a product of our linear formulation of the demand curve.

4 Self-Confirming Equilibrium

A key question to ask of the learning process in our model is whether it will guide suppliers to a self confirming equilibria (SCE), as discussed by Fudenberg and Levine (1993a and b, 2009) and Sargent (1999). In a self-confirming equilibrium, each player's strategy is a best response given their beliefs about their opponents' actions, and each player's beliefs are correct along the equilibrium path. In a self-confirming equilibrium, no player ever observes outcomes that contradict their beliefs, even though beliefs about out of equilibrium actions need not be correct. We show in this section that our model does indeed tend to a well defined SCE. However a novel feature of our model is that the SCE, although well defined, is not constant but changes over time. In particular the SCE depends on the degree of resource scarcity, λ. We follow Williams (2001) and Cho, Sargent and Williams (2002) and derive these results regarding the SCE using stochastic approximation techniques to analyse the *mean dynamics* of the continuous time analogue of our model, thereby tracing out the expected evolution and limit point of supplier beliefs.

4.1 Mean dynamics

The mean dynamics of beliefs can be derived by re-writing the recursive updating scheme as:

$$\frac{\hat{\gamma}_{it+1} - \hat{\gamma}_{it}}{\varepsilon} = R_{it}^{-1} X_{it}' \left(p_t - X_{it} \gamma_{it} \right), \tag{14}$$

$$\frac{R_{it+1} - R_{it}}{\varepsilon} = \left(X_{it}' X_{it} - R_{it} \right). \tag{15}$$

Equations (14) and (15) describe a discrete-time approximation of a continuous time process perturbed by shocks. Taking the limit as $\varepsilon \to 0$, the approximation error tends to zero and a weak law of large numbers ensures that the stochastic element becomes negligible. In the limit, the mean dynamics of beliefs can therefore be described by a pair of ordinary differential equations:

$$\dot{\gamma}_i = R_i^{-1} X_i' \left(p - X_i \gamma_i\right),$$
$$\dot{R}_i = X_i' X_i - R_i.$$

Expressions for p and $X_i = (1\ q_i)$ can be obtained from equations (1) and (2) determining market price, and equation (11) for market supply. After taking expectations, the mean dynamics of beliefs are given by:

$$\dot{\gamma}_i = R_i^{-1} \begin{pmatrix} a - b\left(\frac{\lambda_i - \gamma_i^0}{2\gamma_i^1} + \frac{\lambda_j - \gamma_j^0}{2\gamma_j^1}\right) - \frac{\gamma_i^0}{2} - \frac{\lambda i}{2} \\ \left(\frac{\lambda_i - \gamma_i^0}{2\gamma_i^1}\right)\left[a - b\left(\frac{\lambda_i - \gamma_i^0}{2\gamma_i^1} + \frac{\lambda_j - \gamma_j^0}{2\gamma_j^1}\right) - \frac{\gamma_i^0}{2} - \frac{\lambda_i}{2}\right] - (b + \gamma_i^1)\sigma \end{pmatrix}, \tag{16}$$

$$\dot{R}_i = \begin{pmatrix} 1 & \frac{\lambda_i - \gamma_i^0}{2\gamma_i^1} \\ \frac{\lambda_i - \gamma_i^0}{2\gamma_i^1} & \left(\frac{\lambda_i - \gamma_i^0}{2\gamma_i^1}\right)^2 + \sigma \end{pmatrix} - R_i. \tag{17}$$

These mean dynamics show how beliefs evolve as a function of current beliefs and the value of λ, the Lagrange multiplier on the supplier's resource constraint. To complete the description of the mean dynamics we need to specify the dynamic evolution of λ. Using a similar approach to Section 2 and the continuous time analogues of (12) and (13) we can derive an implicit expression for the Lagrange multiplier:

$$\left(1 + \log \lambda_i - \log \gamma_i^0\right) \gamma_i^0 + 2\gamma_i^1 Q_i \log \beta = \lambda_i, \tag{18}$$

from which it follows by total differentiation that:

$$\left(\log \lambda_i - \log \gamma_i^0\right) \dot{\gamma}_i^0 + \gamma_i^0 \frac{\dot{\lambda}_i}{\lambda_i} + 2\gamma_i^1 \dot{Q}_i \log \beta + 2\dot{\gamma}_i^1 Q_i \log \beta = \dot{\lambda}_i.$$

Given that $\dot{Q}_i = -q_i$ and using (18) to substitute out for Q_i gives:

$$\frac{\dot{\lambda}_i}{\lambda_i} = \gamma_i^0 \left(\frac{\log \lambda_i - \log \gamma_i^0}{\lambda_i - \gamma_i^0}\right) \frac{\dot{\gamma}_i^0}{\gamma_i^0} + \left(1 - \gamma_i^0 \frac{\log \lambda_i - \log \gamma_i^0}{\lambda_i - \gamma_i^0}\right) \frac{\dot{\gamma}_i^1}{\gamma_i^1} - \log \beta. \tag{19}$$

Equations (16), (17) and (19) fully characterise the mean dynamics of our model. Equation (19) differs from our results without learning through the presence of the first two terms on the right hand side. These additional terms capture how learning about the demand curve impacts on an agent's assessment of the degree of scarcity. Equation (19) clearly shows the potential for learning to add greater volatility to commodities in finite supply. Without learning the multiplier grows at a constant rate $-\log\beta$ but with learning the fact agents update their beliefs provides additional volatility. As the market gets close to exhaustion we have $\lambda_i \to \gamma_i^0$ (because $q_T^i = 0$) so asymptotically the original Hotelling trend of Section 2 reasserts itself but at points away from exhaustion learning provides extra volatility.

4.2 Self-Confirming Equilibrium without scarcity

Proposition 1 *When there is no resource scarcity, the self-confirming equilibrium is equivalent to the* ***non-cooperative*** *rational expectations equilibrium of the model.*

Proof. We look for a self-confirming equilibrium in which beliefs are stable so $\dot{\gamma}_i = \dot{\gamma}_j = 0$. When there is no resource scarcity, the Lagrange multipliers λ_i and λ_j are zero and simple algebraic manipulation of (16) and (17) shows the SCE has the form:

$$\begin{aligned} \gamma_i &= \gamma_j = \begin{pmatrix} \frac{2}{3}a \\ -b \end{pmatrix}, \\ R_i &= R_j = \begin{pmatrix} 1 & \frac{a}{3b} \\ \frac{a}{3b} & \left(\frac{a}{3b}\right)^2 + \sigma \end{pmatrix}. \end{aligned}$$

From the supply equation (11) and the market demand curve (1), it follows that:

$$\begin{aligned} q_i &= q_j = \frac{a}{3b}, \\ p &= \frac{a}{3}. \end{aligned}$$

Comparing these expressions with the NCE outcome of Section 2 in the model with no learning we can see that the SCE of the mean dynamics is the same as the price and quantity outcomes in the NCE when $\lambda = 0$, i.e. there is no resource scarcity. ■

4.3 Self-Confirming Equilibrium with scarcity

Proposition 2 *As the resource becomes infinitely scarce the self-confirming equilibrium is equivalent to the* ***cooperative*** *rational expectations equilibrium of the model.*

Proof. We begin the proof by asserting that a self-confirming equilibrium exists in which beliefs are stable so $\dot{\gamma}_i = \dot{\gamma}_j = 0$ in the limit as the depletable resource becomes scarce. From equation (19) it therefore follows that:

$$\frac{\dot{\lambda}_i}{\lambda_i} = -\log\beta,$$

and the standard Hotelling result applies. From (17) it follows that as $\lambda_i \to \infty$ i.e. the resource becomes infinitely scarce, we obtain well-defined limits for the (appropriately scaled) second moment matrix R_i:

$$\begin{aligned}
\lim_{\lambda_i\to\infty}\left(\frac{R_i^{12}}{\lambda_i}\right) &= -\frac{1}{2\gamma_i^1(\log\beta - 1)}, \\
\lim_{\lambda_i\to\infty}\left(\frac{R_i^{22}}{\lambda_i^2}\right) &= -\frac{1}{4(\gamma_i^1)^2(2\log\beta - 1)}.
\end{aligned}$$

Re-writing the first line of equation (16) in terms of R_i^{22}/λ_i^2 and R_i^{12}/λ_i gives:

$$\dot{\gamma}_i^0 = \frac{1}{(R_i^{22}/\lambda_i^2) - (R_i^{12}/\lambda_i)^2}\left[\begin{array}{c} \left(\frac{R_i^{22}}{\lambda_i^2}\right)\left(a - b\left(\frac{\lambda_i-\gamma_i^0}{2\gamma_i^1} + \frac{\lambda_j-\gamma_j^0}{2\gamma_j^1}\right) - \frac{\gamma_i^0}{2} - \frac{\lambda i}{2}\right) \\ -\left(\frac{R_i^{12}}{\lambda_i}\right)\left(\left(\frac{\lambda_i-\gamma_i^0}{2\gamma_i^1\lambda_i}\right)\left[a - b\left(\frac{\lambda_i-\gamma_i^0}{2\gamma_i^1} + \frac{\lambda_j-\gamma_j^0}{2\gamma_j^1}\right) - \frac{\gamma_i^0}{2} - \frac{\lambda_i}{2}\right] - \frac{(b+\gamma_i^1)\sigma}{\lambda_i}\right)\end{array}\right],$$

which has a unique solution for $\dot{\gamma}_i = \dot{\gamma}_j = 0$ when $\lambda_i, \lambda_j \to \infty$:

$$\gamma_i = \gamma_j = \begin{pmatrix} a \\ -2b \end{pmatrix}.$$

The supply and price in self-confirming equilibrium are defined by:

$$\begin{aligned}
q_i &= \frac{a-\lambda_i}{4b}, \\
q_j &= \frac{a-\lambda_j}{4b}, \\
p &= \frac{a}{2} + \frac{\lambda_i+\lambda_j}{4}.
\end{aligned}$$

Assuming symmetry we have $q_i = q_j$ and $\lambda_i = \lambda_j$ and these equations yield the same expressions as for the CE outcome in Section 2 under no learning. ■

These results show that the SCE changes as the depletable resource becomes more scarce. As λ increases the SCE shifts from the high output and low price case of the non-cooperative rational expectations equilibrium towards the low output and high price of the cooperative equilibrium. The dependence on λ arises because of the interactions between scarcity and learning. Each agent estimates the elasticity of demand by assuming their competitor's output remains unchanged and only their own output is changing. However under scarcity the output of both producers is decreasing over time. Each producer therefore overestimates by a factor of 2 the responsiveness of prices to output and unilaterally decides to cut back production. Therefore scarcity acts as a coordinating device and unwittingly leads producers to the cooperative outcome. As a consequence the trend dynamics of output and prices are no longer driven purely by the Hotelling factor $\log \beta$ but by shifts in the SCE. It should be stressed that the shifts that occur in the SCE do not reflect issues of multiple equilibria, such as in the learning model of Evans and Honkapohja (1993, 1994), but changes in the underlying level of scarcity.

4.4 Stability of the SCE and the market at exhaustion

The fact the SCE changes in our model due to shifts in scarcity makes the issue of stability important. As scarcity increases the market shifts away from the SCE at the non-cooperative equilibria - in this sense the NCE is not stable. However we can show (see Appendix A and B) that both the SCE with no scarcity (the NCE) and with infinite scarcity (the CE) are *e-stable.* In other words, small perturbations of beliefs around these SCE are only temporary when evaluated at $\lambda = 0$ and $\lambda = \infty$ respectively. As a consequence once scarcity starts to increase and λ rises the cooperative SCE behaves like an attractor for the mean dynamics, adding additional dynamics to the Hotelling trend. Due to the depletable nature of the commodity an interesting feature of the model is that whilst the CE acts as an attractor the mean dynamics never actually reach the new SCE. The model can get arbitrarily close but beliefs never quite converge to a new SCE. In particular we have:

Proposition 3 *Beliefs are still evolving and tending towards the CE at the point of exhaustion and they never fully converge.*

Proof. From (16) and using the fact that at the point of exhaustion $q = 0$ so $\lambda = \gamma^0$, if the market is at the cooperative SCE then $\gamma^0 = a$ and we have:

$$\dot{\gamma}_i^0 = \frac{1}{(R_i^{22}/\lambda_i^2) - (R_i^{12}/\lambda_i)^2} \left[\begin{array}{c} \left(\frac{R_i^{22}}{\lambda_i^2}\right)\left(a - b\left(\frac{\lambda_i - \gamma_i^0}{2\gamma_i^1} + \frac{\lambda_j - \gamma_j^0}{2\gamma_j^1}\right) - \frac{\gamma_i^0}{2} - \frac{\lambda i}{2}\right) \\ -\left(\frac{R_i^{12}}{\lambda_i}\right)\left(\left(\frac{\lambda_i - \gamma_i^0}{2\gamma_i^1\lambda_i}\right)\left[a - b\left(\frac{\lambda_i - \gamma_i^0}{2\gamma_i^1} + \frac{\lambda_j - \gamma_j^0}{2\gamma_j^1}\right) - \frac{\gamma_i^0}{2} - \frac{\lambda_i}{2}\right] - \frac{(b+\gamma_i^1)\sigma}{\lambda_i}\right) \end{array} \right],$$

$$\dot{\gamma}_i^1 = \frac{1}{(R_i^{22}/\lambda_i^2) - (R_i^{12}/\lambda_i)^2} \left[\begin{array}{c} -\left(\frac{R_i^{12}}{\lambda_i^2}\right)\left(a - b\left(\frac{\lambda_i - \gamma_i^0}{2\gamma_i^1} + \frac{\lambda_j - \gamma_j^0}{2\gamma_j^1}\right) - \frac{\gamma_i^0}{2} - \frac{\lambda i}{2}\right) \\ +\left(\frac{R_i^{11}}{\lambda_i}\right)\left(\left(\frac{\lambda_i - \gamma_i^0}{2\gamma_i^1\lambda_i}\right)\left[a - b\left(\frac{\lambda_i - \gamma_i^0}{2\gamma_i^1} + \frac{\lambda_j - \gamma_j^0}{2\gamma_j^1}\right) - \frac{\gamma_i^0}{2} - \frac{\lambda_i}{2}\right] - \frac{(b+\gamma_i^1)\sigma}{\lambda_i}\right) \end{array} \right],$$

where it can be shown that:

$$\begin{aligned} \lim_{\lambda \to a} \frac{R_{11}}{\lambda} &= \frac{1}{a(1 - \log \beta)}, \\ \lim_{\lambda \to a} \frac{R_{12}}{\lambda} &= \frac{a - \gamma^0}{2\gamma^1 a} \frac{1}{1 - \log \beta}, \\ \lim_{\lambda \to a} \frac{R_{22}}{\lambda^2} &= \frac{(\frac{a - \gamma^0}{2\gamma^1 a})^2 + \frac{\sigma}{a^2}}{1 - 2\log \beta}. \end{aligned}$$

Under the cooperative SCE we have $\gamma^0 = a$ and $\gamma^1 = -2b$. Substituting in this at the point of exhaustion gives $\dot{\gamma}_i^0 = 0$ but $\dot{\gamma}_i^1$ is non-zero involving a term in $\frac{b\sigma}{\lambda_i}$. Therefore at exhaustion producers are still updating their beliefs about the perceived elasticity of demand. ■

The reason convergence in beliefs never fully occurs is because the zero output at exhaustion implies from the producer's first order condition that $\lambda = \gamma_0$, which equals a at the CE SCE. Therefore the non-negativity constraint for output stops $\lambda \to \infty$. Intuitively, after the supplier has produced the very last unit of output they realise their estimates of the elasticity of demand were incorrect and would like to revise them. However with no more output to produce this change in beliefs has no further effect on the market. Obviously the larger is a then the larger is λ and the closer beliefs get to the cooperative equilibrium that is the SCE with infinite scarcity. As $a \to \infty$ the market gets arbitrarily close to the cooperative equilibrium.

4.5 Numerical example

To understand better how the SCE shifts with the level of scarcity and how at exhaustion beliefs are close to but not quite converged we consider a numerical example. We set $a = 2$, $b = 0.1$ and the standard deviation of control errors $\sigma = 0.1$. The initial resource stock is set to $Q_0 = 5.3746$ for each supplier so as to normalise the initial time to exhaustion T_0 to unity. For the discount factor we consider an annual model with 100 years to exhaustion, so $\beta = 0.95^{100}$. We stress that this calibration is not intended as a serious empirical exercise but merely to illustrate certain features we have already examined analytically. The first figure we present is the evolution of γ.

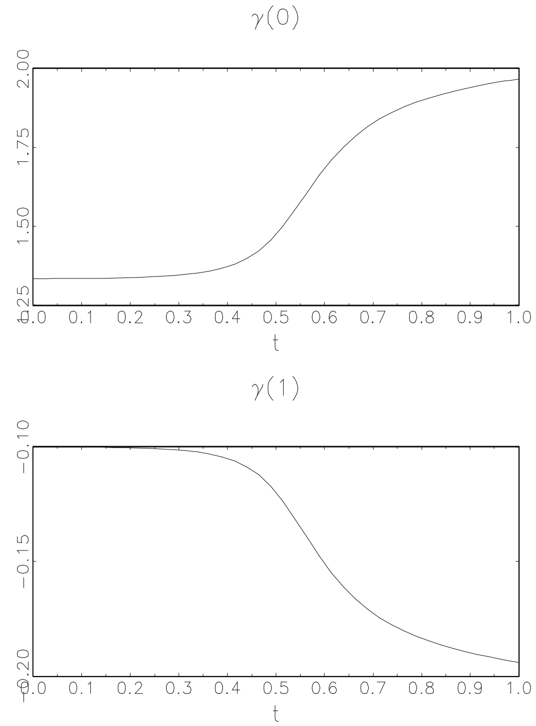

Figure 2: Evolution of beliefs as depletable resource becomes scarce

Figure 2 confirms the result that scarcity interacts with learning to change agent's beliefs and leads to higher estimates of the price inelasticity of demand. In the first panel, belief γ^0 is $(2a/3) = 4/3$ when the resource is relatively abundant and then shifts towards $a = 2$ as the resource becomes scarce. In the second panel, belief γ^1 changes from $-b = -0.1$ to

$-2b = -0.2$, as predicted by Propositions 1 and 2. When the resource is abundant there is no coordination between suppliers and changes in supply are dominated by random control errors that are independent across suppliers. Random control errors feed through to price with multiplier $-b$ in the true linear demand curve (1), so on average demand appears to be relatively price elastic. When the resource becomes scarce, changes in supply become increasingly dominated by scarcity considerations. A supplier then sees its own contraction in supply having a large effect on price - because the other supplier is also contracting supply - and concludes that demand is relatively price inelastic. Close examination of Figure 2 also confirms Proposition 3; at the point of exhaustion beliefs have not quite reached those of the cooperative SCE. Our e-stability result implies that they are tending towards the SCE and the simulation results show that beliefs have nearly converged but at the point of exhaustion the market is frozen in time and beliefs never quite reach the SCE.

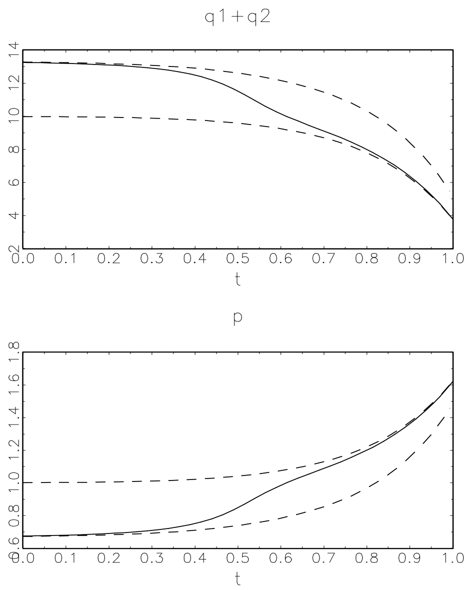

Figure 3: Evolution of self-confirming equilibrium as depletable resource becomes scarce

Figure 3 plots the evolution of the mean dynamics of supply and price as the resource becomes scarce. For comparison, we include the supply and prices that would prevail if the suppliers were in non-cooperative and cooperative rational expectations equilibrium. As expected, in the top panel supply matches that in non-cooperative rational expectations equilibrium when the resource is abundant, but as the resource becomes scarce there is a change towards the cooperative rational expectations equilibrium. The bottom panel of the figure shows the corresponding evolution of market price. The presence of learning means that output and price no longer trend at the constant Hotelling rate of $-\log\beta$; the gradual shift in the SCE provides additional low frequency volatility.

5 Escape episodes

Williams (2001) and Cho, Williams and Sargent (2002) show how the existence of escape episodes away from the SCE are capable of providing rich dynamics. In this section we continue to use the techniques of stochastic approximation to analyse the continuous time analogue of our model and characterise the escape dynamics. Without learning these escape dynamics do not exist. With learning these escape dynamics open up the ability to produce substantial volatility in commodity prices, far in excess of either the control errors or the longer term Hotelling trend.

5.1 Escape dynamics

The formal analysis of escape dynamics in economic models is laid out in Williams (2001), where the dominant escape path is characterised by solving an optimal control problem. Mathematically, the dominant escape path solves:

$$\bar{S} = \inf_{\dot{v}} \frac{1}{2} \int_{0}^{t} \dot{v}(s)' Q(\gamma(s), R(s), \lambda(s))^{-1} \dot{v}(s) ds$$

$$s.t.$$

$$\dot{\gamma} = R^{-1}\bar{g}(\gamma, \lambda) + \dot{v}$$

$$\dot{R} = \bar{M}(\gamma, \lambda) - R$$

$$\dot{\lambda} = \bar{f}(\gamma, R, \lambda)$$

$$\gamma(0) = \bar{\gamma},\ M(0) = \bar{M},\ \lambda(0)\ given,\ \gamma(t) \notin G\ for\ some\ 0 < t < T \tag{20}$$

The optimal control problem works by perturbing the mean dynamics of the model (16) - (17) by a factor $\dot{v}$ and asking which series of perturbations is most likely to cause beliefs to escape. The function $\dot{\lambda} = f(\gamma, R, \lambda)$ is a direct summary of equation (19). In the objective $Q(\gamma, R, \lambda)$ is a weighting function that measures the likelihood of the shocks needed to perturb beliefs by $\dot{v}$. We initialise beliefs at their self-confirming values and define a neighbourhood G around the self-confirming equilibrium that beliefs must escape from. The outcome of the optimal control problem is the series of belief perturbations that occur along the dominant escape path.

5.2 Dominant Escape Path

To solve for the dominant escape path we define a Hamiltonian, where a, b and c are co-state vectors for the evolution of γ, R and λ:[6]

$$\mathcal{H} = aR^{-1}\bar{g}(\gamma, \lambda) - \frac{1}{2}a'Q(\gamma, R, \lambda)a + b \cdot (\bar{M}(\gamma, \lambda) - R) + c\bar{f}(\gamma, R, \lambda). \tag{21}$$

The Hamiltonian is convex so first order conditions (22) - (27) necessarily hold along the

[6] An analytic expression for $a'Q(\gamma, R, \lambda)a$ is given in Appendix C.

dominant escape path:

$$\dot{\gamma} = R^{-1}\bar{g}(\gamma,\lambda) - Q(\gamma,R,\lambda)a, \tag{22}$$

$$\dot{R} = \bar{M}(\gamma,\lambda) - R, \tag{23}$$

$$\dot{\lambda} = f(\gamma,R,\lambda), \tag{24}$$

$$\dot{a} = -\mathcal{H}_{\gamma}, \tag{25}$$

$$\dot{b} = -\mathcal{H}_{R}, \tag{26}$$

$$\dot{c} = -\mathcal{H}_{\lambda}. \tag{27}$$

The first order conditions form a system of ordinary differential equations that characterise a family of escape paths, with each path being indexed by different initial values of the co-state vectors. The dominant escape path is the member of this family that achieves the escape with the most likely series of belief perturbations. A solution to the optimal control problem can therefore be obtained by searching over all possible initial values of a, b and c, applying equations (22) - (27), and choosing initial values that imply beliefs perturbations that are most likely in terms of the $Q(\gamma,R,\lambda)$ metric.

5.3 Numerical example

To illustrate the nature of the dominant escape path we return to our numerical example with $a = 2$, $b = 0.1$, $\sigma = 0.1$ and $\beta = 0.95^{100}$. The first dominant escape path we report is for the case of no resource scarcity. It is shown in the top-left panel of Figure 4. The path is similar to those of Williams (2001), with beliefs spending a long time near the mean dynamics before escaping rapidly at $t \approx 2$ to new values close to $\gamma_i = (a, -2b)$, the cooperative equilibrium. Intuitively, an escape happens when a sequence of control errors causes both suppliers to simultaneously contract supply. The resulting increase in price causes each supplier to start to believe that demand is price inelastic, which gives further incentives for suppliers to contract supply. Price rises again and there is even more reason to contract supply in the belief that demand is inelastic. This reinforces beliefs in inelastic demand and price rises rapidly as supply contracts. In the case with no scarcity the mean dynamics start with the non-cooperative equilibrium as a SCE and so the escape dynamics give us the same pattern of intermittent fluctuations between a NCE and a CE that is the focus of Green and Porter (1984). As in Green and Porter these movements are not the result of genuinely cooperative behaviour although in contrast to their mechanism ours is not based on strategic behaviour on the part of firms. The interaction of learning with scarcity is what drives the shifts in our case.

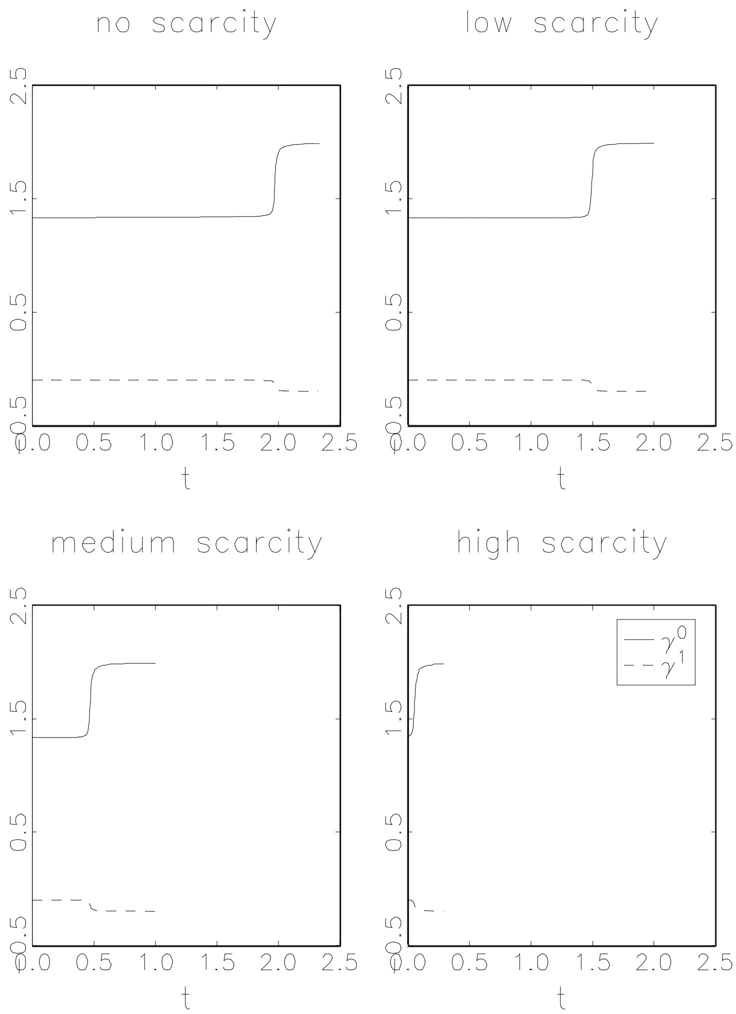

Figure 4: Dominant escape paths with different levels of resource scarcity

Escape dynamics show fluctuations from the mean dynamics to the cooperative outcome. Because of rising scarcity the mean dynamics in our model are also changing. As a consequence fluctuations in depletable commodity prices will be strongly non-stationary, with their volatility and persistence depending on the degree of scarcity. The remaining panels of Figure 4 focus on this by showing how the presence of scarcity affects escape dynamics. Starting from the top two panels we see the introduction of scarcity brings forward the timing of the escape along the dominant escape path. The pattern continues in the bottom two panels, the latter of which showing that the escape becomes almost instantaneous when scarcity reaches high levels. Table 1 shows the same behaviour in an alternative format by reporting the relationship between the time to escape and the level of resource scarcity (as measured by the time to ultimate exhaustion of the depletable resource).

Time to exhaustion	Time to escape
4	2.02
3.5	2.05
3	2.03
2.5	1.53
2	1.04
1.5	0.51
1	0.09

Table 1: Relationship between time to exhaustion and time to escape

The fall in time to escape is explained by scarcity creating a background of falling supply and increasing price[7]. In such circumstances, it is more likely that the special combination of control errors needed to trigger an escape will occur. Recall that simultaneous negative control errors were needed to trigger an escape in the no scarcity case. This is also true here, but the natural uncoordinated contraction of supply induced by scarcity means that smaller control errors are needed to trigger an escape.

[7]McGough (2006) shows how shocks to the natural rate of unemployment alter the time to escape in Sargent's (1999) model. The effect of scarcity in our model is similar, although the fact that our model has a SCE that shifts provides additional features.

To better understand how escape dynamics interact with scarcity to influence the dynamics of commodity prices, return to Figure 1 in the introduction which shows simulations of our basic model. The simulations naturally illustrate the various features we have analysed so far:

1. Industry output declines over time because the commodity is in finite supply
2. Commodity prices display an upward trend based around the rate of time preference.
3. As scarcity starts to bind, the SCE slowly changes and the trend behaviour of prices and output shift from the level of the non-cooperative equilibrium towards the cooperative outcome. The trend in prices is faster than the pure Hotelling rate of the time preference and commodity prices display greater low frequency volatility.
4. Commodity prices are vulnerable to escape dynamics. Escape dynamics take the form of the market jumping from the mean dynamics towards the cooperative equilibrium.
5. Escape episodes become more frequent as scarcity binds.
6. The magnitude of escapes declines over time because the mean dynamics shift towards the cooperative equilibrium.

Figure 1 suggests that learning in the context of commodities in finite supply can add substantial volatility and complex dynamics to commodity prices. Under learning and scarcity, commodity prices will rise over time but will be characterised by sharp breakouts of prices during escape dynamics. During these periods it is as if producers believe that the commodity has become much scarcer and prices rise rapidly to reflect this.

6 Making producers smarter

A key assumption so far has been that each firm thinks its competitor has not systematically changed its recent output, i.e. $q_{jt} = q_j + \eta_{jt}$. This was critical for our result that the self confirming equilibrium changes over time as the mean dynamics shift towards the cooperative equilibria. In this section we replace this assumption with an alternative whereby each producer thinks its competitor's output follows a trend. In making this assumption we shift our model closer to rational expectations - under the Hotelling model this is exactly what is happening.

6.1 Model with perceived trends in competitor's output

The demand curve for our commodity market remains as before:

$$p_t = a - b(q_{it} + q_{jt}), \tag{28}$$

but we modify beliefs such that firm i believes firm j's output is trending according to:

$$q_{jt} = \bar{q}_j + \tilde{q}_j \tilde{\beta}^{-t} + \eta_{jt}, \tag{29}$$

where $\bar{q}_j > 0$ and $\tilde{q}_j < 0$ would imply competitor supply is trending downwards for $\tilde{\beta} < 1$. $\tilde{\beta}$ is the rate at which a firm perceives its competitor's supply is trending. If $\tilde{\beta} = \beta$ then the firm perceives the same trend in its own and competitor's supply. Note that in Section 2 we showed that a symmetric cooperative rational expectations equilibrium was characterised by $q_{jt} = \frac{a-\lambda_t}{4b}$ and $\lambda_t = \lambda_0 \beta^{-t}$, in which case $\bar{q}_j = a/4b$, $\tilde{q}_j = -\lambda_0/4b$ and $\tilde{\beta} = \beta$. The perception (29) is therefore consistent with a minimum state variable representation of a rational expectations equilibrium, but critically we assume the producer i has to infer the parameters $\bar{q}_j$ and $\tilde{q}_j$ from the data. Substituting (29) in (28) gives:

$$p_t = a - b(q_{it} + \bar{q}_j + \tilde{q}_j \tilde{\beta}^{-t} + \eta_{jt}),$$

and agents update their beliefs on the basis of the regression:

$$p_t = \hat{\gamma}^0_{it} + \hat{\gamma}^1_{it} q_{it} + \hat{\gamma}^2_{it} \tilde{\beta}^{-t} + \eta'_{jt}.$$

The first order condition for profit maximisation is:

$$\beta^k (\hat{\gamma}^0_{it} + 2\hat{\gamma}^1_{it} q_{1t+k} + \hat{\gamma}^2_{it} \tilde{\beta}^{-(t+k)}) = \lambda_{it},$$

and period t output is given by:

$$\hat{q}_{it} = \frac{\lambda_{it} - \hat{\gamma}^0_{it} - \hat{\gamma}^2_{it}\tilde{\beta}^{-t}}{2\hat{\gamma}^1_{it}}. \quad (30)$$

Allowing a producer to perceive a downward trend in the supply of its competitor has both static and dynamic effects on supply. The *static effect* comes about because the term $\hat{\gamma}^2_{it}\tilde{\beta}^{-t} < 0$ in equation (30) leads directly to an increase in supply in the current period. The presence of the term $\hat{\gamma}^2_{it}\tilde{\beta}^{-t}$ here is analogous to changing the perceived size of the current market $\hat{\gamma}^0_{it}$. A negative value of $\hat{\gamma}^2_{it}\tilde{\beta}^{-t}$ implies that the competitor is supplying less to the market now and so there is more of the market for the producer to capture by increasing supply in the current period. Hence supply increases. The *dynamic effect* acts through λ_{it} in equation (30) and the extent to which scarcity binds when a competitor is perceived to have output trending downwards. To see how this mechanism works note that at the point of exhaustion:

$$\lambda_{it} = \beta^{T_i}\left(\hat{\gamma}^0_{it} + \hat{\gamma}^2_{it}\tilde{\beta}^{-(t+T_i)}\right),$$

so $\hat{\gamma}^2_{it}\tilde{\beta}^{-(t+T_i)}$ being negative decreases λ_{it} for given $\hat{\gamma}^0_{it}$ and β^{T_i}. The dynamic effect serves to decrease λ_{it} and therefore supply by equation (30). Intuitively, if a producer perceives that its competitor will reduce supply in the future then it is optimal for the producer to reduce supply now in order to supply more in the future when competitor output is less.

To solve for T_i we proceed as before and sum the first order conditions:

$$T_i\hat{\gamma}^0_{it} + 2\hat{\gamma}^1_{it}Q_{it} + \hat{\gamma}^2_{it}\tilde{\beta}^{-t}\frac{1-\tilde{\beta}^{-T_i}}{1-\tilde{\beta}^{-1}} = \beta^{T_i}(\hat{\gamma}^0_{it} + \hat{\gamma}^2_{it}\tilde{\beta}^{-(t+T_i)})\frac{1-\beta^{-T_i}}{1-\beta^{-1}}.$$

The time to exhaustion T_i is affected by both the static effect $\hat{\gamma}^2_{it}\tilde{\beta}^{-t}\frac{1-\tilde{\beta}^{-T_i}}{1-\tilde{\beta}^{-1}}$ and the dynamic effect $\beta^{T_i}\hat{\gamma}^2_{it}\tilde{\beta}^{-(t+T_i)}\frac{1-\beta^{-T_i}}{1-\beta^{-1}}$. In the special case $\beta = \tilde{\beta}$, the static and dynamic effects exactly offset each other and the equation simplifies to:

$$T_i\hat{\gamma}^0_{it} + 2\hat{\gamma}^1_{it}Q_{it} = \beta^{T_i}\hat{\gamma}^0_{it},$$

and T_i is independent of $\hat{\gamma}^2_{it}$. Output is given by:

$$q_{it} = \frac{\beta^{T_i}\hat{\gamma}^0_{it} - \hat{\gamma}^0_{it}}{2\hat{\gamma}^1_{it}},$$

as before and is also independent of $\hat{\gamma}_{it}^2$, although it does depend on scarcity through the term β^{T_i}. In the case of learning, the introduction of a perceived trend affects the estimates of $\hat{\gamma}_{it}^0$ and $\hat{\gamma}_{it}^1$ and so will affect market dynamics, but conditional on estimates of $\hat{\gamma}_{it}^0$ and $\hat{\gamma}_{it}^1$ the estimated parameter on the trend has no effect on supply. Conditioning estimation on a perceived trend in this way removes one of our earlier results that the self-confirming equilibrium shifts as scarcity increases: under this assumption the SCE remains unchanged at the non-cooperative equilibrium and scarcity no longer acts as a coordinating device. However, despite this we find that escape dynamics become more important and are responsible for even more volatility in prices. The logic for this is that the interaction of escape and mean dynamics makes it hard for producers to accurately pin down the true trend so changes in beliefs add extra volatility to the market, as in Figure 5.

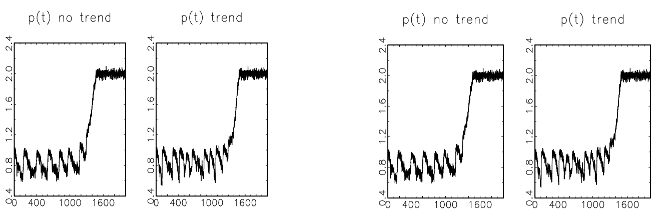

Figure 5: Recurring escapes with and without producers perceiving trends in the supply of their competitors

To see why this is the case consider when prices are falling as the market leaves an escape episode. Firms are revising up their estimate of how elastic the demand curve is and so are increasing their output whilst market price falls as a result. At this point each firm now infers a *positive* trend in competitor output, which leads them to further revise up their estimate of the elasticity of demand. As output levels approach towards those of the non-cooperative equilibrium both firms slow the increase in their output. However, each firm is still inferring a faster upward trend in their competitor's output and so in response to their own easing in output growth they witness a larger than expected increase in market price. As a result, they once again start to believe the demand curve is less elastic than they previously estimated and both firms begin to cut their production in an uncoordinated manner. The price rises sharply as both firms cut their production and once again through learning each firm mistakenly reduces their estimate of the elasticity of demand. Therefore escapes happen more quickly as they do not rely so strongly on the trigger of a particular sequence of stochastic shocks. Instead, they are boosted by the natural increases in production that follow the end of an escape and the simultaneous ending of production increases by both firms as they reach the non-cooperative equilibrium. Therefore, while making the more realistic assumption that competitor's output follows a trend removes the shifting self-confirming equilibrium effect it does paradoxically add even great volatility to prices through the increasing the frequency of escape dynamics.

7 Extraction costs and stochastic demand shifts

Our model of depletable commodity markets omits many important real world considerations. In this section we show how to extend our learning analysis to allow for extraction costs and stochastic shifts in demand. We show how extending the model requires few changes in our analysis and maintains our main qualitative results, although obviously it alters the quantitative nature of shifts in the SCE and escape dynamics.

7.1 Extraction costs

Assume as in Hansen, Epple and Roberds (1985) that total extraction costs are given by:

$$c_{it+k}q_{it+k} = a + bq_{it+k} + cq_{it+k}^2 + d(Q_{i0} - Q_{it+k+1}^*),$$

where:

$$Q^*_{it+k+1} = \sum_{j=k+1}^{\infty} q_{it+j}$$

denotes all future production, i.e. the current stock of the commodity left to be extracted. Total costs therefore depend on a fixed cost of production a, variable costs which depend linearly and quadratically on current production $bq_{it+k} + cq^2_{it+k}$, and costs that depend on the stock of the commodity left to be extracted $d(Q_{i0} - Q^*_{it+k+1})$. If $d > 0$ then extraction costs rise as fewer resources are left to be extracted. With this specification of extraction costs the firm's first order condition becomes:

$$\beta^k(\gamma^0_{it} - b + 2(\gamma^1_{it} - c)q_{it+k}) - \beta^{k-1}d - \beta^{k-2}d - ... - d = \lambda_{it}.$$

Supply in the first period is:

$$q_{it} = \frac{\lambda_{it} - \gamma^0_{it} + b - \frac{d}{1-\beta}}{2(\gamma^1_{it} - c)},$$

and at the time of exhaustion the first order condition of the firm implies:

$$\beta^T(\gamma^0_{it} - b + \frac{d}{1-\beta}) - d\frac{(1-\beta^T)}{1-\beta} = \lambda_{it} + \frac{d}{1-\beta}.$$

Summing the first order conditions between t and T_i gives:

$$T_i(\gamma^0_{it} - b + \frac{d}{1-\beta}) + 2(\gamma^1_{it} - c)Q_{it} = (\lambda_{it} + \frac{d}{1-\beta})\frac{(1-\beta^{-T_i})}{1-\beta^{-1}},$$

so if we redefine $\gamma^{0*}_{it} = \gamma^0_{it} - b + \frac{d}{1-\beta}$, $\gamma^{1*}_{it} = \gamma^1_{it} - c$ and $\lambda^*_{it} = \lambda_{it} + \frac{d}{1-\beta}$ then our previous analysis carries through exactly as before but with the transformed variables such that:

$$\frac{\dot{\lambda}^*_i}{\lambda^*_i} = \left(\frac{\log\lambda^*_i - \log\gamma^{0*}_i}{\lambda^*_i - \gamma^{0*}_i}\right)\dot{\gamma}^{0*}_i + \left(1 - \gamma^{0*}_i\frac{\log\lambda^*_i - \log\gamma^{0*}_i}{\lambda^*_i - \gamma^{0*}_i}\right)\frac{\dot{\gamma}^{1*}_i}{\gamma^{1*}_i} - \log\beta.$$

7.2 Stochastic demand shifts

Our model so far has only had one source of stochastic uncertainty - the control errors that influence each firm's production. However, another likely area of uncertainty is stochastic shifts in the demand curve. Such shifts will complicate inference over the structural parameters a

and b and clearly affect the dynamics of the model and in particular the likelihood of escape dynamics. The simplest way of introducing demand shifts is to assume:

$$p_t = a - b(q_1 - q_2) + e_t,$$

where e_t is i.i.d. with mean zero. Repeating our previous analysis we find that the presence of i.i.d. demand shocks makes no difference to our mean dynamics and so does not alter the behaviour of the SCE. Demand shocks are zero in expectation and so exert no influence on the mean dynamics of the system. However, the presence of demand shocks does influence the likelihood of escape dynamics.

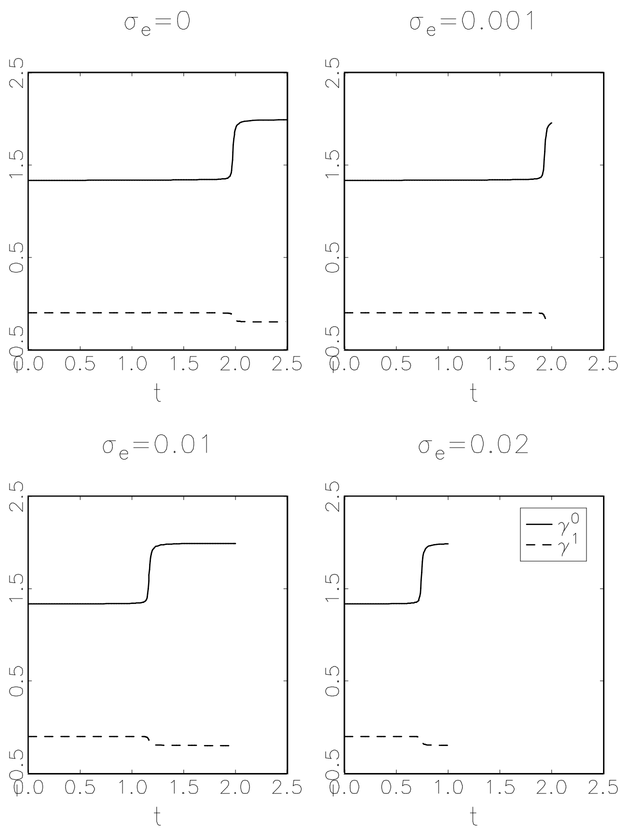

Figure 6: Dominant escape paths with different levels of stochastic demand shifts

As shown in Figure 6, the time to escape shortens as we increase the volatility of demand shocks. If there are no demand shocks then the time to escape in the non-finite resource case is 2.02 (as in the first line of Table 1). If the standard deviation of demand shocks is 0.001 then the time to escape falls to 1.97, if the standard deviation of demand shocks is 0.01 then the time to escape is 1.21. To understand why this happens recall our earlier result that scarcity acts as a coordinating device and as scarcity increases escapes become more prevalent. Demand shocks also act as a coordinating device - they act to raise prices for both firms and lead each firm to believe that price is now less responsive to output than before and so both firms simultaneously cut back production. Escapes to the cooperative outcome are therefore more prevalent.

8 Conclusions

Depletable resource markets are characterised by a finite stock and frequently significant market power. Since Hotelling (1931) the implications of this for the optimal rate of resource exhaustion and the behaviour of prices have been analysed. However, the focus of this analysis has been on the low frequency behaviour of prices that distinguishes nonrenewable resources from renewable commodities. By contrast our focus is on how uncertainty and shifting perceptions can influence high frequency shifts in commodity prices. In an environment characterised by uncertainty over the extent of scarcity and the degree of market power possessed by each agent, we show how learning and changes in agents' perceptions can introduce substantial market volatility over and above that suggested by fundamentals.

We show that in a world characterised by no concerns over scarcity our duopoly model reaches a self confirming equilibrium which is the non-cooperative Cournot-Nash equilibrium. In the case where scarcity is apparent to producers then the self confirming equilibrium is instead characterised by the cooperative Cournot-Nash equilibrium whereby each producer restricts output and the market price is high. In effect scarcity acts as a coordinating device and even though producers do not collude they restrict output and arrive at the cooperative outcome. Therefore as scarcity increases for a commodity the price experiences a shift to the higher cooperative solution.

We show that the existence of these two different self confirming equilibria, each dependent on the perception of scarcity, opens up the possibility of escape dynamics. During the period of escape dynamics prices of depletable resources are characterised by sharp upward increases and then more gradual declines. In this way learning dynamics provide substantial volatility over and above that driven by fundamentals. However, fundamentals do exert an influence on the nature of these dynamics. As scarcity becomes more apparent in reality then the jumps in prices towards the cooperative solution become smaller and these higher prices are more robust and persistent. Our analytical findings confirm the existence of these escape dynamics but our simulations confirm not just that these exist but they occur with a frequency and magnitude that adds considerable volatility to the price dynamics of finite resources.

Ours is obviously a highly stylised model of depletable commodity markets. However, the relevance of the learning mechanism we have outlined hinges on three key characteristics which we believe have empirical resonance - the commodity is depletable, there are a few producers and the market is characterised by demand uncertainty and lack of complete knowledge of the actions of other producers.

References

[1] Adam, K., Marcet, A. and J-P. Nicolini, 2007, Stock Market Volatility and Learning, *CEPR Discussion Paper* No. 6518

[2] Brock, W.A and C.H. Hommes, 1997, A Rational Route to Randomness, *Econometrica* 65(5) 1059-95

[3] Bullard, J. and K. Mitra, 2002, Learning about Monetary Policy Rules, *Journal of Monetary Economics* 49(6), 1105-1129

[4] Cho I.-K., Williams, N. and T.J. Sargent, 2002, Escaping Nash Inflation, *Review of Economic Studies* 70(2), 1-40

[5] Dasgupta, P. and G. Heal, 1974, Optimal Depletion of Exhaustible Resources, *Review of Economic Studies*, 3-24

[6] Deaton, A. and G. Laroque, 1996, Competitive Storage and Commodity Price Dynamics, *Journal of Political Economy* 104(5), 896-923

[7] Evans, G. and S.M.S. Honkapohja, 1993, Adaptive forecasts, hysteresis and endogenous fluctuations *Economic Review Federal Reserve Bank of San Francisco* 3-13

[8] Evans, G. and S.M.S. Honkapohja, 1994, Learning, convergence and stability with multiple rational expectations equilibria, *European Economic Review* 37(2-3), 595-602

[9] Evans, G. and S.M.S. Honkapohja, 1995, Local convergence of recursive learning to steady states and cycles in stochastic nonlinear models, *Econometrica* 63, 195-206

[10] Fudenberg, D. and D. Levine, 1993a, Self-Confirming Equilibria, *Econometrica* 61(3), 523-45

[11] Fudenberg, D. and D. Levine, 1993b, Steady State Learning and Nash Equilibrium *Econometrica* 61(3), 547-573

[12] Fudenberg, D. and D. Levine, 2009, Self Confirming Equilibria and the Lucas Critique *Journal of Economic Theory* forthcoming

[13] Green, E.J. and R.H. Porter, 1984. Non-cooperative collusion under imperfect price information, *Econometrica* 52, 87-100

[14] Hansen, L.P., Epple, D. and W. Roberds, 1985, Linear-Quadratic Duopoly Models of Resource Depletion, in Sargent (ed.), *Energy, Foresight and Strategy*, Resources for the Future, Inc. 101-142

[15] Hotelling, H., 1931, The Economics of Exhaustible Resources, *Journal of Political Economy*, 39(2), 137-75

[16] Kandori, M, Mailath, G.J., and R. Rob, 1993, Learning, Mutation and Long Run Equilibria in Games, *Econometrica* 61(1) 29-56

[17] Kasa, K., 2004, Learning, Large Deviations and Recurrent Currency Crises, *International Economic Review* 45, 141-174

[18] Kreps, D.M., 1998, Anticipated utility and dynamic choice, in *Frontiers of research in economic theory: the Nancy L. Schwartz memorial lectures, 1983-1997*, Cambridge University Press

[19] Lettau, M. and H. Uhlig, 1999, Rules of Thumb versus Dynamic Programming, *American Economic Review* 89(1), 148-174

[20] Loury, G., 1986, A Theory of Oil'gopoly: Cournot Equilibria in Exhaustible Resource Markets with Finite Supply, *International Economic Review* 27, 285-301

[21] Marcet, A. and T.J. Sargent, 1989, Convergence of Least Squares Learning mechanisms in self-referential linear stochastic models, *Journal of Political Economy* 97(6), 1306-22

[22] McGough, B., 2006, Shocking Escapes, *Economic Journal* 116(511), 507-28

[23] Piazzesi, M. and M. Schneider,2007, Equilibrium Yield Curves *NBER Macroeconomics Annual* 21(1) 389-441

[24] Porter, R.H., 1983, A Study of Cartel Stability: The Joint Executive Committee 1880-1886, *Bell Journal of Economics* 14(2), 301-314

[25] Salant, S., 1976, Exhaustible Resources and Industrial Structure: A Nash-Cournot Approach to the World Oil Market, *Journal of Political Economy* 84, 1079-1093

[26] Sargent, T.J., 1999, *The Conquest of American Inflation*, Princeton University Press

[27] Sargent, T.J, Williams, N., and T. Zha, 2009, The Conquest of South American Inflation, *Journal of Political Economy* 117(2), 211-256

[28] Stiglitz, J.E and P. Dasgupta, 1982, Market Structure and Resource Depletion: A Contribution to the Theory of International Monopolistic Competition, *Journal of Economic Theory* 28, 128-64

[29] Williams, N., 2001, Escape Dynamics in Learning Models, *University of Chicago UMI #3006566*

A E-stability of Self-Confirming Equilibrium without resource scarcity

To show that the self-confirming equilibrium is e-stable it is sufficient to show that the eigenvalues of the Jacobian have negative real parts when evaluated at the self-confirming equilibrium. The Jacobian of the system is defined by:

$$J = \begin{pmatrix} \frac{\partial \dot{\gamma}_i}{\partial \gamma_i} & \frac{\partial \dot{\gamma}_i}{\partial \gamma_j} & \frac{\partial \dot{\gamma}_i}{\partial R_i} & \frac{\partial \dot{\gamma}_i}{\partial R_j} \\ \frac{\partial \dot{\gamma}_j}{\partial \gamma_i} & \frac{\partial \dot{\gamma}_j}{\partial \gamma_j} & \frac{\partial \dot{\gamma}_j}{\partial R_i} & \frac{\partial \dot{\gamma}_j}{\partial R_j} \\ \frac{\partial \dot{R}_i}{\partial \gamma_i} & \frac{\partial \dot{R}_i}{\partial \gamma_j} & \frac{\partial \dot{R}_i}{\partial R_i} & \frac{\partial \dot{R}_i}{\partial R_j} \\ \frac{\partial \dot{R}_j}{\partial \gamma_i} & \frac{\partial \dot{R}_j}{\partial \gamma_j} & \frac{\partial \dot{R}_j}{\partial R_i} & \frac{\partial \dot{R}_j}{\partial R_j} \end{pmatrix},$$

which when evaluated at self-confirming equilibrium reduces to:

$$J\,|_{SCE} = \begin{pmatrix} \left.\frac{\partial \dot{\gamma}_i}{\partial \gamma_i}\right|_{SCE} & \left.\frac{\partial \dot{\gamma}_i}{\partial \gamma_j}\right|_{SCE} & 0 & 0 \\ \left.\frac{\partial \dot{\gamma}_j}{\partial \gamma_i}\right|_{SCE} & \left.\frac{\partial \dot{\gamma}_j}{\partial \gamma_j}\right|_{SCE} & 0 & 0 \\ \left.\frac{\partial \dot{R}_i}{\partial \gamma_i}\right|_{SCE} & 0 & -I & 0 \\ 0 & \left.\frac{\partial \dot{R}_j}{\partial \gamma_j}\right|_{SCE} & 0 & -I \end{pmatrix}.$$

The eigenvalues of the $-I$ identity matrices trivially have negative real parts so a sufficient condition for e-stability is that the following matrix has eigenvalues with negative real parts:

$$\begin{pmatrix} \frac{\partial \dot{\gamma}_i}{\partial \gamma_i}\Big|_{SCE} & \frac{\partial \dot{\gamma}_i}{\partial \gamma_j}\Big|_{SCE} \\ \frac{\partial \dot{\gamma}_j}{\partial \gamma_i}\Big|_{SCE} & \frac{\partial \dot{\gamma}_j}{\partial \gamma_j}\Big|_{SCE} \end{pmatrix}.$$

Simple but tedious calculations show that the matrix is of the form:

$$\begin{pmatrix} -1 & 0 & -\frac{1}{2} & -\frac{a}{3b} \\ 0 & -1 & 0 & 0 \\ -\frac{1}{2} & -\frac{a}{3b} & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix},$$

and the eigenvalues $-0.5, -1, -1, -1.5$ all have negative real parts. The self-confirming equilibrium is e-stable.

B E-stability of Self-Confirming Equilibrium with resource scarcity

The proof of e-stability with resource scarcity relies again on showing that the eigenvalues of the Jacobian have negative real parts when evaluated at the self-confirming equilibrium. The calculations are more complex than in Appendix A because the system under scarcity includes a dynamic equation for λ which feeds back on the dynamics of γ_i^0 and γ_i^1. In particular, the system under investigation is:

$$\begin{aligned}
\dot{\gamma}_i^0 &= \frac{1}{R_i^{22} - (R_i^{12})^2} \left[\begin{array}{c} R_i^{22} \left(a - b \left(\frac{\lambda_i - \gamma_i^0}{2\gamma_i^1} + \frac{\lambda_j - \gamma_j^0}{2\gamma_j^1} \right) - \frac{\gamma_i^0}{2} - \frac{\lambda i}{2} \right) \\ -R_i^{12} \left(\left(\frac{\lambda_i - \gamma_i^0}{2\gamma_i^1} \right) \left[a - b \left(\frac{\lambda_i - \gamma_i^0}{2\gamma_i^1} + \frac{\lambda_j - \gamma_j^0}{2\gamma_j^1} \right) - \frac{\gamma_i^0}{2} - \frac{\lambda_i}{2} \right] - \frac{(b+\gamma_i^1)\sigma}{\lambda_i} \right) \end{array} \right], \\
\dot{\gamma}_i^1 &= \frac{1}{R_i^{22} - (R_i^{12})^2} \left[\begin{array}{c} -R_i^{12} \left(a - b \left(\frac{\lambda_i - \gamma_i^0}{2\gamma_i^1} + \frac{\lambda_j - \gamma_j^0}{2\gamma_j^1} \right) - \frac{\gamma_i^0}{2} - \frac{\lambda i}{2} \right) \\ +R_i^{11} \left(\left(\frac{\lambda_i - \gamma_i^0}{2\gamma_i^1} \right) \left[a - b \left(\frac{\lambda_i - \gamma_i^0}{2\gamma_i^1} + \frac{\lambda_j - \gamma_j^0}{2\gamma_j^1} \right) - \frac{\gamma_i^0}{2} - \frac{\lambda_i}{2} \right] - \frac{(b+\gamma_i^1)\sigma}{\lambda_i} \right) \end{array} \right], \\
\frac{\dot{\lambda}_i}{\lambda_i} &= \left(\frac{\log \lambda_i - \log \gamma_i^0}{\lambda_i - \gamma_i^0} \right) \dot{\gamma}_i^0 + \left(1 - \gamma_i^0 \frac{\log \lambda_i - \log \gamma_i^0}{\lambda_i - \gamma_i^0} \right) \frac{\dot{\gamma}_i^1}{\gamma_i^1} - \log \beta,
\end{aligned}$$

plus equations for $\dot{R}_i^{12}, \dot{R}_i^{22}$ and analogous expressions for $\dot{\gamma}_j^0, \dot{\gamma}_j^1, \dot{\lambda}_j/\lambda_j, \dot{R}_j^{12}$ and $\dot{R}_j^{22}$. We evaluate the Jacobian around the scarcity fixed point $\gamma_i^0 = \gamma_j^0 = a, \gamma_i^1 = \gamma_j^1 = -2b, R_i^{12} = R_j^{12} = 0, R_i^{22} = R_j^{22} = \sigma$ with $\lambda_i = \lambda_j \to a$. Calculations involving repeated application of l'Hôpital's rule show that the Jacobian at self-confirming equilibrium takes the form:

$$\begin{pmatrix}
-\frac{3}{4} & 0 & -\frac{1}{4} & 0 & -\frac{1}{4} & \frac{1}{4} & 0 & 0 & 0 & 0 \\
0 & -\frac{1}{a} & 0 & 0 & -\frac{b}{a^2} & 0 & 0 & 0 & 0 & 0 \\
-\frac{1}{4} & 0 & -\frac{3}{4} & 0 & \frac{1}{4} & -\frac{1}{4} & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & -\frac{1}{a} & 0 & -\frac{b}{a^2} & 0 & 0 & 0 & 0 \\
-\frac{3}{4}+\frac{1}{4a} & 0 & -\frac{1}{4} & 0 & -\frac{1}{4}-\frac{1}{4a}-\log\beta & \frac{1}{4} & 0 & 0 & 0 & 0 \\
-\frac{1}{4} & 0 & -\frac{3}{4}+\frac{1}{4a} & 0 & \frac{1}{4} & -\frac{1}{4}-\frac{1}{4a}-\log\beta & 0 & 0 & 0 & 0 \\
\frac{\partial \dot{R}_i^{12}}{\partial \gamma_i^0} & \frac{\partial \dot{R}_i^{12}}{\partial \gamma_i^1} & \frac{\partial \dot{R}_i^{12}}{\partial \gamma_j^0} & \frac{\partial \dot{R}_i^{12}}{\partial \gamma_j^1} & \frac{\partial \dot{R}_i^{12}}{\partial \lambda_i} & \frac{\partial \dot{R}_i^{12}}{\partial \lambda_j} & -1 & 0 & 0 & 0 \\
\frac{\partial \dot{R}_i^{22}}{\partial \gamma_i^0} & \frac{\partial \dot{R}_i^{22}}{\partial \gamma_i^1} & \frac{\partial \dot{R}_i^{22}}{\partial \gamma_j^0} & \frac{\partial \dot{R}_i^{22}}{\partial \gamma_j^1} & \frac{\partial \dot{R}_i^{22}}{\partial \lambda_i} & \frac{\partial \dot{R}_i^{22}}{\partial \lambda_j} & 0 & -1 & 0 & 0 \\
\frac{\partial \dot{R}_j^{12}}{\partial \gamma_i^0} & \frac{\partial \dot{R}_j^{12}}{\partial \gamma_i^1} & \frac{\partial \dot{R}_j^{12}}{\partial \gamma_j^0} & \frac{\partial \dot{R}_j^{12}}{\partial \gamma_j^1} & \frac{\partial \dot{R}_j^{12}}{\partial \lambda_i} & \frac{\partial \dot{R}_j^{12}}{\partial \lambda_j} & 0 & 0 & -1 & 0 \\
\frac{\partial \dot{R}_i^{22}}{\partial \gamma_i^0} & \frac{\partial \dot{R}_i^{22}}{\partial \gamma_i^1} & \frac{\partial \dot{R}_i^{22}}{\partial \gamma_j^0} & \frac{\partial \dot{R}_i^{22}}{\partial \gamma_j^1} & \frac{\partial \dot{R}_i^{22}}{\partial \lambda_i} & \frac{\partial \dot{R}_i^{22}}{\partial \lambda_j} & 0 & 0 & 0 & -1
\end{pmatrix}$$

The block diagonal structure of the Jacobian means that four eigenvalues are immediately identified as $\{-1, -1, -1, -1, \}$ having negative real parts. Similarly, two eigenvalues are $\{-1/a, -1/a\}$ with negative real parts as $a > 0$. The remaining eigenvalues are of the matrix:

$$\begin{pmatrix}
-\frac{3}{4} & -\frac{1}{4} & -\frac{1}{4} & \frac{1}{4} \\
-\frac{1}{4} & -\frac{3}{4} & \frac{1}{4} & -\frac{1}{4} \\
-\frac{3}{4}+\frac{1}{4a} & -\frac{1}{4} & -\frac{1}{4}-\frac{1}{4a}-\log\beta & \frac{1}{4} \\
-\frac{1}{4} & -\frac{3}{4}+\frac{1}{4a} & \frac{1}{4} & -\frac{1}{4}-\frac{1}{4a}-\log\beta
\end{pmatrix}.$$

To calculate them we invoke a change in variables and define $\tilde{\lambda}_i = \lambda_i \beta^t$, from which it follows that $\dot{\tilde{\lambda}}_i = \dot{\lambda}_i \beta^t + \tilde{\lambda}_i \log \beta$ and the eigenvalues solve:

$$\begin{vmatrix} -\frac{3}{4}-\mu & -\frac{1}{4} & -\frac{1}{4}\beta^{-t} & \frac{1}{4}\beta^{-t} \\ -\frac{1}{4} & -\frac{3}{4}-\mu & \frac{1}{4}\beta^{-t} & -\frac{1}{4}\beta^{-t} \\ \left(-\frac{3}{4}+\frac{1}{4a}\right)\beta^t & \left(-\frac{1}{4}\right)\beta^t & \left(-\frac{1}{4}-\frac{1}{4a}\right)\beta^t-\mu & \frac{1}{4}\beta^t \\ \left(-\frac{1}{4}\right)\beta^t & \left(-\frac{3}{4}+\frac{1}{4a}\right)\beta^t & \frac{1}{4}\beta^t & \left(-\frac{1}{4}-\frac{1}{4a}\right)\beta^t-\mu \end{vmatrix} = 0.$$

By the properties of determinants we have:

$$\beta^t\beta^t \begin{vmatrix} -\frac{3}{4}-\mu & -\frac{1}{4} & -\frac{1}{4} & \frac{1}{4} \\ -\frac{1}{4} & -\frac{3}{4}-\mu & \frac{1}{4} & -\frac{1}{4} \\ \left(-\frac{3}{4}+\frac{1}{4a}\right) & -\frac{1}{4} & \left(-\frac{1}{4}-\frac{1}{4a}\right)-\mu & \frac{1}{4} \\ -\frac{1}{4} & \left(-\frac{3}{4}+\frac{1}{4a}\right) & \frac{1}{4} & \left(-\frac{1}{4}-\frac{1}{4a}\right)-\mu \end{vmatrix} \beta^{-t}\beta^{-t} = 0,$$

and the eigenvalues are $\left\{-1,-1,-\frac{1}{4a},-\frac{1}{4a}\right\}$, which all have negative real parts for $a > 0$ as required.

C Analytical expression for $a'Q(\gamma, R, \lambda)a$

The cost function $Q(\gamma, R, \lambda)$ is used to weight belief perturbations along potential escape paths. It is equal to the variance-covariance matrix of belief dynamics $\dot{\gamma}$. As beliefs are quadratic forms of Gaussian variables, $Q(\gamma, R, \lambda)$ is a fourth moment matrix. In static models such as ours, Williams (2001) shows that Q reduces to the logarithm of a moment generating function, meaning the Hamiltonian (21) can be derived analytically. We begin by expressing the second term of the Hamiltonian by the corresponding moment generating function:

$$a'Q(\gamma, R, \lambda)a = \log E \exp \left\langle a \cdot R^{-1} \left(g(\gamma, \lambda, \xi) - \bar{g}(\gamma, \lambda) \right) \right\rangle . \tag{C.1}$$

We obtain an explicit analytic expression for the right-hand side of (C.1) by using equation

(16) to define

$$g(\gamma,\lambda,\xi)-\bar{g}(\gamma,\lambda)=$$
$$\left(\begin{array}{c} -(b+\gamma_i^1)W_1-bW_2 \\ (-2\hat{q}_i(b+\gamma_i^1)-\hat{q}_jb+a-\gamma_i^0)W_1-\hat{q}_ibW_2-bW_1W_2-(b+\gamma_i^1)W_1^2+(b+\gamma_i^1)\sigma^2 \\ -bW_1-(b+\gamma_j^1)W_2 \\ -\hat{q}_jbW_1+(-2\hat{q}_j(b+\gamma_j^1)-\hat{q}_ib+a-\gamma_j^0)W_2-bW_1W_2-(b+\gamma_j^1)W_2^2+(b+\gamma_j^1)\sigma^2 \end{array}\right).$$

To economise on notation, let R^{-1} and a be defined by:

$$R_i^{-1}=\begin{pmatrix} R_i^1 & R_i^2 \\ R_i^3 & R_i^4 \end{pmatrix} \quad R_j^{-1}=\begin{pmatrix} R_j^1 & R_j^2 \\ R_j^3 & R_j^4 \end{pmatrix} \quad a_i=\begin{pmatrix} a_i^1 \\ a_i^2 \end{pmatrix} \quad a_j=\begin{pmatrix} a_j^1 \\ a_j^2 \end{pmatrix}.$$

The right-hand side of (C.1) can now be expressed in terms of the underlying shocks W_1 and W_2:

$$\log E\exp\left\langle a\cdot R^{-1}\left(g(\gamma,\lambda,\xi)-\bar{g}(\gamma,\lambda)\right)\right\rangle=\log E\left[e^{d_0+d_1W_2+d_2W_1+d_3W_1W_2+d_4W_1^2+d_5W_2^2}\right]. \quad \text{(C.3)}$$

The constants $d_0,\ldots,d_4$ are simple functions (C.4)-(C.9) of the structural parameters $\{a,b\}$, beliefs γ, the Lagrange multiplier λ, the co-state vectors $\{a_i,a_j\}$ and the precision matrices $\{R_i,R_j\}$.:

$$\begin{aligned} d_0 &= (a_i^1R_i^2+a_i^2R_i^4)(b+\gamma_i^1)\sigma^2+(a_j^1R_j^2+a_j^2R_j^4)(b+\gamma_j^1)\sigma^2 && \text{(C.4)} \\ d_1 &= -(a_i^1R_i^1+a_i^2R_i^2)b\sigma-(a_i^1R_i^2+a_i^2R_i^4)\hat{q}_ib\sigma \\ &\quad -(a_j^1R_j^1+a_j^2R_j^2)(b+\gamma_j^1)\sigma+(a_j^1R_j^2+a_j^2R_j^4)(-2\hat{q}_j(b+\gamma_j^1)-\hat{q}_ib+a-\gamma_j^0)\sigma && \text{(C.5)} \\ d_2 &= -(a_i^1R_i^1+a_i^2R_i^2)(b+\gamma_i^1)\sigma+(a_i^1R_i^2+a_i^2R_i^4)(-2\hat{q}_i(b+\gamma_i^1)-\hat{q}_jb+a-\gamma_i^0)\sigma \\ &\quad -(a_j^1R_j^1+a_j^2R_j^2)b\sigma-(a_j^1R_j^2+a_j^2R_j^4)\hat{q}_jb\sigma && \text{(C.6)} \\ d_3 &= -(a_i^1R_i^1+a_i^2R_i^2)b\sigma^2-(a_j^1R_j^1+a_j^2R_j^2)b\sigma^2 && \text{(C.7)} \\ d_4 &= -(a_i^1R_i^2+a_i^2R_i^4)(b+\gamma_i^1)\sigma^2 && \text{(C.8)} \\ d_5 &= -(a_j^1R_j^2+a_j^2R_j^4)(b+\gamma_j^1)\sigma^2 && \text{(C.9)} \end{aligned}$$

The next step is to integrate W_2 out from the right-hand side of equation (C.3):

$$\begin{aligned} & \log E\left[e^{d_0+d_1W_2+d_2W_1+d_3W_1W_2+d_4W_1^2+d_5W_2^2}\right] \\ = \; & d_0+\log E\left[E\left(e^{(d_1+d_3W_1)W_2+d_5W_2^2}\middle| W_1\right)e^{d_2W_1+d_4W_1^2}\right] \end{aligned} \tag{C.10}$$

The expectation conditional on W_1 can be solved analytically by defining $k_1 = d_1 + d_3W_1$ and $k_2 = d_5 - 0.5$ and completing the square of $k_1x + k_2x^2$. In expression (C.11) we have $A = \sqrt{-2k_2}$, $B = -k_1/A$ and $C = -B^2/2$.

$$\begin{aligned} E\left(e^{(d_1+d_3W_1)W_2}e^{d_5W_2^2}\middle| W_1\right) &= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{+\infty} e^{k_1x+k_2x^2}dx \\ &= \frac{e^{-C}}{\sqrt{2\pi}}\int_{-\infty}^{+\infty} e^{-0.5(Ax+B)^2}dx \\ &= \frac{e^{-C}}{A} \end{aligned} \tag{C.11}$$

Using (C.11) to substitute out for the expectation conditional on W_1 in (C.10) leaves an expression in only the W_1 shock:

$$\begin{aligned} & \log E\left[e^{d_0+d_1W_2+d_2W_1+d_3W_1W_2+d_4W_1^2+d_5W_2^2}\right] \\ = \; & d_0+\log E\left[\frac{e^{(2A)^{-2}(d_1+d_3W_1)^2+d_2W_1+d_4W_1^2}}{A}\right] \\ = \; & d_0+\frac{1}{2A^2}d_1^2-\log A+\log E\left[e^{(A^{-2}d_1d_3+d_2)W_1+((2A)^{-2}d_3^2+d_4)W_1^2}\right] \end{aligned} \tag{C.12}$$

This can be solved analytically by defining $\tilde{k}_1 = A^{-2}d_1d_3 + d_2$, $\tilde{k}_2 = (2A)^{-2}d_3^2 + d_4 - 0.5$ and completing the square of $\tilde{k}_1x + \tilde{k}_2x^2$. With $\tilde{A} = \sqrt{-2\tilde{k}_2}$, $\tilde{B} = -\tilde{k}_1/\tilde{A}$ and $\tilde{C} = -\tilde{B}^2/2$ we have:

$$\begin{aligned} E\left[e^{(A^{-2}d_1d_3+d_2)W_1+((2A)^{-2}d_3^2+d_4)W_1^2}\right] &= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{+\infty} e^{\tilde{k}_1x+\tilde{k}_2x^2}dx \\ &= \frac{e^{-\tilde{C}}}{\sqrt{2\pi}}\int_{-\infty}^{+\infty} e^{-0.5(\tilde{A}x+\tilde{B})^2}dx \\ &= \frac{e^{-\tilde{C}}}{\tilde{A}} \end{aligned} \tag{C.13}$$

The final analytic expression for $a'Q(\gamma, R, \lambda)a$ is:

$$a'Q(\gamma, R, \lambda)a = d_0 + \frac{1}{2A^2}d_1^2 - \log A - \log \tilde{A} - \tilde{C}. \tag{C.14}$$