NBER WORKING PAPER SERIES

CO-OPTIMIZATION OF ENHANCED OIL RECOVERY AND CARBON SEQUESTRATION

Andrew Leach
Charles F. Mason
Klaas van't Veld

Working Paper 15035
http://www.nber.org/papers/w15035

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2009

Co-optimization of Enhanced Oil Recovery and Carbon Sequestration
Andrew Leach, Charles F. Mason, and Klaas van't Veld
NBER Working Paper No. 15035
June 2009
JEL No. Q32,Q40,Q54

ABSTRACT

In this paper, we present what is to our knowledge the first theoretical economic analysis of CO2-enhanced oil recovery (EOR). This technique, which has been used successfully in a number of oil plays (notably in West Texas, Wyoming, and Saskatchewan), entails injection of CO2 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 CO2 remain sequestered in the reservoir. If CO2 emissions are regulated, oil producers using EOR should therefore be able to earn sequestration credits in addition to oil revenues. We develop a theoretical framework that analyzes the dynamic co-optimization of oil extraction and CO2 sequestration, through the producer's choice at each point in time of an optimal CO2 fraction in the injection stream (the control variable). We find that the optimal fraction 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. They show also that cumulative sequestration is positively related to the oil price, and is in fact much more responsive to oil-price increases than to increases in the carbon tax. Only at very high taxes does a tradeoff between oil output and sequestration arise.

Andrew Leach
University of Alberta
School of Business
3-40K Business Building
Edmonton, Alberta T6G 2R6
Canada
andrew.leach@ualberta.ca

Charles F. Mason
Department of Economics and Finance
University of Wyoming
1000 E. University Avenue
Laramie, WY 82071-3985
bambuzlr@uwyo.edu

Klaas van't Veld
Department of Economics and Finance
University of Wyoming
1000 E. University Avenue
Laramie, WY 82071-3985
klaas@uwyo.edu

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.[1] 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 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 seems to be broad agreement, moreover, that large-scale deployment of geological sequestration is likely to start with projects that apply CO_2-enhanced oil recovery (EOR).[2]

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 in a manner that causes 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.[3] 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, all the injected CO_2 is sequestered.

EOR is a "game-changing" technology for the recovery of oil 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

[1] For example, the Vice President of Environmental Policy for Duke Energy stated in a Washington Post article, "Our viewpoint is that it's going to happen. There's scientific evidence of climate change. We'd like to know what legislation will be put together so that, when we figure out how to increase our load, we know exactly what to expect." (Steven Mufson and Juliet Eilperin, "Energy Firms Come to Terms With Climate Change," Washington Post, Saturday, November 25, 2005, p. A01)

[2] 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).

[3] 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.

the United States by over 80 billion barrels (ARI; 2006). Similarly, Shaw and Bachu (2003) claim that 4,470 fields, just over half of the known oil reservoirs in Alberta, are amenable to CO_2 injection for enhanced oil recovery. Babadagli (2006) states that enhanced oil recovery applied in these reservoirs could translate in 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 present what is to our knowledge the first theoretical economic analysis of CO_2-enhanced oil recovery. In the tradition of Hotelling (1931), an oil field contains a physical quantity of oil which the producer seeks to extract at a particular rate over time so as to maximize the economic rents from extraction. The ability to enhance oil extraction rates through injection of CO_2 alters this extraction 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 remains sequestered in the reservoir. Separating the remaining fraction that resurfaces with the produced oil, and then dehydrating and recompressing it, is costly as well.

Second, even at a constant injection rate, the amount of oil recovered declines over time, as does the fraction of injected CO_2 that remains sequestered 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's objective would then be the maximization of the combined revenue streams from both oil production and CO_2 sequestration, net of CO_2 purchase and recycling costs.

Fourth, carbon taxes affect these revenue and cost streams in multiple ways. While a carbon tax effectively reduces the input cost for EOR and increases the net present 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[4] indicate that both oil recovery and CO_2 sequestration are maximized when a mix of CO_2 and water is injected (whereby the CO_2 fraction that maximizes oil recovery typically differs from that maximizing sequestration).

Our paper provides an evaluation of the role of all five factors. We start by developing a theoretical framework that analyzes the dynamic co-optimization of oil extraction and CO_2 sequestration, through the producer's choice at each point in time of an optimal CO_2 fraction in the injection stream (the control variable). The decision to cease extraction is determined by a transversality condition. Both the injection and termination decisions depend in part on the anticipated price of oil and the carbon tax or credit price. The paper concludes with a series of simulations that are based on an ongoing project, namely the Lost Soldier-Tensleep field in Wyoming. These 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 positively related to the oil price, and is in fact much more responsive to oil-price increases than to increases in the carbon tax.

[4] 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).[5] In addition, the rate of oil production is linked to remaining reserves by the so-called "decline curve."[6] 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 $c(t)$, and the rate of water injection be $h_i(t)$. We assume the total rate of injection, $I \equiv c(t) + h_i(t)$, is constant across time. This reflects the fact that CO_2-EOR projects are usually operated at "minimum miscibility pressure," 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 rate of oil production at time t be $q(t)$, the rate of CO_2 production (or "leakage") at time t be $\ell(t)$, and the rate of water production at time t be $h_p(t)$. Materials balance then requires that

$$q(t) + l(t) + h_p(t) = c(t) + h_i(t)$$

at each point in time (with both sums equaling I). In practice, the leaked CO_2 could be vented or recycled. Let the price of a unit of newly purchased CO_2 equal w_s and the unit cost of recycling CO_2 equal w_ℓ. We assume that $w_s > w_\ell$, so that it is always cheaper for the

[5] Reservoir engineers refer to this as "materials 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.

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

firm to recycle than to vent.[7] As a result, all leakage is re-injected and total CO_2 injection is the sum of new purchases and leakage, or

$$c(t) = s(t) + \ell(t). \tag{1}$$

The rate at which CO_2 is sequestered depends on the linkage between injected CO_2 and produced 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. As a result, we have $s(t) = c(t)q(t)/I$. As we are focusing on an individual firm and a particular oil reservoir, we may normalize so that $I = 1$. Accordingly,

$$s(t) = c(t)q(t). \tag{2}$$

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. As usual, this variable plays the role of the state variable in our analysis, and it evolves via

$$\dot{R} = -q. \tag{3}$$

In keeping with the physical reality of oil recovery, we assume that the rate of production can be described by a decline curve: $q(t) = \delta R(t)$. In our setting, however, the ratio of output to reserves—which plays the role of the decline rate—is linked to the rate of injection: $\delta = \delta(c)$.

[7] In practice, the purchase price of CO_2 is several times higher than the cost of recycling. 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 $w_s - \tau$ (as the seller of the CO_2 avoids the carbon tax or receives a credit for sequestering), while the opportunity cost of recycling becomes $w_\ell - \tau$ (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. In light of the fact (discussed further in the concluding section of the paper) that aggregate CO_2 emissions currently far outstrip all estimates of aggregate EOR sequestration, this assumption seems reasonable.

We therefore have the relation

$$q(t) = \delta(c)R(t). \tag{4}$$

Combining (3) and (4), we obtain

$$\dot{R} = -\delta(c)R. \tag{5}$$

Consistent with results from the reservoir-engineering studies cited in the introduction, we assume that the $\delta(c)$ function relating the rate of injection 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.[8] We assume that all costs other than those of CO_2 purchases and CO_2 recycling are tied to the overall amount of fluids injected and the amount of fluids produced. As both amounts are constant and equal to I, these other costs are a constant F. Accordingly, the firm earns a rate of profits equal to

$$\pi = pq - (w_s - \tau)s - w_\ell \ell - F.$$

Using (1), (2), and (5), we may rewrite the profit rate as

$$\begin{aligned} \pi &= p\delta(c)R - [w_s - \tau]c\delta(c)R - w_\ell c[1 - \delta(c)R] - F \\ &= p\delta(c)R - [w_s - \tau - w_\ell]c\delta(c)R - w_\ell c - F. \end{aligned} \tag{6}$$

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, it will be convenient to isolate this effect in the expression of profits. To that end, we denote the induced tax liability for a one-dollar

[8] Because our focus in this paper is on the optimal operation of a CO_2 flood, we abstract from up-front investments required to make an oil field "CO_2 ready." Such investments include changes in well equipment, additions of metering equipment and pipelines in the field, and the construction of a CO_2 recycling plant that separates produced CO_2 from the oil and then dehydrates and recompresses it. We discuss likely implications of introducing up-front investment costs in the concluding section of the paper.

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. Adjusting (6) to take account of these aspects, we may write the rate of profits as

$$\pi = (p - \beta\tau)\delta(c)R - [w_s - \tau - w_\ell]c\delta(c)R - w_\ell c - F.$$

To save on notation, we will typically summarize the combination $p - \beta\tau$ as Y and the combination $w_s - \tau - w_\ell$ as Z. Using this notational convention, the profit rate is

$$\pi = Y\delta(c)R - Zc\delta(c)R - w_\ell c - F.$$

3. Analysis

The goal of the firm is to choose a time path of the injection rate $c(t)$ so as to maximize its present discounted value, subject to the state equation (5), the initial value of the state variable, R_0, and the constraints $0 \leq c \leq 1$, $R \geq 0$. Both the terminal time T and the terminal stock $R(T)$ are free, and so the optimal choices of these values will be governed by transversality conditions. To solve this dynamic optimization problem, we first define the current-value Hamiltonian

$$H = \pi - mq = Y\delta(c)R - Zc\delta(c)R - w_\ell c - F - m\delta(c)R, \tag{7}$$

where m is the current-value multiplier (shadow price) associated with a unit of oil *in situ*.

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

$$H_c = (Y - Zc - m)\delta'(c)R - Z\delta(c)R - w_\ell = 0, \tag{8}$$

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 (5), and the equation of motion for the shadow price is

$$\dot{m} = rm - (Y - Zc - m)\delta(c). \tag{9}$$

The right-hand side of equation (9) differs from the standard Hotelling representation by virtue of the second set of terms. These terms capture the fact that current carbon injections reduce the productive capability of future injections. 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 (7) that the profit rate must be zero at time T. But for that to happen there must be positive revenues, which in turn requires a positive production rate. It follows that the terminal stock is positive, so 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 (8) to get

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

where $H_{cx} = \partial^2 H/\partial c \partial x$, $x = m, c$, or R. From (8), we see that $H_{cm} = -\delta'(c)R$ and $H_{cR} = (Y - Zc - m)\delta'(c) - Z\delta(c) = w_\ell/R$ (where we use (8) to extract the last relation). Combining these observations with the state equation, we get

$$\dot{c} = [\delta'(c)R\dot{m} + w_\ell \delta(c)]/H_{cc}. \tag{10}$$

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 (8) and (9) to get

$$\delta'(c)R\dot{m} + w_\ell\delta(c) = [r\delta'(c)m - Z\delta(c)^2]R. \tag{11}$$

Comparing (10) and (11), it is apparent that a sufficient condition for $\dot{c}$ to be negative is for the right-hand side of (11) 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 $Z = w_s - \tau - w_\ell$, and that by assumption $w_s - w_\ell > 0$. If $w_s - w_\ell$ 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 (11) will be positive in a range of cases that seem empirically relevant. As such, the rate of injection will commonly be declining.

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

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

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 (5),

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

Combining (9) and (12), taking note of the fact that $c = 0$, we deduce that

$$\dot{H}_c = -[rm\delta'(0) - Z\delta_w^2]R. \tag{13}$$

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 (13) 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

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

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

$$\dot{H}_c = -[rm\delta'(1) - Z\delta(1)^2]R. \tag{14}$$

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, 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 solved and simulated the model numerically to yield optimal time paths of carbon injection rates for various combinations of oil price and carbon tax. Below we first discuss the solution algorithm and then present results.

The optimization problem is reasonably straightforward, in that it involves a single control variable, CO_2 injection c, which is optimized given a single state variable, remaining physical reserves R. We used two approaches to solving the problem, which (up to rounding errors) yielded identical results. The first approach was a brute-force determination of the optimal time path of c. In this approach, a discretized version of the control problem was programmed. Time was divided into discrete periods $t = 0, \ldots, \overline{T}$ for a large time horizon $\overline{T}$, and the function $c(t)$ was approximated by the $\overline{T}+1$ values c_t that maximize the present value of profits, subject to the (discretized) state equation and bounds on c. Simultaneously, the optimal terminal period T was solved for as well.

In the second approach, the problem was solved by again first converting it to discrete time and then using an algorithm that iterates on an approximation to the solution to the Bellman equation. Here, the solution to the dynamic program is computed using a neural-network approximation defined over a finite set of grid points distributed within the state space.[9]

Let $V(R)$ denote the optimal value function:

$$V(R) = \max_c Y\delta(c)R - Zc\delta(c)R - w_\ell c - F + V(R - \delta(c)R). \tag{15}$$

[9] A one-hidden-layer feedforward neural network as is used in this algorithm provides a uniform approximation to any continuous, multivariate function to any desired degree of accuracy. For a detailed discussion of neural networks see Hassoun (1995), page 46.

I	1	overall rate of injection and production ($\times$1 million barrels)
R_0	1	initial stock of oil in reservoir ($\times$1 million barrels)
w_s	4	per-barrel cost of purchased CO_2
w_ℓ	1	per-barrel cost of separating and recycling "leaked" CO_2
F	0.1	fixed costs ($\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

Table 1. Baseline parameter values.

Write $\Phi(R|\phi)$ as a neural-network approximation of $V(R)$ with parameter values ϕ. The algorithm consists of 5 steps:

1. Draw a distribution of grid points in R space.
2. Begin with an initial guess of $\Phi^0(R) = 0, \forall R$ and solve (15) given this guess at each grid point. Denote the solution to this iteration by $V^1(R)$.
3. Compute the approximation for iteration $i = 1, 2, \ldots$ by solving $\min_\phi\{\Phi(R|\phi) - V^i(R)\}^2$ and denote the solution $\Phi^i(R)$.
4. Solve $V^{i+1} = \max_c Y\delta(c)R - Zc\delta(c)R - w_\ell c - F + \Phi^i(R - \delta(c)R)$.
5. Return to step 3 unless $||V^i(R) - V^{i-1}(R)|| < 10^{-6}$.

The final approximation, $\Phi(R, \phi)$, represents an approximate solution to the dynamic program.

We compute the solution for scenarios with oil prices of \$100, \$200, and \$300 per barrel (bl), taking \$100/bl as our baseline price, 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 in units of 1 million "reservoir" barrels (rb) per year (1 barrel= 42 gallons (US) $\approx 0.16m^3$), meaning barrels 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 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.[10] Oil producers commonly measure CO_2 in units of 1,000 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 LSTP, the compression factor (referred to by reservoir engineers as the "formation volume factor for CO_2") is 0.471 rb/mcf (which, since 1 *mcf* corresponds to about 178 barrels, amounts to a compression rate of about 380 times). Rounding this factor up to 0.5, we end up with a gross CO_2 purchase price w_s of \$4/*rb*. We take the unit cost w_ℓ of separating and recycling "leaked" CO_2 that is mixed in with the produced oil to be on the order of \$0.50/*mcf*, or \$1/*rb*.

It is important to note at this point that, 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_s and w_ℓ) the parameter τ is expressed in dollars per *rb*. Since one 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.

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 producing or injecting about 800,000 *rb* per year, this works out to fixed costs F of about \$0.1 million dollars per year.

[10] In particular, Charles Fox of Kinder Morgan, Inc., and Mark Nicholas of Nicholas Consulting Group.

As noted above, the parameter β, which describes carbon-tax incidence on oil producers, is actually a combination of tax incidence effects and unit conversions that transform a unit of produced oil into carbon units. Tax incidence effects depend on demand and supply elasticities. Based on estimates by Gately and Huntington (2002) of the long-run elasticity of oil demand and by Gately (2004) of the long-run elasticity of non-OPEC oil supply, we set the tax incidence on producers of a given tax expressed in dollars per barrel of oil at 55%. Based on data reported in EPA (2007), we estimate the quantity of CO_2 generated by combusting one barrel of oil at around $0.4tCO_2$, or $4rb$. 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 $\beta = 2.2$.[11]

Lastly, the parameters of $\delta(c)$ function are based on a combination of production experience at LSTP and simulation results in the literature. 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 six years of injecting pure water, recovery after six years of injecting a mix of half CO_2, half water is higher by a factor of two, while recovery after six 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 $\delta_2 = 0.16$.

Figure 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 injection rates above 0.625 million barrels/year are counter-productive to oil recovery, but still increase sequestration.

[11] 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 $\beta = 0$ also, and where the incidence is 100%, so that $\beta = 4$, our results are essentially unchanged.

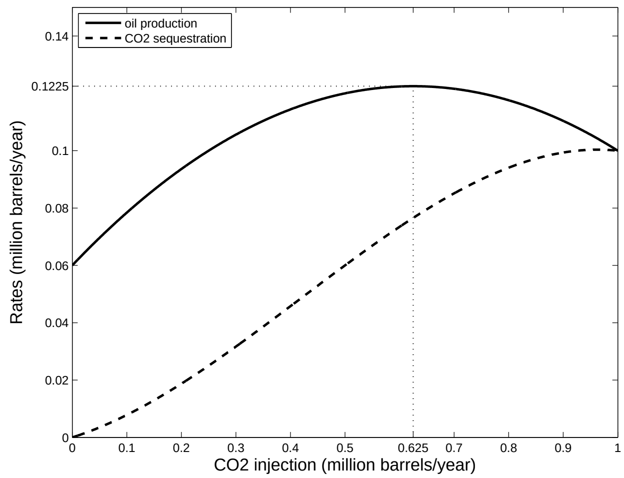

FIGURE 1. Initial oil production and CO_2 sequestration as a function of the initial CO_2 injection rate.

It should be emphasized that, although there are good reasons to believe that the $\delta(c)$ function is concave,[12] 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. The particular parameterization used here should therefore be viewed as only illustrative.

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 Figure 2 shows the optimal paths of CO_2 injection (c), CO_2 leakage (ℓ), oil production (q), and flow CO_2 sequestration (s).

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

[12] See the reservoir-engineering studies cited in footnote 4.

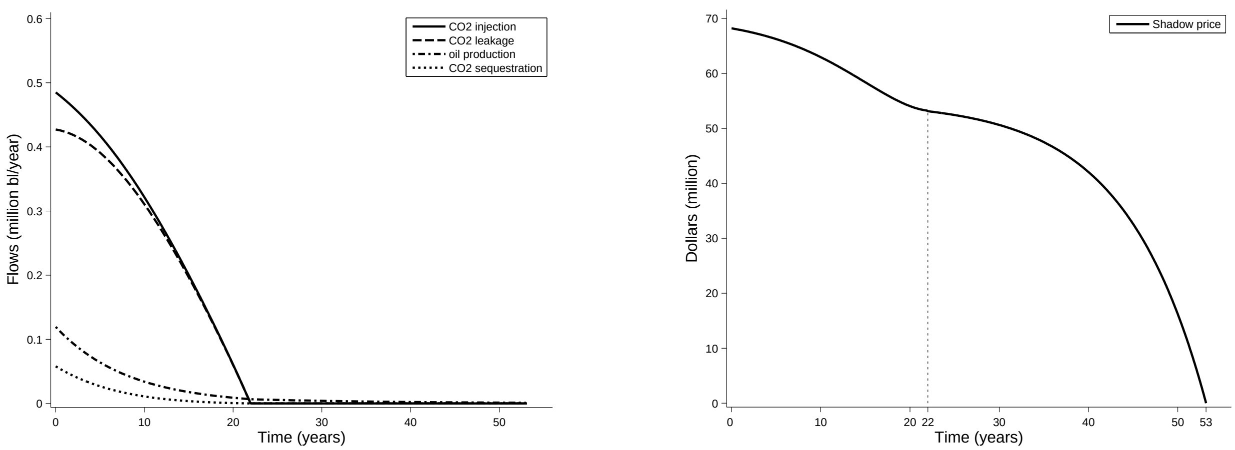

(a) (b)

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

vacated by the oil. As a result, CO_2 leakage, 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 injection 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 Figure 3 shows how the optimal CO_2 injection path changes as the oil price level is raised from \$100 to \$300/bl. Panel (c) of Figure 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

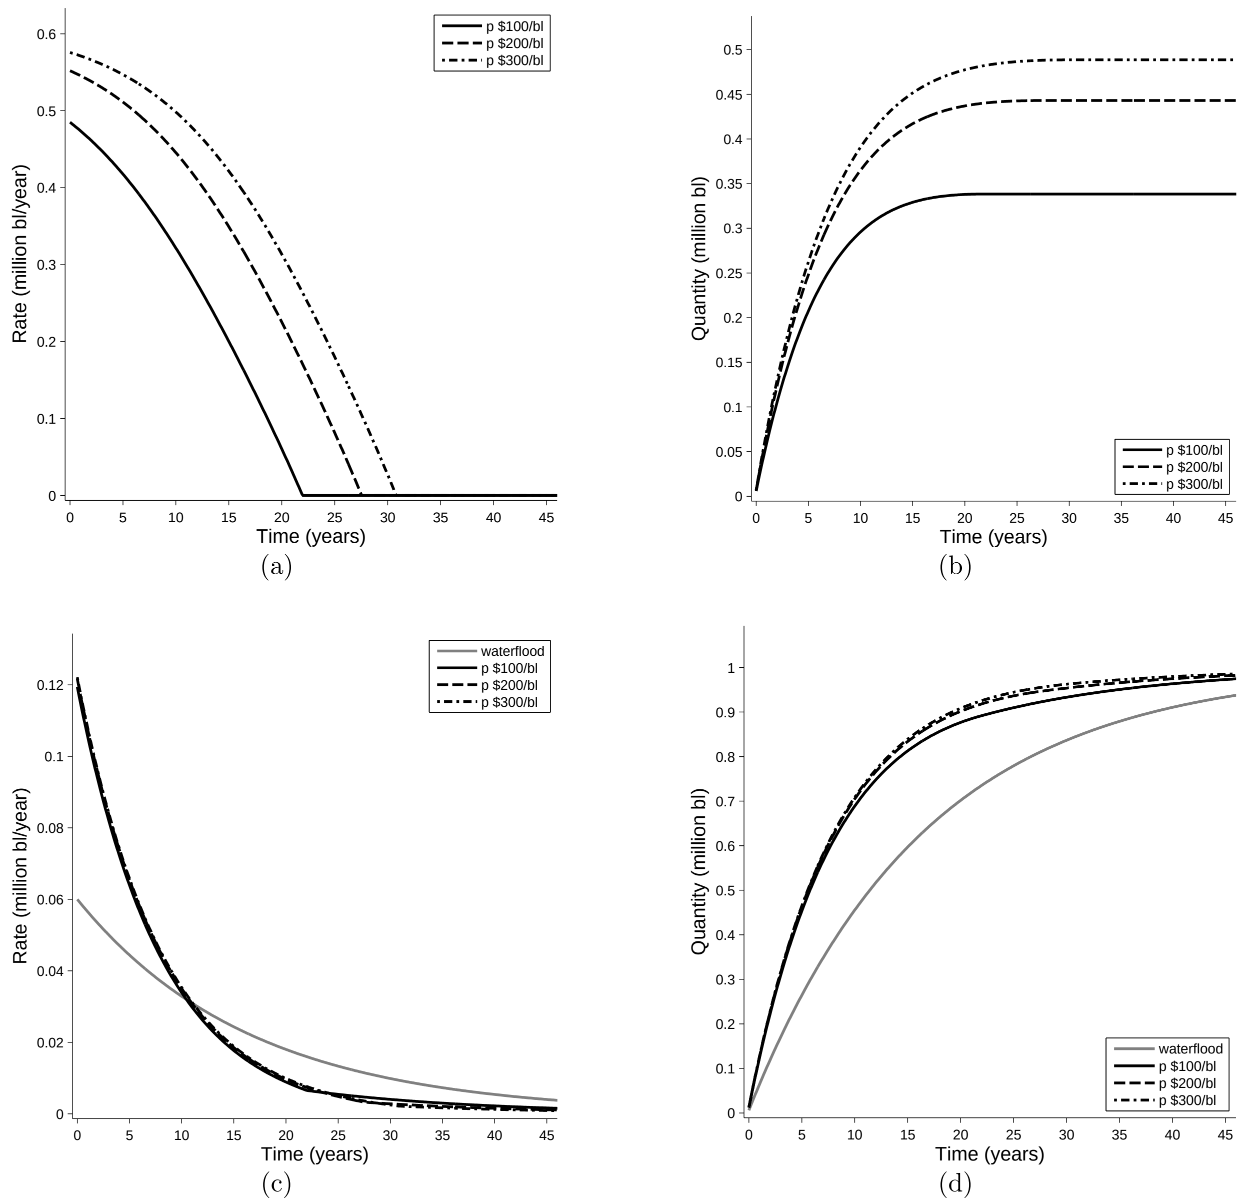

FIGURE 3. Change in (a) CO_2 injection, (b) cumulative sequestration, (c) oil production, and (d) cumulative oil production paths as a result of oil price changes.

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

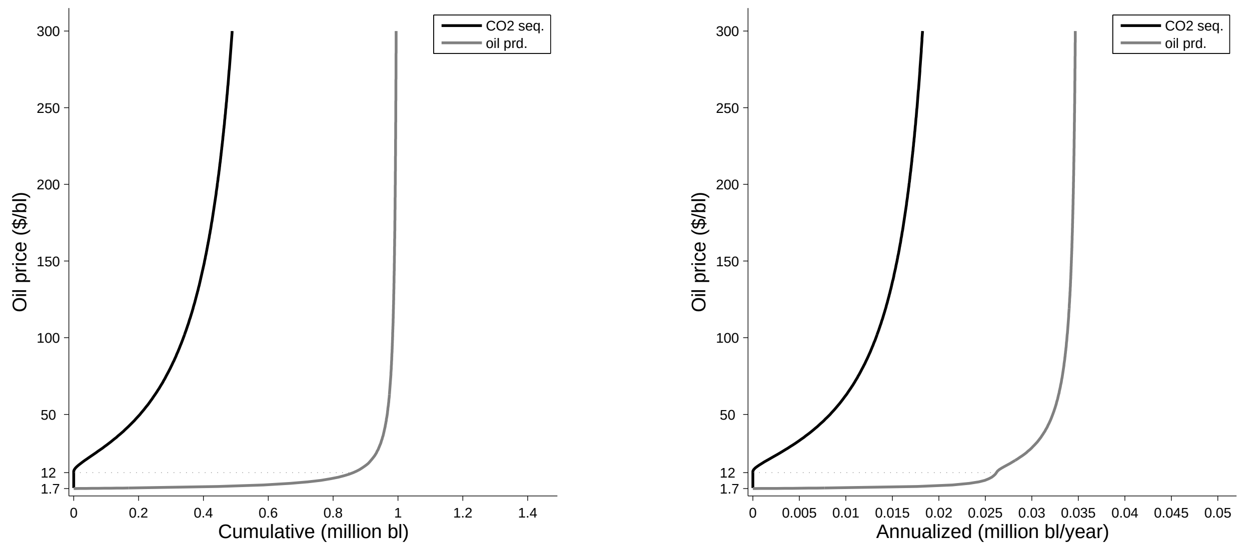

(a) (b)

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

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).

Figure 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.[13] 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.[14] 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.

[13] The annualized values are calculated as the constant rate $\overline{s}$ or $\overline{q}$ 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, $\overline{s}$ is implicitly defined by

$$\int_0^{\infty} e^{-rt}\tau\overline{s}\,dt = \int_0^{\infty} e^{-rt}\tau s(t)\,dt,$$

and $\overline{q}$ is defined analogously.

[14] 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.

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 Figure 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 Figure 3, higher oil prices do induce substantially higher rates of CO_2 injection, and thereby sequestration. As a result, the sequestration curves in Figure 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, at a given injection rate $c(t)$, the change in oil revenues from a marginal tax change is is $-q(t)\,d\tau$, whereas the change in input costs is $-c(t)q(t)\,d\tau$. As long as $c(t)$ is below its upper bound of 1, the revenue effect therefore dominates, in which case the firm is motivated to move the injection schedule forward in time. Indeed, panel (a) of Figure 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) 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 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

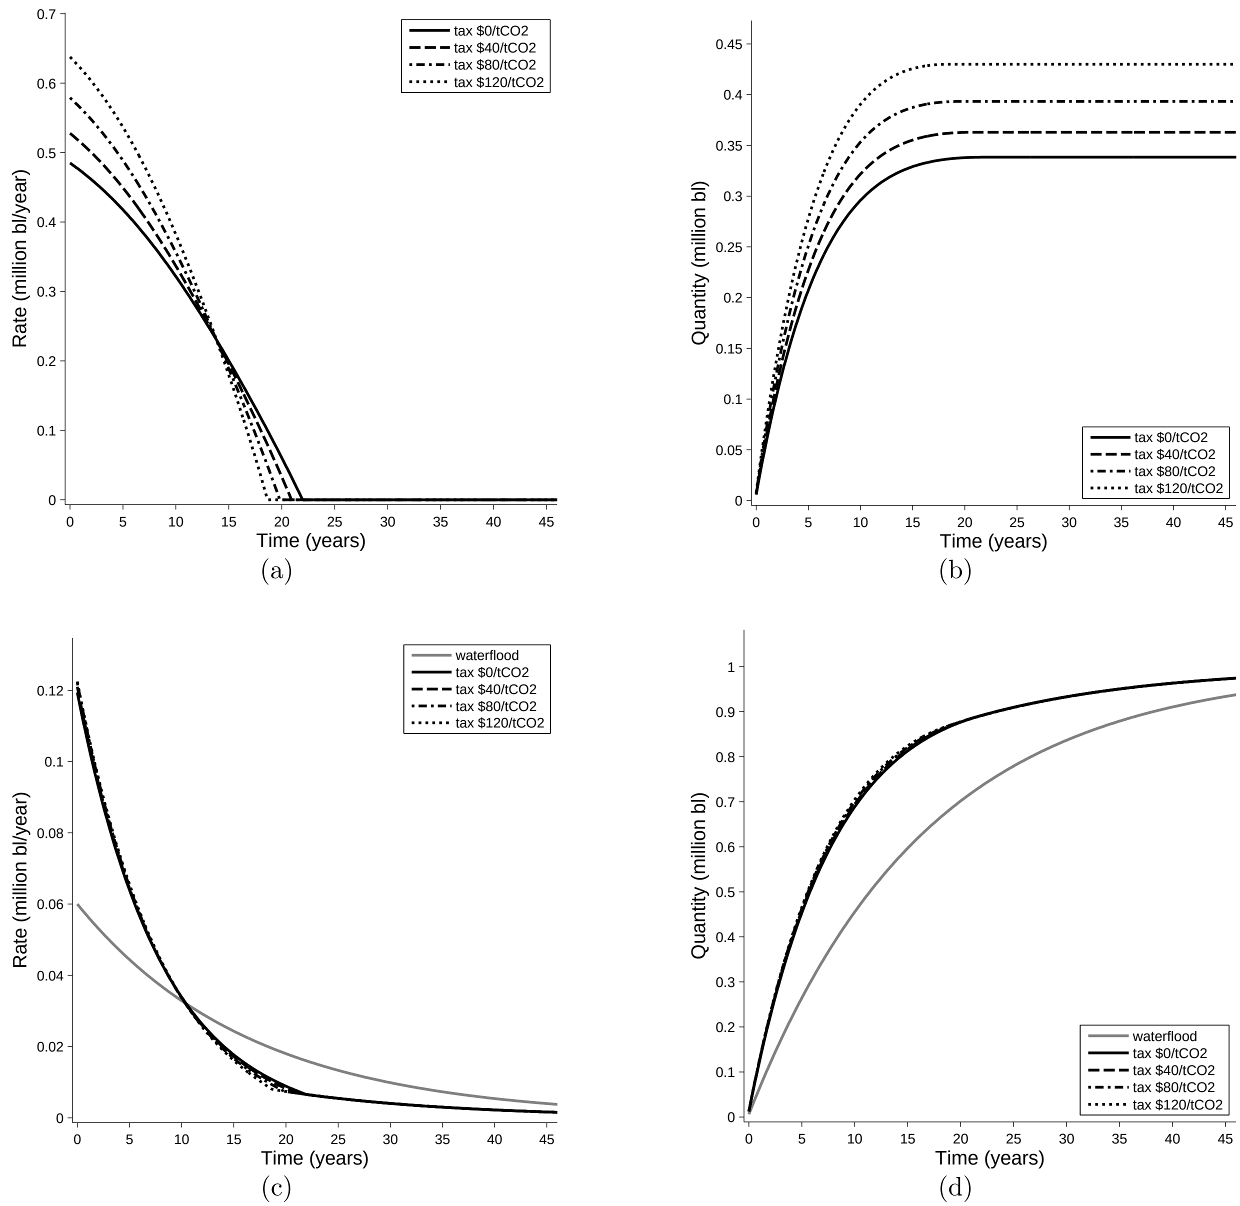

FIGURE 5. Change in (a) CO_2 injection, (b) cumulative sequestration, (c) oil production, and (d) cumulative oil production paths as a result of carbon tax changes.

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. Very high tax rates, in other words, induce a tradeoff between maximizing oil revenues and sequestration revenues.

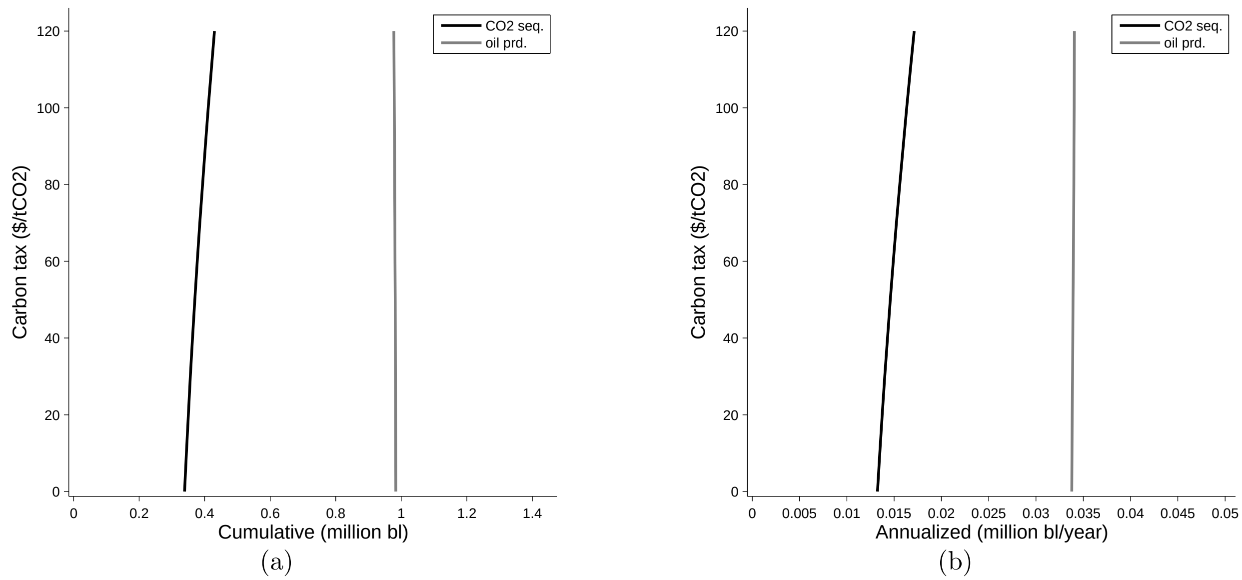

FIGURE 6. CO_2 sequestration and oil production, both cumulative (a) and annualized (b), as a function of the carbon tax.

Figure 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 subsection 5.1 suggest that the rates of CO_2 injection and sequestration 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 subsection 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 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 Figure 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. At the time of writing, current oil prices are about \$100/bl, while the current carbon price in the European market is about \$40/tCO_2. Thus, one can interpret the variations as corresponding to a doubling of current prices. The plots in panel (a) suggest a negative interaction effect: at both oil prices, a doubling of the carbon 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 CO_2-EOR projects, high oil prices are much more potent incentives for sequestration than are high carbon taxes. The reason is that 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 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 subsection 5.2, it is only at very high carbon-tax levels that sequestration revenues start to compete with oil revenues, driving the firm to increase CO_2 injection beyond the output-maximizing rate.

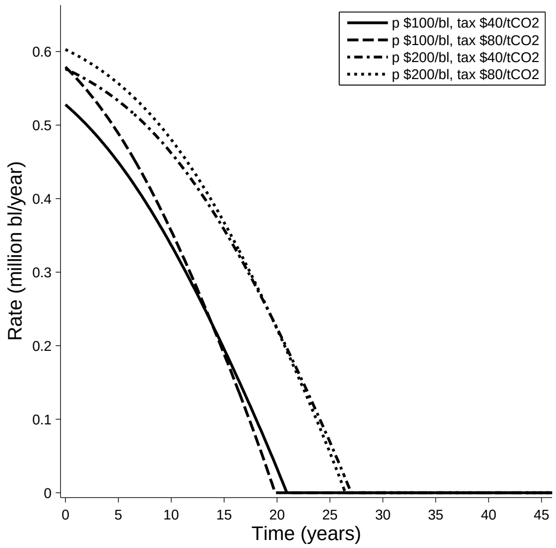

(a)

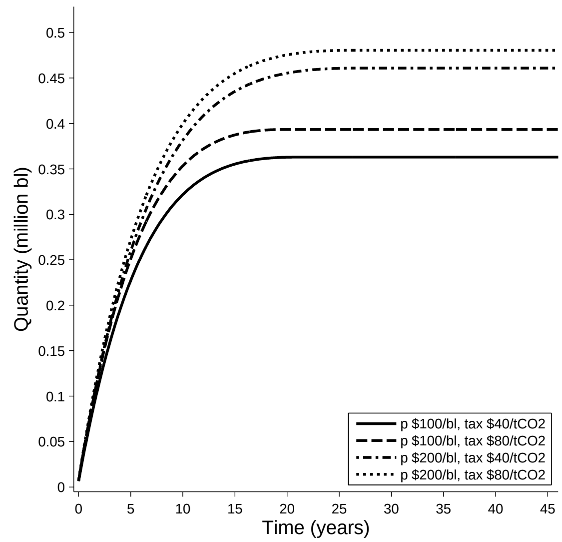

(b)

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

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 injection rate over time 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 injection rate will typically decline over time, and may eventually drop to zero before it becomes optimal to terminate the extraction process. 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 EOR. The bad news is that EOR-based carbon sequestration appears to be highly inelastic to carbon taxes. As such, there is little hope that policies raising the cost of CO_2 emissions will 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 well 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 reservoir, but not others, for a given combination of economic parameters.[15] In such a scenario, the supply of sequestration services for the oil reservoir 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.

An important caveat to the good news—the significant responsiveness of sequestration to oil prices—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.

In addition, the overall sequestration capacity of EOR projects, while quite large in absolute terms, is quite small in comparison to both overall CO_2 emissions and the capacity of other geological sequestration options. For example, the 12 $GtCO_2$ that Dooley et al. (2006) estimate as the theoretical sequestration capacity of all depleted U.S. oil reservoirs

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

(including those depleted through EOR) amounts to just two years' worth of U.S. CO_2 emissions (EPA; 2008). In contrast, the same study estimates the theoretical sequestration capacity of U.S. saline aquifers to be as large as 3,630 $GtCO_2$, making clear that in the long run, the main contribution to geological sequestration will have to come from such aquifers.

Nevertheless, 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.[16]

A key question we plan to address in future work is how large these EOR profits are likely to be. Clearly, analysis of this question will require expanding our model to account 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 (in the absence of a carbon tax) at which incurring these costs would be justified lies around \$50/bl.

A further extension concerns the effect of rising (rather than constant) oil prices and carbon taxes on both the optimal management of a CO_2 flood and the decision to initiate such a flood. In a stationary economic environment, there is never an incentive to delay switching to EOR—if doing so is not profitable at time 0, it will never be. Increasing prices may well induce such delay, however. Geologically heterogeneous projects will make the switch at different cutoff prices, thereby shifting out the aggregate supply of sequestration over time. Moreover, even projects that could profitably switch immediately at time 0 may optimally

[16] 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 ten 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 2.

choose to delay, if doing so increases the net present value of switching.[17] Interestingly, it seems likely that the latter type of delay may have the effect of *reducing* sequestration in a given reservoir. This is because the switch would occur at a lower remaining reserve stock, leaving less oil to be replaced by CO_2.

A final complication left for future work concerns the likely endogeneity of the reserve stock to 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. That is, oil that is impossible to flush out with a waterflood—referred to by reservoir engineers as "stranded" oil—may become recoverable once it mixes with CO_2. Clearly, 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 second state variable, thereby significantly complicating the analysis. How, if at all, this might alter the qualitative conclusions of the present paper is an open question.

[17] This is conceivable even if oil prices rise at rates below the discount rate, because the ability to delay up-front investments implies that rents may increase faster.

References

Al-Shuraiqi, H. S., Muggeridge, A. H. and Grattoni, C. A. (2003). Laboratory investigations of first contact miscible wag displacement: The effects of wag ratio and flow rate, *Technical 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.

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

Dooley, J. J., Dahowski, R. T., Davidson, C. L., Wise, M. A., Gupta, N., Kim, S. H. and 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. 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. and 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. and Fu, Y. (2006). Optimization of tertiary water-alternate-CO_2 flood in Jilin oil field of China: Laboratory and simulation studies, *Technical Paper 99616*, Society of Petroleum Engineers (SPE).

Hassoun, M. H. (1995). *Fundamentals of Artificial Neural Networks*, MIT Press, Cambridge, MA.

Hotelling, H. (1931). The economics of exhaustible resources, *Journal of Political Economy* **39**(2): 137–175.

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

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

McCoy, S. T. (2008). *The Economics of CO_2 Transport by Pipeline and Storage in Saline Aquifers and Oil Reservoirs*, Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA.

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.

Shaw, J. C. and 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. and 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.

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. and 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.