Resource and Energy Economics 33 (2011) 893–912

Contents lists available at ScienceDirect

Resource and Energy Economics

journal homepage: www.elsevier.com/locate/ree

Co-optimization of enhanced oil recovery and carbon sequestration

Andrew Leach [a], Charles F. Mason [b,*], Klaas van 't Veld [b]

[a] University of Alberta School of Business, 3-40K Business Building, Edmonton, Alberta, Canada T6G 2R6
[b] Department of Economics & Finance, University of Wyoming, 1000 E. University Ave., Laramie, WY 82071, United States

ARTICLE INFO

Article history:
Received 26 September 2008
Received in revised form 2 November 2010
Accepted 19 November 2010
Available online 26 November 2010

JEL classification:
Q32
Q40
Q54

Keywords:
Enhanced oil recovery
Carbon sequestration
Climate change

ABSTRACT

In this paper, we present an economic analysis of CO_2-enhanced oil recovery (EOR). This technique entails injection of CO_2 into mature oil fields in a manner that reduces the oil's viscosity, thereby enhancing the rate of extraction. As part of this process, significant quantities of CO_2 remain sequestered in the reservoir. If CO_2 emissions are regulated, oil producers using EOR should therefore be able to earn revenues from sequestration as well as from oil production. We develop a theoretical framework that analyzes the dynamic co-optimization of oil extraction and CO_2 sequestration, through the producer's choice of the fraction of CO_2 in the injection stream at each moment. We find that the optimal fraction of CO_2 is likely to decline monotonically over time, and reach zero before the optimal termination time. Numerical simulations, based on an ongoing EOR project in Wyoming, confirm this result. We also find that cumulative sequestration is less responsive to the carbon tax than to the oil price. Only at very high taxes does a tradeoff between revenues from oil output and sequestration arise.

1. Introduction

There is a growing consensus in both policy circles and in the energy industry that within the next few years, the US Federal government will adopt some form of regulation of CO_2 emissions. At the same time, it is widely believed that much of the nation's energy supply over the coming decades will continue to come from fossil fuels, coal in particular (MIT, 2007). Many analysts believe the only way

* Corresponding author. Tel.: +1 307 766 2178; fax: +1 307 766 5090.
E-mail addresses: andrew.leach@ualberta.ca (A. Leach), bambuzlr@uwyo.edu (C.F. Mason), klaas@uwyo.edu (K.v. Veld).

0928-7655/$ – see front matter
doi:10.1016/j.reseneeco.2010.11.002

to reconcile the anticipated growth in the use of coal with anticipated limits on CO_2 emissions is through the development and deployment of carbon capture and geological sequestration (CCS). There also seems to be agreement that deployment of geological sequestration is likely to start with projects that apply CO_2-enhanced oil recovery (EOR).[1] This technique, which has been used successfully in a number of oil plays (notably in West Texas, Wyoming, and Saskatchewan), entails injection of CO_2 into mature oil fields at pressures high enough to cause the CO_2 to mix with some fraction of the oil that still remains underground. Doing so reduces the oil's viscosity, thereby improving its ability to flow through the reservoir rock, and enhancing the rate of extraction.[2] Although some of the CO_2 resurfaces with the oil, it can be separated from the output stream, recompressed, and reinjected. Eventually, when the EOR project is terminated, much of the injected CO_2 remains sequestered.[3] EOR has the potential of greatly increasing oil output from depleted reserves. Estimates suggest that recovery rates for existing reserves could be approximately doubled, while the application of EOR on a broad scale could raise domestic recoverable oil reserves in the United States by over 80 billion barrels (ARI, 2006). Similarly, Shaw and Bachu (2003) claim that 4470 fields, just over half of the known oil reservoirs in Alberta, Canada, are amenable to CO_2 injection for enhanced oil recovery. Babadagli (2006) states that enhanced oil recovery applied in these reservoirs could translate into an additional 165 billion barrels of oil recovered and over 1 Gt of CO_2 sequestration. Snyder et al. (2008) estimate that at current oil and carbon prices and with current technology, approximately half of this capacity is economically viable.

In this paper, we analyze a price-taking producer's problem of optimally extracting oil from a given field when extraction is subject to a realistic, physical constraint on the rate at which oil can be produced. In this sense, the setup of our model is similar to that in Cairns and Davis (2001) as well as to the manner in which oil-industry practitioners tend to model the producer's decision problem.[4] Cairns and Davis focus, however, on physical constraints that arise during the so-called primary or secondary stages of oil recovery from a field. The primary stage utilizes the reservoir's natural pressure to bring oil to the surface, and during this stage (which typically extracts about 5–20% of the total oil originally present in the reservoir) the decline in pressure over time constrains the rate of oil production. The secondary stage involves injection of water into the reservoir to augment or maintain pressure. During this stage (which typically extracts another 10–30% of the oil) maximum production declines over time because water is progressively less able to "flush" additional oil out of tiny connected pore spaces in the reservoir rock. In both stages, once a producer has developed a field (i.e., put wells and other equipment in place), the decline rate of maximum oil production is essentially exogenous. By contrast, during EOR, which is typically applied in a tertiary stage, the producer is able to affect the rate at which oil production declines by varying the rate of CO_2 injection. This ability effectively endogenizes the physical constraint and complicates the producer's problem in a number of non-trivial ways.

First, CO_2 is not a costless input. Significant up-front investments are required to make production and injection wells suitable for CO_2 use. In addition, maintaining a given injection rate over time requires continuous purchases to make up for the fraction of injected CO_2 that is retained in the reservoir. Separating the remaining fraction that resurfaces with the produced oil, and then dehydrating and recompressing it, is costly as well.

[1] This view was expressed in recent Congressional testimony by George Peridas, Science Fellow at the National Resources Defense Council, who said that CO_2-EOR "has a substantial immediate- to long-term role to play in both increasing domestic oil production in a responsible way, and in sequestering CO_2" (Peridas, 2008). See also the Congressional testimony by William L. Townsend (Townsend, 2007) and the National Petroleum Council report *Hard Truths: Facing the Hard Truths About Energy* (NPC, 2007).

[2] Other methods of enhanced oil recovery exist as well, including injection of steam, nitrogen, methane, and various polymers. Because these methods are not the focus of this paper, we use the term enhanced oil recovery, or EOR, as shorthand for CO_2-enhanced oil recovery.

[3] Producers may use a process called "blowdown" to recover part of the sequestered CO_2 after termination, for use in other EOR projects. Possible leakage of CO_2 during and after the EOR project is a concern as well.

[4] See for example the text by Mian (2002), which is aimed specifically at practicing reservoir engineers. Davis and Cairns (1999) and Thompson (2001) find empirical evidence supporting Adelman's (1990, 1993) contention that physical constraints on the extraction rate always bind in practice.

Second, even at a constant injection rate, the decline in oil recovery over time implies a similar decline in the fraction of injected CO_2 that is retained in the reservoir. Both the producer's revenue stream and cost stream are therefore time varying.

Third, while sequestration of CO_2 currently yields no economic benefits in jurisdictions without carbon emissions restrictions, future regulations of CO_2 emissions in the context of climate-change policies may generate such benefits if EOR projects are allowed to earn credits for units of CO_2 sequestered. The producer will then receive revenue streams from both oil production and CO_2 sequestration.

Fourth, carbon policies affect the producer's revenue and cost streams in multiple ways. While a carbon tax (or an equivalent tradeable permit scheme – hereafter we take this equivalency as understood) effectively reduces the input cost for EOR and increases the value of the CO_2-storage potential of the oil field, the incidence of the tax on the price of oil reduces the value of the traditional use of the asset.

Fifth, in addition to these economic tradeoffs, fluid-dynamic interactions of CO_2, water, and oil inside the reservoir give rise to a further, physical tradeoff faced by the producer. Whereas injecting pure CO_2 maximizes oil recovery from the area of the reservoir that the CO_2 sweeps through, that area itself may be small, as pure CO_2 tends to "finger" or "channel" between injection and production wells, bypassing some of the oil. In comparison, injecting pure water increases the area that is swept, but reduces recovery from that area. Reservoir-engineering studies indicate that both oil recovery and CO_2 sequestration are maximized when a mix of CO_2 and water is injected (though the CO_2 fraction that maximizes oil recovery typically differs from that which maximizes sequestration).[5]

These five factors combined present the producer with a problem of co-optimizing oil production and CO_2 sequestration. Three earlier studies have investigated different components of this problem. Kovscek and Cakici (2005), for example, use a reservoir simulation model to investigate the impact of varying the fraction of CO_2 in the injection stream. However, their analysis is limited by the choice of objective function (an arbitrary weighted sum of cumulative oil recovery and cumulative CO_2 sequestration) and by the imposition of an exogenous termination time. McCoy and Rubin (2009), on the other hand, explicitly optimize the termination time to maximize an EOR project's net present value at different oil and CO_2 prices. However, they do not allow for an endogenously selected fraction of CO_2 in the injection stream, assuming instead that it is exogenously set at 100%. Finally, in a recent working paper, Bandza and Vajjhala (2008) compare four CO_2 injection strategies that producers might choose: focus on profits from oil production alone; on profits from sequestration alone; on profits from oil production with sequestration profits considered only later, as an "afterthought"; or on combined profits from oil production and sequestration from the outset. While Bandza and Vajjhala pay attention to the role of prices, there is no explicit model of dynamic optimization – a shortcoming shared by all three papers. By considering how the operator of an EOR project might maximize the project's net present value through its choice of the optimal time path of the fraction of CO_2 in the injection stream, as well as the optimal time to terminate the project, our paper addresses this shortcoming in the extant literature.

We start by analyzing the dynamic co-optimization problem theoretically, using an optimal control framework in which the CO_2 fraction in the injection stream is the control variable, the stock of oil remaining in the reservoir is the state variable, and the terminal time is determined by a transversality condition. Using this model, we find that the optimal CO_2 fraction will typically decline over time, and may eventually drop to zero before it becomes optimal to terminate the project.

Next, we conduct a series of simulations of the model, after calibrating it to an ongoing EOR project, namely the Lost Soldier-Tensleep project in Wyoming. The simulations generate time paths of CO_2 injection and implied time paths of oil production and CO_2 sequestration, for a range of oil prices and carbon taxes. A key finding is that cumulative sequestration is more responsive to the oil price than to the price of carbon.

[5] See, e.g., Al-Shuraiqi et al. (2003), Jessen et al. (2005), Juanes and Blunt (2006), Guo et al. (2006), and Trivedi and Babadagli (2007).

2. The model

Our model of oil production is based on the physical reality that input injections (water, gas, or some mixture of the two) must balance with fluid output (oil, water, and gas).[6] In addition, the rate of oil production is linked to remaining reserves by the so-called "decline curve."[7] This relation specifies output as a particular fraction of remaining reserves, where that fraction is itself linked to the fraction of CO_2 in the injection stream. Upon specifying the relation between rate of CO_2 injection and oil production we may write down the formal dynamic optimization model, which we then use to describe the time path of CO_2 injection. Ultimately, this allows us to describe the rate of CO_2 that is sequestered at every point in time, and thereby to determine the total amount sequestered.

We begin with some notation. Let the rate of CO_2 injection at time t be $q_i^c(t)$, and the rate of water injection be $q_i^w(t)$. We assume the total rate of injection, $q_i \equiv q_i^c(t) + q_i^w(t)$, is constant across time. This reflects the fact that CO_2-EOR projects are usually operated at "minimum miscibility pressure" (MMP), which is the minimum pressure required to make the CO_2 mix with the oil. Maintaining that pressure requires a roughly constant overall injection rate.

Let the rates of oil, CO_2, and water production at time t be $q_p^o(t)$, $q_p^c(t)$ and $q_p^w(t)$. The physical requirement that, at constant pressure conditions, the outflow of liquids from the reservoir must equal the inflow of liquids then implies that $q_p^o(t) + q_p^c(t) + q_p^w(t) = q_i^c(t) + q_i^w(t)$ at each point in time.

In principle, the produced CO_2 could be vented or recycled. Let the price of a unit of newly purchased CO_2 equal w_m and the unit cost of recycling CO_2 equal w_r. We assume that $w_m > w_r$, so that it is always cheaper for the firm to recycle than to vent. In practice, the purchase price of CO_2 is in fact several times higher than the cost of recycling.[8] Thus, firms undertaking EOR do generally recycle CO_2. Importantly, the presence of a carbon price τ does not change the relevant comparison: the cost of a newly purchased unit becomes $w_m - \tau$ (as the seller of the CO_2 avoids the carbon tax or receives a credit for sequestering), while the opportunity cost of recycling becomes $w_r - \tau$ (as venting would obligate the producer to pay the carbon tax or purchase a credit). Note, however, the implicit assumption that competition between CO_2 sellers will induce them to pass on the full savings on the tax or value of the credit to the EOR buyer. As we discuss in the conclusion, aggregate CO_2 emissions currently far outstrip all estimates of aggregate EOR sequestration, so this assumption seems reasonable.

If indeed $w_m > w_r$, then all produced CO_2 is recycled, so $q_r^c(t) = q_p^c(t)$, and total CO_2 injection is the sum of new purchases and recycling, so $q_i^c(t) = q_m^c(t) + q_r^c(t)$. An immediate implication is that CO_2 purchases are needed only to make up for CO_2 that is *not* recovered from the production stream and recycled, but instead remains in the reservoir. Denoting the rate of CO_2 retention by q_s^c, we therefore have $q_m^c(t) = q_s^c(t)$.

Denote the amount of underground CO_2 at time t by $S(t)$. By construction, $q_s^c(t)$ equals the change in $S(t)$, so

$$\dot{S}(t) = q_s^c(t) = q_i(t) - q_p^c(t). \tag{1}$$

One may interpret $S(t)$ as the amount of CO_2 sequestered as of time t, and so we shall often refer to $q_s^c(t)$ as the rate of sequestration. This rate depends on the linkage between injected CO_2 and produced

[6] Reservoir engineers refer to this as "material balance." It should be noted that this requirement applies at the temperature and pressure conditions that obtain inside the reservoir. At these conditions, CO_2 exists in a highly compressed, "supercritical" state and behaves much like a liquid.

[7] Fetkovitch (1980) provides an in-depth discussion of decline curves, and of the justification for their widespread use in predicting oil production from reservoirs.

[8] The purchase price w_m covers both flow costs and a return on capital required to extract and compress CO_2 and to then transport it to the EOR site. Extraction currently occurs either from natural geological deposits of CO_2 or from the production stream of various manufacturing processes (e.g., natural-gas processing, ammonia production, fertilizer production) that generate CO_2 as a by-product. Included also in the price is a pipeline tariff that provides a return on investment to pipeline owners. The recycling cost w_r, on the other hand, covers just the flow cost of operating a (usually on-site) plant that separates the produced CO_2 from the water and oil that it is mixed in with, and then recompresses the CO_2 to the required injection pressure. Most of this cost is that of electricity required to power the recycling plant. As discussed in more detail in footnote 10, we treat the EOR operator's up-front capital costs as already sunk at time 0.

oil. We assume that the fraction of oil production displaced by CO_2 (as opposed to that displaced by water) is proportional to the fraction of CO_2 in the total injection stream. We also assume that sequestered CO_2 takes up the underground space vacated by the oil that it displaces. To simplify the exposition, units of oil are chosen such that in the reservoir, vacating the space taken up by one unit of oil creates space for sequestering exactly one unit of CO_2. Moreover, since we are focusing on an individual firm and a particular oil reservoir, we may normalize so that $q_i = 1$. As a result, we have

$$q_s^c(t) = c(t)q_p^o(t), \tag{2}$$

where $c(t) \equiv q_i^c(t)/q_i(t) = q_i^c(t)$ is both the fraction of CO_2 in the total injection stream and, by our normalization, the rate of CO_2 injection. This variable will play a central role in the analysis to follow.

At any point in time, the amount of recoverable oil is $R(t)$; we write the initial amount of oil at the moment the EOR project is undertaken as R_0. In keeping with the physical reality of oil recovery, we assume that the rate of production can be described by a decline curve: $q_p^o(t) = \delta R(t)$. In our setting, however, the ratio of output to reserves – which plays the role of the decline rate – is linked to the CO_2 fraction in the total injection stream: $\delta = \delta(c(t))$. We therefore have the relation[9]

$$\dot{R}(t) = -q_p^o(t) = -\delta(c(t))R(t). \tag{3}$$

Consistent with results from the reservoir-engineering studies cited in the introduction, we assume that the $\delta(c)$ function relating the CO_2 fraction in the injection stream to the decline rate is concave, with an interior maximum. If only water is used (termed a "waterflood"), the decline rate is $\delta_w \equiv \delta(0) > 0$. If only CO_2 is used (termed a "pure CO_2 flood"), the decline rate is $\delta(1)$. We assume, consistent again with reservoir-engineering studies, that $\delta(1) > \delta_w$. In light of the concavity of $\delta(c)$, $\delta'(0) > 0 > \delta'(1)$.

The economic environment depends on three ingredients: the price of oil, p; the carbon tax, τ; and operating costs.[10] We assume that all operating costs other than those of CO_2 purchases and CO_2 recycling are tied to the total rates of fluid injection and production. Recall our assumption that the total rate of fluid injection, q_i, is constant over time; by material balance, the same must be true of the total rate of fluid production. Operating costs tied to either the total rate of fluid injection or the total rate of fluid production must then be constant as well. We denote these atemporal operating costs by F. Absorbed in F are also costs of recycling and pumping water, which, though not strictly constant over time, are sufficiently small to treat as such.[11] Since the combustion of oil generates CO_2 as a by-product, it seems reasonable to expect that there will be a tax liability embedded within the market price. To facilitate further discussions of the role played by the carbon tax, we denote the induced tax liability for a one-dollar increase in the carbon tax by β. This parameter combines tax incidence effects with unit conversions associated with the transformation of a unit of produced oil into carbon units.

Dropping time arguments where no confusion will result, we can then write the firm's profit rate as

$$\pi = [p - \beta\tau]q_p^o - w_m q_m^c - w_r q_r^c + \tau q_s^c - F, \tag{4}$$

[9] In principle, we could write $q_p^o(t) \leq \delta(c(t))R(t)q_i$, allowing for the possibility that the operator might choose an injection rate $q_i < 1$ and thereby produce at less than the maximum physically feasible rate. It is easy to show, however, that the producer's problem then becomes linear in the control q_i, and so the solution is bang-bang. In other words, the upper bound on production always binds if the producer chooses to produce at all.

[10] Because our focus in this paper is on the optimal operation of a CO_2 flood, we abstract from both up-front investments required to make an oil field "CO_2 ready" and post-termination costs of monitoring for possible leakage from the reservoir. Up-front investments include changes in well equipment, additions of metering equipment and pipelines in the field, and the construction of a CO_2 recycling plant. We discuss likely implications of introducing up-front investment costs in the concluding section of the paper. As for monitoring costs, Benson (2006) notes that a large variety of proven monitoring techniques exist, but that no decisions have yet been made about which of them to deploy on a routine basis. Some of these techniques will focus on possible leakage from wellbores, making their cost proportional to the number of wells in a project; others will focus on possible leakage through the cap rock, making their cost proportional to the area of the underground CO_2 plume. Benson et al. (2005) estimate total life-cycle monitoring costs for EOR projects, including periodic seismic surveys, wellhead pressure and injection-rate monitoring, and surface CO_2 flux monitoring, to be very small: on the order of \$0.10–0.30 per ton of CO_2.

[11] Because water is heavier than liquid or supercritical CO_2, gravity operating on the water column in the injection well costlessly provides much of the required injection pressure. Also, any supplemental water required to make up for losses in the reservoir is usually drawn from a nearby saline aquifer, and is therefore free except for a small pumping cost.

or, using (2), (3), the identities $q_m^c = q_s^c$ and $q_r^c = q_i - q_s^c$, and our normalization $q_i = 1$, as

$$\pi = [p - \beta\tau]\delta(c)R - [w_m - \tau - w_r]c\delta(c)R - w_r c - F. \tag{5}$$

In order to keep our baseline model tractable, we have abstracted from many real-world complications that operators of CO_2-EOR projects are faced with. These complications concern (i) operational decisions other than the fraction of CO_2 in the injection stream, (ii) the termination decision and post-termination options, and (iii) potential CO_2 leakage.

Although we model the operator as injecting a direct mixture of CO_2 and water, in reality, operators cycle between pure CO_2 and water – a process referred to as "water alternating gas" (WAG). The variable $c(t)$ should therefore be interpreted as the share of CO_2 in the total injection stream averaged over such cycles.[12] By working with this average, however, we ignore the cycling frequency as a choice variable that may affect field performance.[13] Other operational decisions from which we have abstracted target the tendency of CO_2 to move through highly permeable streaks or zones in the reservoir rock, bypassing oil in less permeable parts, or to float to the top of a reservoir, overriding oil in lower layers. In order to reduce such bypassing and overriding, operators may attempt to plug permeable reservoir zones by injecting surfactant foams or gel polymers into them, perforate only the lower or upper sections of well tubes, shut off mistakenly perforated segments with sand or cement, or completely shut down wells that produce too much CO_2 relative to oil. These so-called "conformance improvements" tend to increase both oil recovery and CO_2 sequestration, however, and therefore do not give rise to any tradeoff between the two. Moreover, Masoner and Wackowski (1994), reviewing conformance improvements and other management decisions made for the huge (and still ongoing) Rangely CO_2-EOR project in Colorado, note that "WAG decisions are the number one method to improve the oil cut and CO_2 utilization. It provides the cheapest and most effective short and long term means to achieve improving the sweep."[14] With respect to decisions surrounding termination of the project, we implicitly treat the project as going full-steam until time T. In reality, oil fields are typically arranged into sub-areas called "patterns," centered on a single injection well and bordered by production wells. Because of heterogeneities of the reservoir rock, patterns in different parts of a field will tend to perform differently. Operators should therefore be expected to shut down under-performing patterns first, and may only gradually bring the entire project to a full stop. In principle, this can be dealt with by interpreting our model as applying to a single project pattern. Such an interpretation ignores, however, the fact that operating costs (mainly of personnel and CO_2 recycling) are partially shared across all patterns in a field, so that the per-pattern cost of still operating patterns may increase somewhat towards the end of a project's life.

Furthermore, although we implicitly assume that the operator shuts in both injection and production wells at time T, in reality the operator may choose to initially shut in only the injection wells. Doing so will cause a gradual release of the reservoir's pressure to below MMP. As a result of this so-called "blowdown" process, some of the sequestered CO_2 mixed in with oil remaining in the reservoir will come out of that mixture, and make its way to the production wells together with some of the remaining oil. In addition to selling the additional oil, the operator can then re-compress and sell the CO_2 to another project.[15] A final complication that is not included in our model is potential CO_2 leakage to the atmosphere, either through fugitive emissions from surface equipment in the course of EOR operations or through seepage out of the reservoir, possibly many years after operations have

[12] In the engineering literature, this share is commonly expressed in terms of the ratio of water to CO_2 injections, referred to as the "WAG ratio." Denoting this ratio W, mathematically $c = 1/(1+W)$.

[13] In the engineering literature, this frequency is commonly expressed in terms of the volume of water or CO_2 "slugs" injected in any single cycle. All else equal, larger slugs take longer to inject, thereby lengthening each cycle and reducing the cycle frequency. See Jessen et al. (2005) for a simulation study of how the cycle frequency may affect both oil recovery and cumulative CO_2 sequestration, finding little effect.

[14] In seeming contradiction to this statement, Cakici (2003) and Kovscek and Cakici (2005) find significant improvements in CO_2 sequestration, without much loss in oil recovery, from a management strategy focused on temporary shutdowns of wells that produce high ratios of CO_2 to oil. They also note, however, that the reservoir used to calibrate their simulation study "is not a prototypical candidate for CO_2-EOR," because it cannot be pressurized to MMP without risking fracture of the cap rock. Perhaps because of this, the reservoir is found to produce more oil with a pure waterflood than with any mix of CO_2 and water.

[15] The interested reader is directed to our working paper Leach et al. (2010), in which we derive and simulate an alternative model formulation which includes an endogenous blowdown decision. We show that the implications of including blowdown are intuitive and do not affect the qualitative results.

ceased. This issue is of great concern to policy makers, both because of public perceptions that geological sequestration may be unsafe, and because any CO_2 credits awarded to EOR operators should be discounted to the extent that sequestration is only temporary.[16] The best available evidence, however, suggests that leakage rates from either source are likely to be negligibly small. Aycaguer et al. (2001) estimate, based on data from the Willard-San Andres EOR project in West Texas, that over the project's 40-year expected lifetime "releases associated with venting, leaks, or discharges from process vents, along with emissions during maintenance activities, accidents, and system upsets" will cumulatively release about 3% of cumulative CO_2 purchases. As for CO_2 seepage from reservoirs, either through imperfect seals in the caprock or through corrosion or other imperfections in wellbores, a wide range of studies reviewed in the Intergovernmental Panel on Climate Change's recent report on carbon capture and storage (IPCC, 2005) estimate such seepage to occur at annual rates of just 0.00001–0.001% of the sequestered CO_2 stock.[17,18]

3. Analysis

In the context of our model, the firm's choice variables are the time path of the CO_2 fraction $c(t)$ in its injection stream (or equivalently, by our normalization that $q_i = 1$, the time path of the CO_2 injection rate $q_i^c(t)$), the terminal time T at which to cease injection altogether, and the resulting terminal state $R(T)$. Note that, by Eqs. (2) and (3), the firm's choice of $c(t)$ directly determines both its oil production rate $q_p^o(t)$ and CO_2 sequestration rate $q_s^c(t)$. The firm's goal is to maximize its present discounted value, subject to the state equation (3), the initial value of the state variable, R_0, the constraints $0 \leq c \leq 1, R \geq 0$, and transversality conditions.[19] To solve this dynamic optimization problem, we first define the current-value Hamiltonian

$$H = \pi - mq_p^o = Y\delta(c)R - Zc\delta(c)R - w_r c - F - m\delta(c)R, \tag{6}$$

where m is the current-value multiplier (shadow price) associated with a unit of oil *in situ* and, to save on notation, we have defined $Y \equiv p - \beta\tau$ and $Z \equiv w_m - \tau - w_r$. The optimal path of extraction satisfies the maximum principle, which consists of the state equation, an equation for identifying the optimal extraction rate at a given point in time, and an equation of motion for the shadow price. If the optimal extraction rate is described by an interior solution, we have

$$H_c = (Y - Zc - m)\delta'(c)R - Z\delta(c)R - w_r = 0, \tag{7}$$

where $H_c = \partial H/\partial c$. Irrespective of whether the optimal value of c is described by an interior solution, the state equation is given by (3), and the equation of motion for the shadow price is

$$\dot{m} = rm - (Y - Zc - m)\delta(c). \tag{8}$$

The right-hand side of Eq. (8) differs from the standard Hotelling representation by virtue of the second set of terms. These terms capture the fact that current carbon injections, by reducing the stock

16 See Herzog et al. (2003) for an analysis of how the appropriate discount factor depends not just on leakage rates, but also the expected path of the carbon price and the time (if any) at which a carbon-free backstop energy source might effectively cap that price.

17 Zhou et al. (2005), for example, estimate that *cumulatively* over the next 5000 years, CO_2 seepage through the 1000 or so wellbores of the Weyburn EOR project in Canada will amount to maximally 0.14% of the CO_2 in place at the end of EOR operations, with cumulatively 0.001% the more representative estimate. An outlier is Klusman's (2003) study of the Rangely EOR project in Colorado, which bounds annual seepage at between 0.00076 and 0.02% of the stock. The study has difficulty, however, distinguishing seepage from CO_2 generated by soil microbial oxidation of methane, and notes that the lower bound of on actual seepage is "likely even lower."

18 See Benson (2006) for further discussion of the literature on leakage, including issues of monitoring and verification. Benson remarks on an unfortunate tendency in this literature to confuse normative and descriptive results. Hepple and Benson (2005), for example, estimated an annual leakage rate of 0.1% to be "acceptable" for geological storage to remain effective as a climate-change mitigation technique, but later studies have misinterpreted this normative result as descriptive, i.e., as an estimate of the actual leakage rate.

19 The variable S, representing the accumulated stock of CO_2, does not affect the firm's marginal conditions as long as complications such as blowdown or eventual leakage are not considered. It is, however, a variable of interest, and we compare its evolution and terminal value $S(T)$ across various price scenarios in Section 5.

and thereby moving the firm further down its oil production decline curve, reduce the capability of future injections to enhance oil production. The value associated with this induced diminution of the future production rate is the product of the marginal impact of the production rate on the current-value Hamiltonian $(Y - Zc - m)$ and the current decline rate $(\delta(c))$.

Because the end time is free, the value of the current-value Hamiltonian at the terminal time T must be zero. As the end state is free, the product of the shadow price and the state variable at the terminal time must also be zero: $m(T)R(T) = 0$. As the extraction rate is proportional to the stock, we infer from Eq. (6) evaluated at T that the firm optimally terminates the project when the profit rate drops to zero. From Eq. (4), this will be true when combined flow revenues from oil extraction and CO_2 sequestration just barely cover the flow costs of CO_2 purchases and recycling, as well as the flow fixed costs. Because those fixed costs are positive, however, the equivalent equation (5) implies that the terminal stock must be positive, and thereby that the terminal value of the shadow price must be zero.

We now turn to a discussion of the time path of injection. Assuming an interior solution over an interval, we may time-differentiate (7) to get

$$\dot{c} = \frac{[-H_{cm}\dot{m} - H_{cR}\dot{R}]}{H_{cc}},$$

where $H_{cx} = \partial^2 H/\partial c \partial x$, $x = m$, c or R. From (7), we see that $H_{cm} = -\delta'(c)R$ and $H_{cR} = (Y - Zc - m)\delta'(c) - Z\delta(c) = w_r/R$ (where we use (7) to extract the last relation). Combining these observations with the state equation, we get

$$\dot{c} = \frac{[\delta'(c)R\dot{m} + w_r\delta(c)]}{H_{cc}}. \tag{9}$$

At an interior solution, the denominator is negative and the second term within square brackets is positive. It follows that injection is falling at any moment where $\delta'(c)\dot{m}$ is positive; if it is negative, the sign of $\dot{c}$ is ambiguous.

To further explore the time path of c, we combine (7) and (8) to get

$$\delta'(c)R\dot{m} + w_r\delta(c) = [r\delta'(c)m - Z\delta(c)^2]R. \tag{10}$$

Comparing (9) and (10), it is apparent that a sufficient condition for $\dot{c}$ to be negative is for the right-hand side of (10) to be positive. This will occur, for example, if $\delta' > 0$ and Z is not large and positive, or if Z is negative and large in magnitude. Heuristically, $\delta' > 0$ is consistent with the notion of restraining current production so as to allow rents to rise over time, which seems plausible. For Z to be small is a bit less obvious. Recall that $Z = w_m - \tau - w_r$, and that by assumption $w_m - w_r > 0$. If $w_m - w_r$ is small, which is the case in our simulations and seems to be the empirically important case, then Z will be small irrespective of the size of the carbon tax. On the other hand, if the carbon tax is particularly large, then Z will be negative. On balance, then, the right-hand side of (10) will be positive in a range of cases that seem empirically relevant. As such, the rate of injection will commonly be declining.

It is instructive to compare these results to the traditional Hotelling model. There, one expects extraction to fall over time so as to inter-temporally equate discounted rents. That effect is complicated somewhat if extraction costs shift up over time, as happens in our model, but the fundamental logic remains. However, the logic is fundamentally changed if the extraction problem is subject to important constraints. Our problem entails two such constraints: the rate of oil production is subject to a decline curve effect, which ties the maximal extraction rate to the stock of remaining oil; there is also a physical limit on the amount of fluids (oil, water or CO_2) that can be extracted at any point in time. Because of the decline curve effect, production necessarily falls over time, though not because the firm has managed to equate discounted rents. Indeed, with constant price, an unconstrained firm would extract as rapidly as possible.

Why then does the rate of CO_2 injection fall over time? There are three reasons. First, one can view the sequestration component of the co-optimization problem as an exhaustible resource problem of sorts. There is a finite amount of pore space available in which to sequester. Because the carbon price τ associated with sequestration activity does not increase over time, sequestration is more valuable early on. Accordingly, the firm wishes to sequester as rapidly and as early as possible. Second, the decline curve effect implies a reduction in the marginal product of CO_2 as time goes by; all else equal this would induce the firm to lower its usage of CO_2 on a moment-by-moment basis. The third reason is related: the second motive is reinforced for combinations where $\delta'(c) > 0$, since, as the injection rate c declines, the marginal product of CO_2 is depressed still further. Again, all else equal this effect would push the firm to inject most rapidly in early periods.[20] Together, these three effects induce the firm to select a time path for c that decreases over time.

The preceding discussion focuses on interior solutions. While these will be common, there are circumstances under which corner solutions obtain. We now discuss those conditions. First, suppose the optimal rate of CO_2 injection is zero (i.e., it is optimal to undertake a waterflood); in that case $H_c \leq 0$ when evaluated at $c=0$. The condition of interest is

$$(Y - m)\delta'(0)R - Z\delta_w R - w_m \leq 0.$$

Because m and R do not change discontinuously, if this condition holds with strict inequality at a particular moment t, it must hold for an interval of time following t. Accordingly, during this interval the optimal level of c remains equal to zero. It follows that during this interval

$$\dot{H}_c = -\delta'(0)[R\dot{m} - (Y - m)\dot{R}] - Z\delta_w \dot{R},$$

or upon using (3),

$$\dot{H}_c = -\{\delta'(0)[\dot{m} + \delta_w(Y - m)] - Z{\delta_w}^2\}R. \tag{11}$$

Combining (8) and (11), taking note of the fact that $c=0$, we deduce that

$$\dot{H}_c = -[rm\delta'(0) - Z{\delta_w}^2]R. \tag{12}$$

The important thing to note here is that for negative values of Z, or values of Z that are positive but relatively small in magnitude, the right-hand side of (12) will be non-positive; as we noted above, this restriction does not seem to be terribly demanding. In such a scenario, once H_c becomes negative, it tends to stay negative. We conclude that it will be typical for the corner solution $c=0$ to remain in effect once it is initiated.

Now suppose the optimal rate of CO_2 injection is one (i.e., it is optimal to undertake a pure CO_2 flood); in that case $H_c \geq 0$ when evaluated at $c=1$. The condition of interest is

$$(Y - Z - m)\delta'(1)R - Z\delta(1)R - w_m \geq 0.$$

As with the $c=0$ corner solution, if this condition holds with strict inequality, it must apply for an interval of time; during that interval we have

$$\dot{H}_c = -[rm\delta'(1) - Z\delta(1)^2]R. \tag{13}$$

As noted above, the only way this corner solution can obtain is if Z is negative and large in magnitude. On the other hand, $\delta'(1) < 0$. Thus, depending on the relative magnitudes of Z and m, H_c can either be rising or falling. Importantly, as m is likely to fall over time, eventually $\dot{H}_c$ will become negative. It follows that the pure CO_2 flood cannot last indefinitely: at some point, it will be optimal to adopt an interior solution.

Our model thus predicts that under most conditions, from the point at which a CO_2 flood is initiated, CO_2 injection will be non-increasing over time until it reaches zero. After this, a pure waterflood will continue until the flow profits are equal to the flow fixed costs, at which time extraction activity ceases. This endogenous endpoint occurs when the shadow value reaches zero. The initial value of carbon injection, the rate of decline of injection over time, the point at which a pure waterflood begins and ends,

[20] This third effect works in the opposite direction for combinations where $\delta'(c) < 0$. Such a scenario can arise when the unit price for sequestered carbon is sufficiently high; in our simulations this only takes place at implausibly large prices.

Table 1
Baseline parameter values.

I	1	Overall rate of injection and production ($\times$1 million barrels)
R_0	1	Initial stock of oil in reservoir ($\times$1 million barrels)
w_m	4	Per-barrel cost of purchased CO_2
w_r	1	Per-barrel cost of separating and recycling "leaked" CO_2
F	0.1	Fixed costs ($\times$\$1 million)
β	2.2	Incidence of the carbon tax on the oil producer
r	0.05	Discount rate
δ_w	0.06	Intercept of $\delta(c)$ function
δ_1	0.20	First coefficient of $\delta(c)$ function
δ_2	0.16	Second coefficient of $\delta(c)$ function

and the total amounts of oil production and carbon sequestration will be determined by field-specific physical characteristics that determine the δ function and the initial state R. Of course, they will also be affected by the oil price and carbon tax, which define the economic environment.

4. Simulation framework

In order to add greater context to the results derived above, we have implemented the model numerically to solve for optimal time paths of carbon injection and oil production at various combinations of the oil price and carbon tax. Specifically, we compute the solution for scenarios with oil prices of \$50, \$100, \$150, and \$200 per barrel (bl), taking \$100/bl as our baseline price,[21] and for carbon taxes of \$0, \$40, \$80, and \$120 per tonne of CO_2 (tCO_2), taking the absence of any tax as our baseline. Table 1 shows the baseline parameter values of the numerical model.

All quantity flows are measured in millions of "reservoir" barrels (rb) per year (1 reservoir barrel = 42 gallons (US) $\approx 0.16\,m^3$ at the temperature and pressure conditions that obtain inside the reservoir). Overall injection I is normalized to 1 million such barrels. The initial stock of recoverable oil in the reservoir, R_0, is set at 1 million barrels as well. For comparison, the Lost Soldier-Tensleep (LSTP) EOR project in Wyoming injects about 44 million barrels per year, and extrapolating the decline curve for its oil production since starting the CO_2 flood suggests that ultimately about 36 million barrels of oil would be recovered over the course of that flood were it to be continued forever. In effect, then, our simulation applies a scaling factor of about 1/40 to the LSTP project.

The various cost parameters of the model are based on a variety of sources, including data presented in McCoy (2008) and EIA (2007), as well as personal communication with industry experts.[22]

As noted above, the parameter β describes the effect that a carbon tax or equivalent cap-and-trade policy will have on the oil price received by all producers, including those that apply EOR. This parameter combines the tax incidence on producers with unit conversions that transform a unit of produced oil into carbon units. Assuming that the tax applies worldwide, the elasticities relevant to calculating its incidence are those of world oil demand and supply.[23] Lastly, the parameters of $\delta(c)$ function are based on a combination of production experience at LSTP and simulation results in the literature. As we discuss in Appendix A, these data are consistent with a quadratic function[24]

$$\delta(c) = 0.06 + 0.2c - 0.16c^2.$$

[21] The Energy Information Administration's *Annual Energy Outlook 2009* projects \$130 as its reference-case oil price for 2030, with \$50 and \$200 as its low and high cases.

[22] In particular, Charles Fox of Kinder Morgan, Inc., and Mark Nicholas of Nicholas Consulting Group. A discussion of the derivation of our baseline cost parameters is contained in Appendix A.

[23] Because estimates of demand and supply elasticities for oil, and thereby of the tax incidence, are subject to considerable uncertainty, we have performed some sensitivity analysis on this parameter. We find that, even in the two extreme cases where the incidence is zero, so that β = 0 also, and where the incidence is 100%, so that β = 4, our results are essentially unchanged.

[24] Although the reservoir-engineering studies cited in footnote 1 suggest that the $\delta(c)$ function is concave, its precise shape for any given reservoir is likely to strongly depend on geological properties such as permeability, thickness, and heterogeneity of the reservoir rock. Accordingly, the particular parameterization used here is best viewed as only illustrative.

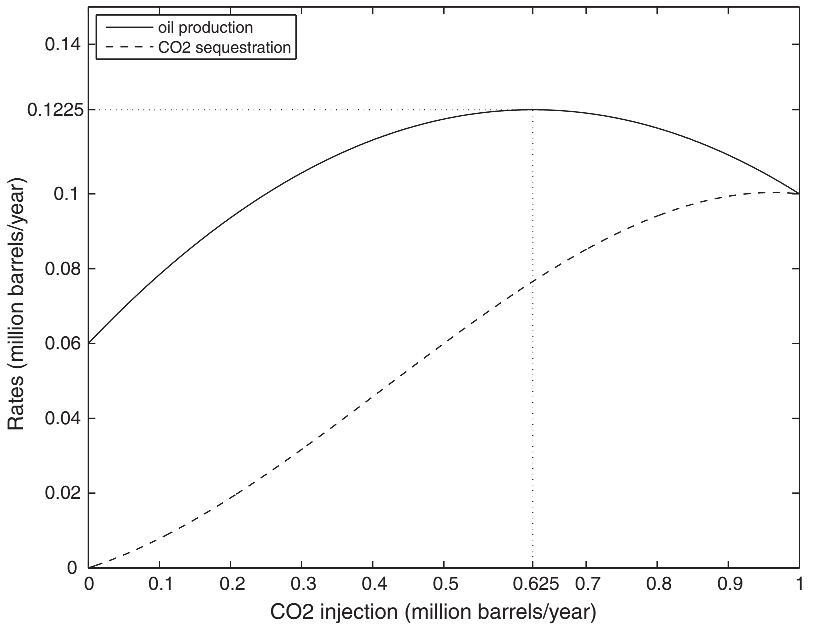

Fig. 1. Initial oil production and CO_2 sequestration as a function of the initial CO_2 injection rate.

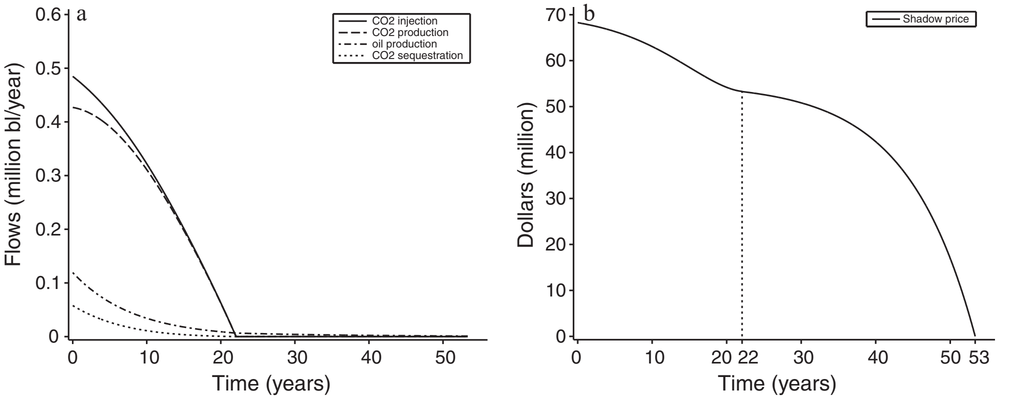

Fig. 2. CO_2 injection, production, and sequestration, and oil production (a) and the shadow price (b) over time.

Fig. 1 shows the initial rates of oil production, $\delta(c)R_0$, and CO_2 sequestration, $c\delta(c)R_0$, implied by this parameterization. The important thing to note is that initial CO_2 injection rates up to 0.625 million barrels/year increase both oil recovery and CO_2 sequestration; it is only beyond this critical value that CO_2 injection becomes counter-productive to oil recovery, while still increasing sequestration. In fact, these relationships continue to hold also after time 0: oil production at any time $t > 0$ equals $\delta(c)R(t)$, while CO_2 sequestration equals $c\delta(c)R(t)$, so that Fig. 1 continues to apply up to a vertical scaling by $R(t)/R_0$.

5. Simulation results

The first element of behavior that we wish to define is the optimal extraction and sequestration path for our benchmark assumptions. Here, we use an oil price of $100 with no carbon tax. Panel (a) of

Fig. 2 shows the optimal paths of CO_2 injection (q_i^c), CO_2 production (q_p^c), oil production (q_p^o), and flow CO_2 sequestration (q_s^c).

The optimal initial injection rate is 0.485 million barrels/year, somewhat smaller than either the instantaneous oil-production maximizing rate of 0.625 or the myopic profit-maximizing injection rate of 0.582. This reflects the producer's tradeoffs of current against future extraction, and of oil revenues against CO_2 injection costs. Note also that a large fraction (initially about 88%) of the injected CO_2 resurfaces with the produced oil and must be recycled. As oil production declines over time from its initial rate of 0.119 million barrels/year, the producer's revenues decline as well, as does CO_2 sequestration in the space vacated by the oil. As a result, CO_2 production, and thereby recycling costs, would increase over time even if the producer chose to hold CO_2 injection constant. This changing balance between oil revenues and recycling costs makes it optimal for the producer to instead gradually reduce the injection rate over time, as predicted by the theory.

After 22 years, the optimal CO_2 injection rate drops to zero, at which point the producer switches to a pure waterflood, thereby completely avoiding CO_2 purchase and recycling costs. From that point in time forward, profits consist of the (declining) oil revenues less fixed costs. These remain positive for another 31 years, after which the field is shut down.

Panel (b) of the figure shows the corresponding path of the shadow price. Consistent with our analysis in the previous section, the shadow price declines throughout, reaching its terminal value of zero after 53 years.

5.1. Effects of oil price

Our investigation of the comparative dynamics of the model starts with the effect of higher oil prices. Panel (a) of Fig. 3 shows how the optimal CO_2 injection path changes as the oil price level is raised from \$100 to \$200/bl. Panel (c) of Fig. 3 shows how the optimal time path of oil production changes with price; for reference, we also plot the time path under a pure water flood. The higher resulting oil revenues make it optimal to initially raise the CO_2 injection rate, bringing it closer to the output-maximizing level. However, because even at the baseline price of \$100, initial revenues are already very high relative to CO_2-related costs, baseline oil production is already very close to its revenue-maximizing rate at each point in time. Raising the price therefore has a negligible effect on the oil production path, as is evident from panel (c) of the graph; it also has a negligible effect on cumulative oil production, as shown in panel (d). Nevertheless, the fact that the oil is produced with a more CO_2-rich injection mix implies that cumulative sequestration over the productive lifetime of the field increases, as shown in panel (b).

Fig. 4 shows oil supply and resulting CO_2 sequestration as a function of the oil price. Panel (a) shows cumulative levels of both, whereas panel (b) shows the annualized equivalent.[25] Note that CO_2 sequestration supply drops to zero at at a price of \$12/bl, below which incremental oil revenues from CO_2 injection no longer justify the higher variable costs.[26] At lower prices, the producer therefore optimally operates the field as a waterflood, resulting in zero sequestration and in slower oil extraction. Once the price drops below \$1.70/bl, oil revenues no longer cover the operating costs of a waterflood either, making it optimal to not operate the field at all.

Note also that at prices of \$50/bl and higher, changes in the price of oil have almost no effect on oil output: at the baseline price of \$100, the elasticity of cumulative output is 0.01, while that of annualized output is 0.04. This point is consistent with the result shown in panels (c) and (d) of Fig. 3: at these high oil prices, operating costs become so small relative to oil revenues that the optimal oil extraction path is very close to the optimal path that would obtain if costs were zero (i.e., the revenue-maximizing path). Nevertheless, as shown in panels (a) and (b) of Fig. 3, higher oil prices do induce substantially higher rates of CO_2 injection, and thereby sequestration. As a result, the sequestration

[25] The annualized values are calculated as the constant rate $\bar{s}$ or $\bar{q}$ that, when multiplied by the relevant price τ or p, would over an infinite time horizon yield the same present value as the actual, time-varying rate $s(t)$ or $q(t)$. That is, $\bar{s}$ is implicitly defined by $\int_0^\infty e^{-rt}\tau\bar{s}\,dt = \int_0^\infty e^{-rt}\tau s(t)\,dt$ and $\bar{q}$ is defined analogously.

[26] Recall that we abstract from up-front capital costs associated with CO_2 injection. Implicitly, we assume that at time 0 these costs have already been incurred, and are sunk.

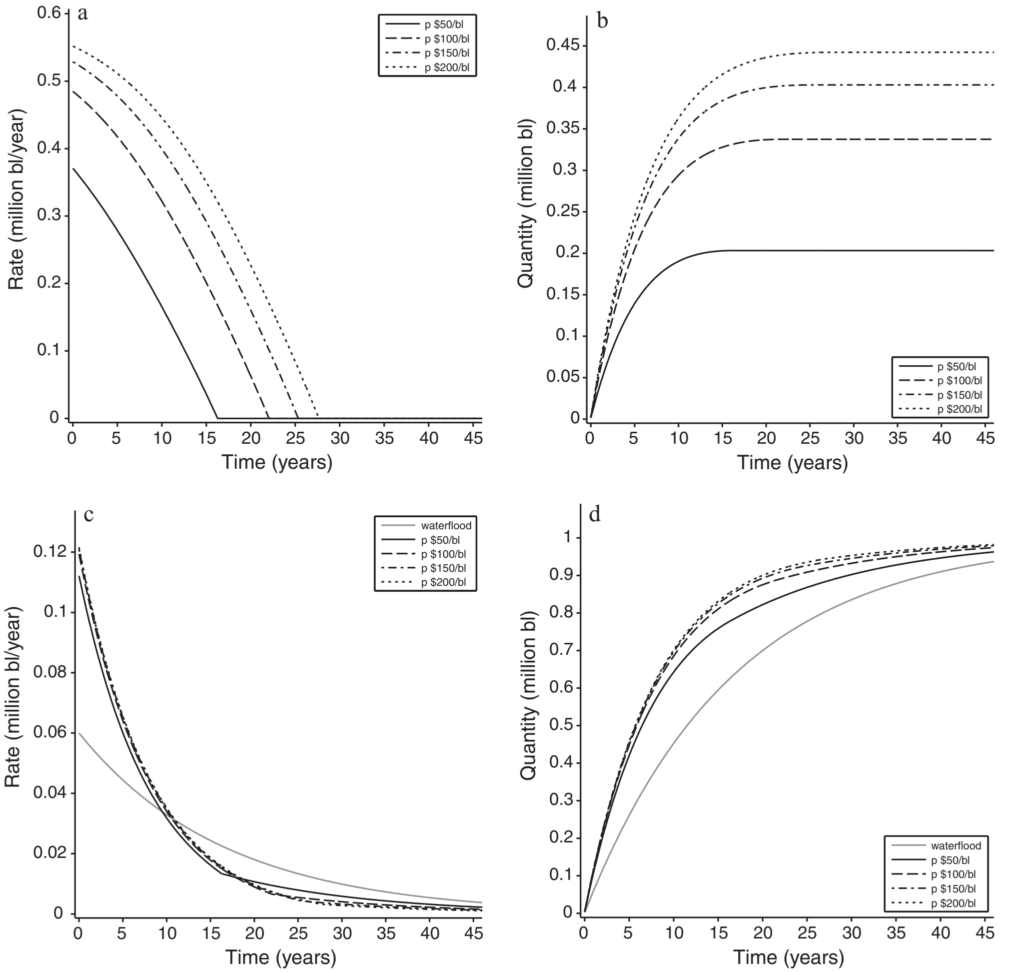

Fig. 3. Change in (a) CO_2 injection, (b) cumulative sequestration, (c) oil production, and (d) cumulative oil production paths as a result of oil price changes.

curves in Fig. 4 are substantially more elastic: at the baseline price of $100, the cross-price elasticity of cumulative sequestration is 0.52, while that of annualized sequestration is 0.47.

5.2. Effects of carbon tax

We continue our investigation of the comparative dynamics of the model with the effect of higher carbon taxes. Such taxes reduce both the net-of-tax oil price received by the producer and the net-of-tax input price of CO_2.

However, Eq. (4) shows that at a given injection rate $c(t)$ the change in oil revenues from a marginal tax change is $-\beta q_p^o(t)\,d\tau$, whereas the change in input costs is $-c(t)q_p^o(t)d\tau$. Since in our model $c(t) \leq 1$ and β=2.2, the revenue effect dominates, inducing the firm to move the injection schedule forward in time. Indeed, panel (a) of Fig. 5 shows that the optimal initial injection rate increases in the tax rate. However, it also shows that the optimal time to switch to pure water injection is accelerated. As a result, injection rates decline more rapidly over time the higher is the tax. Even so, the overall effect of higher carbon taxes on cumulative sequestration is positive, as shown in panel (b). Panels (c) and (d)

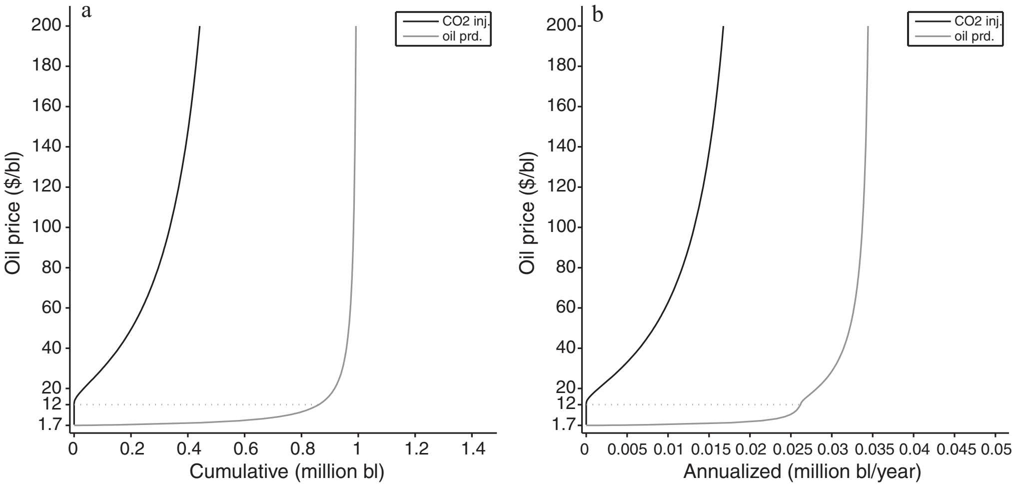

Fig. 4. CO_2 sequestration and oil production, both cumulative (a) and annualized (b), as a function of the oil price.

show that oil production tends to be insensitive to the level of the carbon tax, which coincides with our earlier observation that revenue effects from oil sales tend to be more important than input costs in driving the firm's output decisions.

Nevertheless, it is worth noting that at the highest tax level considered in the figure, namely $120/tCO$_2$, the initial injection rate in panel (a) slightly exceeds the oil-output-maximizing rate of 0.625 million bl/year; further increases in the tax level would raise the initial injection rate to even higher levels. Because such high tax levels make the net CO_2 price strongly negative, CO_2 injection is optimally pushed to levels where its marginal effect on oil production becomes negative as well, thereby reducing initial oil production. In effect, the producer sacrifices oil output and revenues early on in return for higher sequestration revenues that result from higher initial CO_2 injection rates. At very high tax rates, in other words, higher sequestration revenues come at the expense of lower oil revenues.

Fig. 6 shows CO_2 sequestration supply and associated oil output as a function of the carbon tax. Because the oil extraction paths are very close to their revenue-maximizing values regardless of the level of the carbon tax, oil output is almost perfectly inelastic with respect to the carbon tax. More surprising is that the sequestration supply curves are quite inelastic as well. At the current European tax level of about $40/tCO$_2$, the elasticity of cumulative sequestration is only 0.05, and that of annualized sequestration only 0.06. Even at much higher taxes, up to $400/tCO$_2$, these elasticities never exceed 0.55.

5.3. Effects of oil price and carbon tax combined

To recap, the results of Section 5.1 suggest that the rates of CO_2 injection and sequestration by a given EOR project are both relatively responsive to higher oil prices. The result is larger levels of cumulative sequestration at higher prices. On the other hand, the results of Section 5.2 indicate that higher carbon-tax levels increase the optimal CO_2 injection rate early on, but reduce it later, with the same qualitative effects on the induced CO_2 sequestration rate. The initial increase in sequestration dominates, however, resulting in higher overall levels of cumulative sequestration. Even so, the net impact is relatively small, so that cumulative sequestration by the project is relatively unresponsive to higher carbon taxes.

In this subsection, we briefly consider combined changes in oil price and carbon tax, to look for possible interaction effects. Panel (a) of Fig. 7 shows optimal CO_2 injection at four oil price/carbon tax combinations, namely oil prices of $100/bl and $200/bl, with carbon taxes of $40/tCO$_2$ and $80/tCO$_2$. The plots in panel (a) suggest a negative interaction effect: at both oil prices, a doubling of the carbon

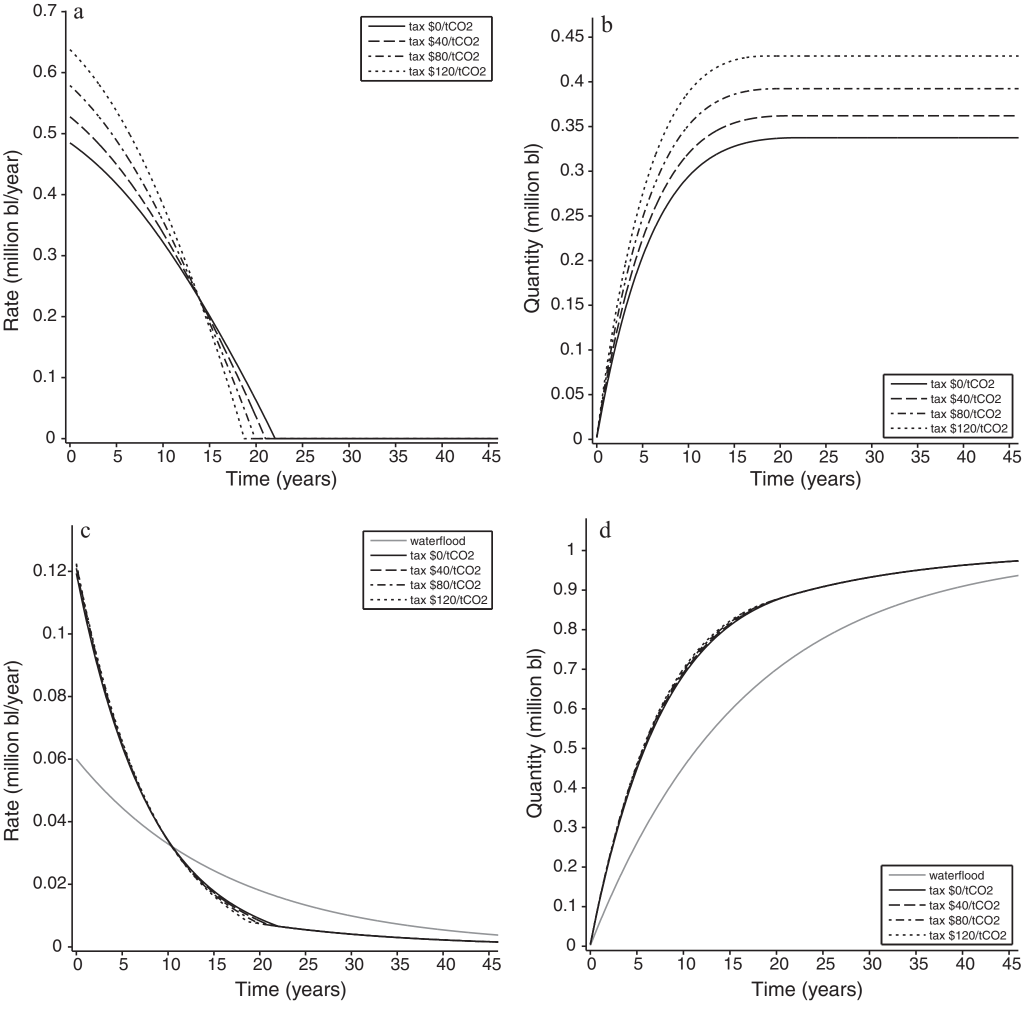

Fig. 5. Change in (a) CO_2 injection, (b) cumulative sequestration, (c) oil production, and (d) cumulative oil production paths as a result of carbon tax changes.

tax tilts the injection path forward in time, but the effect is smaller at the higher oil price. As a result, the increment in cumulative sequestration, shown in panel (b), is smaller as well.

The more obvious point to take away from this figure, however, is that both CO_2 injection and cumulative sequestration are far less sensitive to the carbon tax than to the oil price. While a doubling of the carbon tax does tilt the CO_2 injection path forward, and does increase the ultimate amount of sequestered carbon, these effects pale by comparison with the impacts due to a doubling of the oil price. The interesting – and somewhat paradoxical – implication is that for EOR projects, high oil prices are much more potent incentives for sequestration than are high carbon taxes. The explanation for the paradox lies in Fig. 1, which shows that for CO_2 injection rates up to 0.625 million barrels/year EOR projects do not face any tradeoff between oil production and CO_2 sequestration: increases in CO_2 injection increase both. Higher oil prices induce the firm to significantly increase CO_2 injection throughout the lifetime of the CO_2 flood, so as to bring oil production even closer to its physical maximum rate than it already is. As a result, a greater fraction of the space in the reservoir vacated by

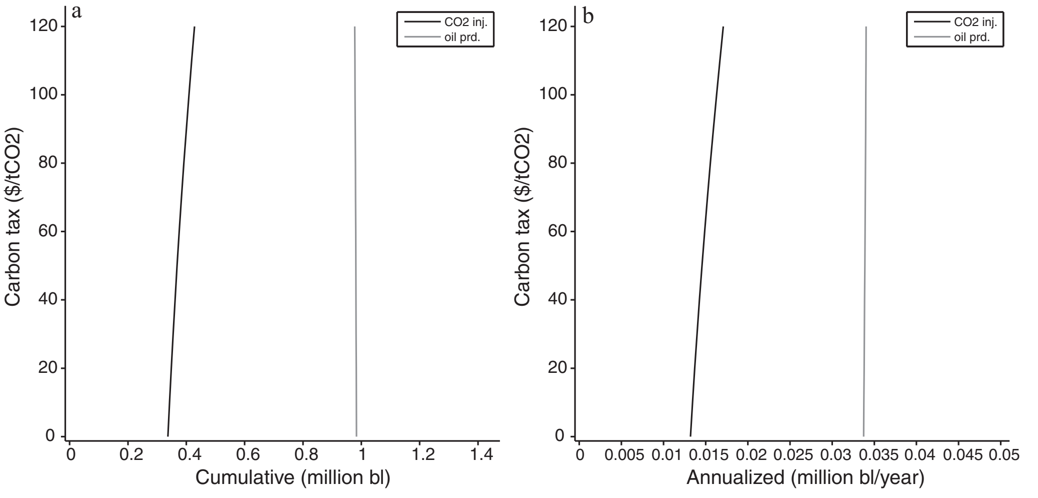

Fig. 6. CO_2 sequestration and oil production, both cumulative (a) and annualized (b), as a function of the carbon tax.

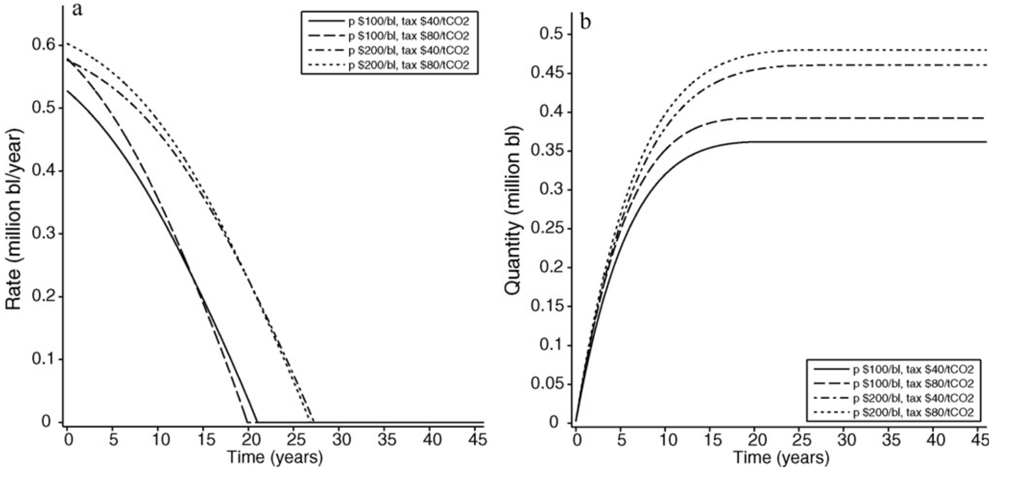

Fig. 7. CO_2 time paths of injection (a) and cumulative sequestration (b), for various combinations of oil price and carbon tax.

the produced oil is taken up by CO_2 rather than water. The higher sequestration, in other words, arises essentially as an unintended by-product, or side effect, of the higher oil production.

As noted in Section 5.2, it is only at very high carbon-tax levels – above $120/t$CO_2$ – that sequestration revenues start to compete with oil revenues, as the taxes induce the firm to increase CO_2 injection beyond the oil-output-maximizing rate of 0.625 million barrels/year.

6. Conclusion

In this paper, we have examined how the standard resource-economics problem of optimizing the rate of oil extraction from a field is altered when the producer has the option of increasing the rate of oil extraction through continuous injections of a mix of CO_2 and water into the reservoir. Our focus in the paper is on the producer's problem of determining the optimal CO_2 fraction in the injection stream and thus the effects of carbon taxes and oil prices on oil production and carbon sequestration.

Our theoretical analysis of this problem indicates that the optimal CO_2 fraction will typically decline over time, and may eventually drop to zero before it becomes optimal to terminate injection altogether. Numerical simulations confirm these results and allow us to further investigate comparative dynamics of the model.

Our simulation results suggest good news and bad news for potential carbon sequestration from existing EOR projects. The bad news is that the use, and resulting sequestration, of CO_2 by such projects appears to be quite inelastic to carbon taxes. Policies raising the cost of CO_2 emissions, whatever effect they may have on other forms of sequestration or abatement, will therefore not induce large increases in EOR-based sequestration. The good news is that market conditions favoring high oil prices *are* likely to induce such increases in sequestration, essentially as a by-product of producers' attempts to increase oil output.

Of course, the apparent inelasticity of EOR-based carbon sequestration at the individual project level need not imply inelastic supply at a more aggregated level. Moreover, because oil reservoirs are generally not spatially homogeneous with respect to relevant physical parameters such as thickness, permeability, and integrity of the cap rock, it is conceivable that EOR would be attractive in some sections of a project's reservoir, but not others, for a given combination of economic parameters.[27] In such a scenario, the supply of sequestration services by the project might be less inelastic to the carbon price than our results indicate. Additionally, if one imagines comparing across different reservoirs, it seems likely that EOR projects would come on-line at different combinations of oil price and carbon price. Again, this observation suggests that the sequestration supply curve for a broader geographic entity, such as a state or country as a whole, would likely be less inelastic than is true for the single unit that we study.

The good news – the significant responsiveness of sequestration to oil prices – comes with three important caveats. The first caveat concerns a counter-balancing effect that applies at the larger geographic level, but is insignificant at the single-unit level. Because large-scale deployment of EOR will generally raise aggregate oil production, it will tend to reduce the market price of oil for any given level of the carbon tax. This in turn will increase the consumption of petroleum-based products, such as motor vehicle fuel, which will generate increased carbon emissions in its own right. It is not clear how these additional emissions compare to the sequestration associated with EOR, but it is conceivable that, on balance, EOR could lead to a net *increase* in carbon emissions at the state or national level.[28] The second caveat concerns the possibility that EOR operators may choose to blow down their reservoirs after terminating injection, in an attempt to extract a small amount of additional oil while also releasing some fraction of the CO_2 sequestered underground back to the surface. We tested a version of our numerical model modified to allow for blowdown, and found that the results are generally robust to its inclusion. At low carbon taxes – below $30/t$CO_2$ in our numerical model – the released CO_2 generates an additional incentive for blowdown, as it can be sold to other EOR projects at a profit. At higher taxes, the released CO_2 must be sold at a loss, but blowdown may still be attractive if additional oil revenues outweigh that cost. We find that overall (undiscounted) sequestration from the CO_2 flood and blowdown phase combined becomes inelastic in the oil price. The reason is that high oil prices are as potent an incentive for releasing CO_2 during blowdown as they are for sequestering it during the CO_2 flood. In contrast, because high carbon taxes are somewhat effective at discouraging CO_2 releases, overall sequestration becomes somewhat more elastic in the carbon tax.[29] The third caveat concerns the overall sequestration capacity of EOR projects. While seemingly large in absolute terms, this capacity is quite small in comparison to both overall CO_2 emissions and the capacity of other geological sequestration options. For example, Dooley et al. (2006) estimate the potential sequestration capacity of all depleted U.S. oil reservoirs, including those

[27] This is true, for example, of the Salt Creek field in Wyoming, one of the largest EOR projects currently operating in the US.

[28] Although beyond the scope of this paper, this issue has received considerable attention in the debate over including CCS projects in the Kyoto Protocol's Clean Development Mechanism (CDM). This mechanism allows countries that have committed to reducing their greenhouse-gas emissions under the Kyoto Protocol to receive credit for reductions achieved in developing countries, through sponsoring certain projects. A key question thereby is the relevant "baseline" relative to which reductions should be measured. For reviews of this debate see IEA GHG (2007), Philibert et al. (2007), and de Coninck (2008).

[29] We should emphasize that, given the paucity of data on blowdown and conflicting evidence on how commonly EOR operators have applied it in the past, our blowdown results should be seen as speculative at best.

depleted through EOR, as 12 $GtCO_2$. That amounts to just 2 years' worth of U.S. CO_2 emissions (EPA, 2008). In contrast, Dooley et al. (2006) estimate the sequestration capacity of U.S. saline aquifers to be as large as 3630 $GtCO_2$.

These points do not imply that EOR has no positive social role to play in promoting geological carbon sequestration. The societal importance of EOR lies in the widely held expectation that it can provide a bridge to that long run. That is, profits from CO_2-enhanced oil output can be used to "jump-start" the building of pipelines and other infrastructure required for ultimately much larger-scale sequestration in non-oil-bearing formations.[30] Three potential extensions are left to future work. The first is the quantification of EOR profits. Analysis of this question will require accounting for up-front investment costs associated with converting a field to CO_2 injection. Preliminary estimates suggest that, for a field of the scale used in our numerical simulations, these costs would amount to several million dollars, and that the cutoff oil price at which incurring these costs would be justified lies around \$50/bl in the absence of a carbon tax. The second extension concerns the effect of rising (rather than atemporal) oil prices and carbon taxes on both the optimal management of a CO_2 flood and the decision to initiate such a flood. Increasing prices may well induce a delay in both oil extraction and the decision to switch to a CO_2 flood, and geologically heterogeneous projects will make the switch at different cutoff prices. Interestingly, it seems likely that a delay in the switch to a CO_2 flood may end up *reducing* sequestration in a given reservoir, since a lower remaining reserve stock leaves less oil to be replaced by CO_2. Finally, it is possible that different results would be obtained it reserves at any point were taken to be jointly determined with CO_2 injections. By reducing the viscosity of reservoir oil, CO_2 injections may not only enhance the rate at which a given reserve stock can be extracted, but also increase the stock itself. This reserve-enhancing effect of CO_2 injections is likely to not be instantaneous, however, but rather tied to the cumulative amount of CO_2 injected. Modeling it would therefore require introducing cumulative injection as a second state variable, thereby significantly complicating the analysis.

Acknowledgement

This research was funded in part by the U.S. Department of Energy and the National Energy Technology Laboratory through Award number: DE-FC26-05NT42587. The authors thank, without implicating, participants at the NBER Summer Institute workshop on Energy and Environmental Economics, an anonymous referee and the Editor for constructive comments.

Appendix A. Derivation of model parameters

Oil producers commonly measure CO_2 in units of 1000 cubic feet (mcf) at standard surface temperature and pressure conditions. In Wyoming, the purchase price of CO_2 is currently about \$2 per mcf. To convert this price to reservoir barrels, we have to take account of the fact that the CO_2 is greatly compressed when it is injected into the reservoir. At Wyoming's Lost Soldier-Tensleep field (LSTP), the compression factor (referred to by reservoir engineers as the "formation volume factor for CO_2") is 0.471 rb/mcf. Rounding this up to 0.5 yields a gross CO_2 purchase price w_m of \$4/rb. The unit cost w_r of separating and recycling CO_2 that is mixed in with the produced oil is on the order of \$0.50/mcf, or \$1/rb.

Operating costs unrelated to injection or recycling of CO_2 amount to about \$24,000 per well per year in non-injection or production-related expenses, plus about \$0.0125 per barrel of overall injection or production. Applying our scaling factor of 1/40 to LSTP's total of about 110 active wells, each of which produce or inject about 800,000 rb per year, this works out to fixed costs F of about \$0.1 million dollars per year.

[30] As noted by William L. Townsend, CEO of a major CO_2-pipeline company, in recent Congressional testimony: "It is clear that the long-term geologic sequestration answer to single-point, industrial CO_2 emissions capture and storage is in saline aquifers, not EOR projects. That being said, there is a very strong, cost-effective interim answer for the next 10 years that employs the oil-based revenues in EOR to subsidize the infrastructure build-out and prepare the foundation of a carbon highway for the next generation of cost-effective CCS in power generation." Townsend (2007). The same view was expressed also in the testimony by George Peridas (Peridas, 2008) and in the National Petroleum Council report (NPC, 2007) cited in footnote 1.

Although we express carbon taxes throughout the paper in terms of dollars per tCO_2, for conformity with the other prices in the model (p, w_m and w_r) the parameter τ is expressed in dollars per rb. Since 1 tCO_2 corresponds to about 19.05 mcf, the above-mentioned conversion factor for LSTP of 0.5 mcf/rb results in a combined conversion factor of 9.5 rb/tCO_2, which we round up to 10. In other words, a carbon tax of \$40/tCO_2 translates to a per-barrel tax of \$4.

Based on estimates by Gately and Huntington (2002) of the long-run elasticity of world oil demand and by Gately (2004) of the long-run elasticity of non-OPEC oil supply, we set the tax incidence on producers at 55%. Weighting Gately and Huntingon's (2002) demand elasticity estimates of -0.60 for OECD and -0.18 for non-OECD countries by the respective OECD and non-OECD demand shares of about two-thirds and one-third yields a world demand elasticity ε of -0.46. Gately's (2004) median estimate for the elasticity η of non-OPEC world supply is 0.365. Combining these yields a producer incidence of $|\varepsilon|/(|\varepsilon|+\eta)\approx 0.55$. Note the implicit assumption that OPEC suppliers will respond to a carbon tax in the same manner as non-OPEC suppliers, rather than responding strategically.

Based on data reported in EPA (2007), we estimate the quantity of CO_2 generated by combusting one barrel of oil at around 0.4 tCO_2, or 4 rb. Multiplying this by the incidence of 55%, and recalling that τ in the numerical model is expressed in dollars per rb, we end up with a combined incidence parameter of $\beta=2.2$.

The decline rate of overall oil production at LSTP since it started its CO_2 flood in 1989 is about 11.5%, whereby the fraction of CO_2 in overall injection has been held roughly constant over time at 0.35. Also, simulation data based on data from an oil field in China indicate that, compared to cumulative oil recovery after 6 years of injecting pure water, recovery after 6 years of injecting a mix of half CO_2, half water is higher by a factor of two, while recovery after 6 years of injecting pure CO_2 is higher by a factor of five-thirds (Guo et al., 2006). These data are consistent with a quadratic $\delta(c)$ function

$$\delta(c) = \delta_w + \delta_1 c - \delta_2 c^2$$

with parameters $\delta_w = 0.06$, $\delta_1 = 0.2$ and $\delta_2 = 0.16$.

References

Adelman, M.A., 1990. Mineral depletion, with special reference to petroleum. Review of Economics and Statistics 72 (1), 1–10.

Adelman, M.A., 1993. Modeling world oil supply. Energy Journal 14 (1), 1–32.

Al-Shuraiqi, H.S., Muggeridge, A.H., Grattoni, C.A., 2003. Laboratory investigations of first contact miscible wag displacement: the effects of wag ratio and flow rate. Working Paper 84894, Society of Petroleum Engineers (SPE).

ARI, 2006. Evaluating the Potential for "Game Changer" Improvements in Oil Recovery Efficiency from CO_2 Enhanced Oil Recovery. Advanced Resources International. Prepared for U.S. Department of Energy Office of Fossil Energy - Office of Oil and Natural Gas.

Aycaguer, A.-C., Lev-On, M., Winer, A.M., 2001. Reducing carbon dioxide emissions with enhanced oil recovery projects: a life cycle assessment approach. Energy & Fuels 15 (2), 303–308.

Babadagli, T., 2006. Optimization of CO_2 injection for sequestration/enhanced oil recovery and current status in Canada. In: Lombardi, S., Altunina, L.K., Beaubien, S.E. (Eds.), Advances in the Geological Storage of Carbon Dioxide: International Approaches to Reduce Anthropogenic Greenhouse Gas Emissions. Springer, Dordrecht, The Netherlands, pp. 261–270.

Bandza, A.J., Vajjhala, S.P., 2008. Long-term risks and short-term regulations: modeling the transition from enhanced oil recovery to geologic carbon sequestration. RFF Discussion Paper 08-29, Resources for the Future, Washington, DC.

Benson, S.M., 2006. Monitoring carbon dioxide sequestration in deep geological formations for inventory verification and carbon credits. Working Paper 10374, Society of Petroleum Engineers (SPE).

Benson, S.M., Hoversten, M., Gasperikova, E., Haines, M., 2005. Monitoring protocols and life-cycle costs for geologic storage of carbon dioxide. In: Proceedings of the 7th International Conference on Greenhouse Gas Control Technologies (GHGT-7), September 5–9, 2004, Vancouver, Canada. Vol. 2 of Contributed Papers and Panel Discussion. Elsevier, Oxford, UK, pp. 1259–1264.

Cairns, R.D., Davis, G.A., 2001. Adelman's rule and the petroleum firm. Energy Journal 22 (3), 31–54.

Cakici, M.D., 2003. Co-optimization of oil recovery and carbon dioxide storage. Engineer Thesis, Department of Petroleum Engineering, Stanford University.

Davis, G.A., Cairns, R.D., 1999. Valuing petroleum reserves using current net price. Economic Inquiry 37 (2), 295–311.

de Coninck, H., 2008. Trojan horse or horn of plenty? Reflections on allowing CCS in the CDM. Energy Policy 36 (3), 929–936.

Dooley, J.J., Dahowski, R.T., Davidson, C.L., Wise, M.A., Gupta, N., Kim, S.H., Malone, E.L., 2006. Carbon dioxide capture and geologic storage. A technology report from the second phase of the Global Energy Technology Strategy Program.

EIA, 2007. Oil and gas lease equipment and operating costs 1988 through 2006. Energy Information Administration. Available at: http://www.eia.doe.gov.

EPA, 2007. Inventory of U.S. greenhouse gas emissions and sinks: fast facts 1990–2005. Conversion factors to energy units (heat equivalents) heat contents and carbon content coefficients of various fuel types. U.S. Environmental Protection Agency, Washington, DC, EPA 430-R-07-002.

EPA, 2008. Inventory of U.S. greenhouse gas emissions and sinks: 1990–2006. U.S. Environmental Protection Agency, Washington, DC, EPA 430-R-08-005.

Fetkovitch, M.J., 1980. Decline curve analysis using type curves. Journal of Petroleum Technology 32 (6), 1065–1077.

Gately, D., 2004. OPEC's incentives for faster output growth. Energy Journal 25 (2), 75–96.

Gately, D., Huntington, H.G., 2002. The asymmetric effects of changes in price and income on energy and oil demand. Energy Journal 23 (1), 19–55.

Guo, X., Du, Z., Sun, L., Fu, Y., 2006. Optimization of tertiary water-alternate-CO_2 flood in Jilin oil field of China: laboratory and simulation studies. Working Paper 99616, Society of Petroleum Engineers (SPE).

Hepple, R.P., Benson, S.M., 2005. Geologic storage of carbon dioxide as a climate change mitigation strategy: performance requirements and the implications of surface seepage. Environmental Geology 47 (4), 576–585.

Herzog, H., Caldeira, K., Reilly, J., 2003. An issue of permanence: assessing the effectiveness of temporary carbon storage. Climatic Change 59 (3), 293–310.

IEA GHG, 2007. ERM – Carbon Dioxide Capture and Storage in the Clean Development Mechanism. No. 2007/TR2. International Energy Agency Greenhouse Gas R&D Programme, Cheltenham, UK.

IPCC, 2005. IPCC Special Report: Carbon Dioxide Capture and Storage. Intergovernmental Panel on Climate Change, Cambridge University Press, Cambridge, UK.

Jessen, K., Kovscek, A.R., Orr Jr., F.M., 2005. Increasing CO_2 storage in oil recovery. Energy Conversion and Management 46, 293–311.

Juanes, R., Blunt, M.J., 2006. Analytical solutions to multiphase first-contact miscible models with viscous fingering. Transport in Porous Media 64, 339–373.

Klusman, R.W., 2003. A geochemical perspective and assessment of leakage potential for a mature carbon dioxide-enhanced oil recovery project and as a prototype for carbon dioxide sequestration; Rangely field, Colorado. American Association of Petroleum Geologists Bulletin 87 (9), 1485–1507.

Kovscek, A.R., Cakici, M.D., 2005. Geologic storage of carbon dioxide and enhanced oil recovery. II. Cooptimization of storage and recovery. Energy Conversion and Management 46, 1941–1956.

Leach, A., Mason, C.F., van't Veld, K., 2010. Co-optimization of enhanced oil recovery and carbon sequestration. Climates of Change: Sustainability Challenges for Enterprise Working Paper 005, Smith School of Enterprise and the Environment, University of Oxford, UK.

Masoner, L.O., Wackowski, R.K., 1994. Rangely Weber Sand Unit CO_2 project update: decisions and issues facing a maturing EOR project. Working Paper 27756, Society of Petroleum Engineers (SPE).

McCoy, S.T., 2008. The economics of CO_2 transport by pipeline and storage in saline aquifers and oil reservoirs. PhD Thesis, Carnegie Mellon University, Pittsburgh, PA.

McCoy, S.T., Rubin, E.S., 2009. The effect of high oil prices on EOR project economics. Energy Procedia 1 (1), 4143–4150.

Mian, M.A., 2002. Project Economics and Decision Analysis, vol. I–II. PennWell, Tulsa, OK.

MIT, 2007. The Future of Coal: An Interdisciplinary MIT Study. Massachusetts Institute of Technology, Cambridge, MA.

NPC, 2007. Hard Truths: Facing the Hard Truths About Energy. National Petroleum Council, Washington, DC.

Peridas, G., 2008. Written testimony submitted to the Natural Resources Committee, Subcommittee on Energy and Mineral Resources, Unites States House of Representatives hearing on "Spinning Straw Into Black Gold: Enhanced Oil Recovery Using Carbon Dioxide," June 12, 2008.

Philibert, C., Ellis, J., Podkanski, J., 2007. Carbon Capture and Storage in the CDM. Organization for Economic Co-Operation and Development, International Energy Agency, Paris, France, cOM/ENV/EPOC/IEA/SLT(2007)10.

Shaw, J.C., Bachu, S., 2003. Evaluation of the CO_2 sequestration capacity in Alberta's oil and gas reservoirs at depletion and the effect of underlying aquifers. Journal of Canadian Petroleum Technology 42 (9), 51–61.

Snyder, S., Anderson, I., Keith, D., Sendall, K., Youzwa, P., 2008. Canada's Fossil Energy Future: The Way Forward on Carbon Capture and Storage. Report by the ecoENERGY Carbon Capture and Storage Task Force to the Minister of Alberta Energy and the Minister of Natural Resources Canada.

Thompson, A.C., 2001. The Hotelling principle, backwardation of futures prices and the values of developed petroleum reserves – the production constraint hypothesis. Resource and Energy Economics 23 (2), 133–156.

Townsend, W.L., 2007. Testimony before the Energy Subcommittee of the Finance Committee, United States Senate hearing on "Coal: A Clean Future, Response of the Market to Global Incentives and Mandates for Clean Coal," April 26, 2007.

Trivedi, J.J., Babadagli, T., 2007. Optimal injection strategies for CO_2 and flue gas sequestration during tertiary oil recovery. Oil Gas-European Magazine 33 (1), 22–26.

Zhou, W., Stenhouse, M.J., Arthur, R., Whittaker, S., Law, D.H.-S., Chalaturnyk, R., Jazrawi, W., 2005. The IEA Weyburn CO_2 monitoring and storage project: modeling of the long-term migration of CO_2 from Weyburn. In: Proceedings of the 7th International Conference on Greenhouse Gas Control Technologies (GHGT-7), September 5–9, 2004, Vancouver, Canada. Vol. 1 of Peer-Reviewed Papers and Plenary Presentations. Elsevier, Oxford, UK, pp. 721–730.