Chapter 3: Literature Review

Recent studies have estimated the potential farm-sector impacts of strategies to increase the quantity of carbon sequestered in agricultural soils and biomass. Findings in these studies suggest farm-sector sequestration activities that might be economically feasible at different prices. A review of the accounting and modeling procedures employed in the studies reveals how researchers have addressed permanence, C-stock equilibrium, and leakage, and, where possible, highlights the economic implications of alternative treatments (app. 1).

The general approach in past studies has been to construct a hypothetical situation in which farmers are paid to change land uses and/or production practices to store additional carbon in soils and biomass. Most of the analyses focus on a single carbon-sequestering activity. Parks and Hardie (1995), Alig et al. (1997), Stavins (1999), and Plantinga et al. (1999) assess the sequestration potential of afforesting marginal agricultural lands. Antle et al. (2001) look separately at shifting cropland to grasses and reducing summer fallow. Pautsch et al. (2001) focus on expanding the use of no-till systems. McCarl and Schneider (2001) and McCarl et al. (2003) present more comprehensive assessments in which carbon sequestration is part of agriculture's larger potential to mitigate GHG emissions. In these studies, farmers can adjust land uses, crop choices, and management practices in ways that increase carbon sequestration, decrease GHG emissions (i.e., CO_2, CH_4, and N_2O), or increase production of biofuel crops.

Discussion of Past Studies

In a 1995 study, Parks and Hardie simulate a strategy to afforest marginal agricultural lands patterned on USDA's Conservation Reserve Program (CRP). The strategy assumes farmers enroll lands in year 1 for a period of 10 years, receive a 50-percent cost-share payment to plant trees, and receive an annual payment to offset lost agricultural revenue. Region-specific carbon-sequestration values for lands converted to forests are developed by disaggregating 116.1 million acres of private nonprime farmland deemed suitable for forest into 433 regions, which are then linked to a map of forest types. Sequestration is defined narrowly to include only carbon stored in trees. Agricultural and forest rents reflect the discounted returns to lands in each use in each region. Thus, the decision to move land from agriculture to forest is based on a comparison of the longrun returns in each use.

The study is limited in its treatments of permanence, C-stock equilibrium, and leakage. It is assumed that 10 years after agricultural lands are planted to trees, timber production becomes the profit-maximizing land use. At that point, farmers will maintain the forests without additional Government payments. By assumption, the study sidesteps the permanence issue. Additionally, because the carbon accounting terminates at 10 years—which is shorter than the optimal rotation period in most regions—the study fails to address the issue of permanence raised by potential harvests. The 10-year timeframe of the analysis also sidesteps the C-stock equilibrium issue—since it takes decades for a newly planted forest to reach a carbon equilibrium state. Finally, because the model does not include any carbon-emitting activities that might occur in the farm or forest sectors in response to the afforestation incentives, the analysis does not consider the potential for agricultural- or forest-sector carbon leakage.

Parks and Hardy present results for four scenarios assuming an afforestation program funded at $456.2 million annually. The first two scenarios target both cropland and pasture under the enrollment criteria of minimizing the cost per ton of carbon sequestered and minimizing the cost per acre enrolled (based on a fixed per acre payment). The last two scenarios consider the same enrollment criteria but limit enrollment eligibility to cropland. The most cost-effective strategy bases payments directly on carbon sequestered and has the broadest land eligibility criterion. For the two scenarios that target both cropland and pasture, afforestation is 0.9 million acres less and carbon sequestration is 3.3 MMT more under the criterion of minimizing costs per ton sequestered than under the criterion of minimizing costs per acre (app. 1). More generally, for the strategy that targets carbon directly, the supply curve for sequestered carbon rises gradually from a level of 18.1 MMT and a marginal cost of $85 per mt to a sequestration level of 90.7 MMT and a marginal cost of $465 per MT.[1] For levels of sequestration above 90.7 MMT, the supply curve turns up sharply, suggesting an upper limit to the amount of carbon that can be economically removed from the atmosphere by afforesting agricultural lands.[2] When afforestation incentives are limited to cropland, enrollment and sequestration levels drop by over 55 percent, indicating that most afforestation would occur on lands now in pasture.

Alig et al. also analyze CRP-type incentives that would pay farmers to shift marginal agricultural lands to forest. Their framework—the Forest and Agricultural Sector Optimization Model (FASOM)—is an intertemporal market and spatial equilibrium model in which agriculture and forestry compete for the use of land. FASOM's forestry sector consists of nine geographic regions and various categories reflecting differences in product class, land ownership, forest type, site productivity, and management intensity. The model's agricultural component is a version of McCarl and Schneider's Agricultural Sector Model (ASM). Model simulations run for 90 years, but policy impacts are only reported for 50 years. Producers maximize the return to land by shifting land between agriculture and forestry and by changing forest management intensity (endogenous forest management decisions include harvest age, management intensity, and forest type).

FASOM runs on a decadal time step, allowing land to shift between agriculture and forestry every 10 years as the relative returns to these uses change over time. It addresses the permanence and leakage issues in that decisions on afforestation, deforestation, and forest management are endogenous to the model and the associated activities are tracked over time. Hence, the model accounts for changes in the net levels of carbon sequestration that occur over the years in each simulation. Conceptually, the C-stock equilibrium issue is reflected in the model's timber growth functions, although in the simulations, forests are typically harvested before this equilibrium is reached. In this study, sequestration incentives are based on gross rather than net sequestration—meaning landowners are rewarded for adopting practices that sequester carbon but are not penalized for adopting practices that emit carbon.

Among the model's simulations, the scenario with a fixed carbon flux target of 1.6 gigatons per decade has a relatively large economic impact. Over the 90-year simulation, the present value of losses for U.S. forest and agricultural consumers are $7.5 billion and $76.4 billion, respectively, while the present value of gains for U.S. forest and agricultural producers are $6.6 billion and $39.6 billion, respectively. During the first 50 years of the simulation, many owners of private forestland shift to earlier timber harvests, and some of this land moves into agricultural production. Overall, 59 million acres shift between the two sectors—with a net increase in forest area of 19 million acres. Alig et al. conclude that afforesting agricultural lands could be an effective component of a national GHG-mitigation strategy, but that the net carbon sequestration achieved will be much overstated if the program does not account for emissions from related forest-sector responses.

Plantinga et al. and Stavins also focus solely on the afforestation of agricultural lands. These studies develop econometric land-use models based on behavioral data. The methodologies lack an explicit link between afforestation of agricultural lands and related responses in forest product markets but do capture the impacts of various difficult-to-measure factors (e.g., becoming familiar with forestry) that make decisions to shift agricultural land into forest more "sticky" than if based strictly on maximizing expected profits. As in Alig et al., land moves between agriculture and forestry over time, and so these studies track changes in net carbon sequestration as land shifts between the two sectors—that is, they account for permanence and leakage. Stavins bases incentives on the *net* carbon sequestration from land-use decisions, while Plantinga et al. base incentives on *gross* sequestration. Again, C-stock equilibrium could be addressed in the model's timber growth functions, but forests are generally harvested before equilibrium is reached.

Plantinga et al. estimate separate models for Maine, Wisconsin, and South Carolina using data from various years between 1971 and 1996. The dependent variables are the share of county land in agriculture relative to the share in forest, and the share of county land in other uses relative to the share in forest. Independent variables include agricultural rents per acre (stated as the present discounted value of all net

---

[1] The figures $85 and $465 are obtained by setting T=20 tons (i.e., 18.1 mt) and T=100 (i.e. 90.7 mt), respectively, in the MC equation in Parks and Hardie (p. 130) and dividing the results by 0.8051 (to convert to 1997 dollars).

[2] We do not mean to imply here that any particular level of sequestration is economical. Only that the shape of the MC curve implies an upper limit to quantity of carbon that would be economical to sequester.

returns from agricultural uses over the simulation period), forest rents per acre (stated as the present discounted value of timber revenues), population, and land quality. Afforestation incentives are again patterned on the CRP and offer landowners annual payments to convert agricultural lands to trees for a period of 10 years starting at the beginning of each decade. Simulations start in 2000 and run for 60 years. Although not included in the simulation, program costs are adjusted to reflect a one-time tree establishment payment.

Results are presented for four scenarios reflecting different assumptions about population growth, timber harvest, and payment structure. Across scenarios and regions, results indicate that the marginal cost of carbon sequestration increases almost linearly with the level of sequestration. Marginal costs are generally lowest when population is held constant, sequestration payments are uniform across States, and harvesting is not allowed on enrolled lands. For this scenario and a payment of $45.09 per mt carbon, the States would have total sequestration of 1.36 MMT (Maine), 12.89 MMT (South Carolina), and 27.22 MMT (Wisconsin) (with future sequestration discounted at 5 percent).[3] At the same payment level but allowing for timber harvests, the discounted sequestration levels decrease to 1.04 MMT (Maine), 11.53 MMT (South Carolina), and 22.69 MMT (Wisconsin).

Stavins models the shares of county land in forest and in agriculture as functions of agricultural rents, farm production costs, existing forest area, conversion costs, and net returns to forestry. The model is estimated with panel data for 36 parishes/counties in Louisiana, Arkansas, and Mississippi covering the period 1935-1984. Model simulations run for 90 years, and the sequestration incentives include a payment for afforestation and a charge for deforestation.

For the three Delta States, Stavins finds that the marginal cost of sequestering carbon rises gradually to a sequestration level of about 6.35 MMT per year. At this level, the marginal cost of sequestration is about $79.86 per mt, net afforestation is 4.6 million acres, and the carbon payment/charge is about $109.81 per acre. Marginal costs rise steeply for sequestration levels above 6.35 MMT annually, becoming nearly asymptotic at about 14.52 MMT per year. Stavins concludes that sequestration would be competitive with abatement at lower levels of net emissions reduction.

Shifting from afforestation, Antle et al. and Pautsch et al. assess the economic feasibility of paying farmers to sequester carbon in agricultural soils by eliminating fallow periods from crop rotations, expanding grasslands, and increasing the use of no-till systems. The methodologies differ from the afforestation studies in two key respects. First, the geographic scopes are smaller—Pautsch et al. is limited to Iowa and Antle et al. focus on a subregion of Montana. Second, the crop production models used in these works are estimated with field-level data, which precludes these frameworks from comparing sequestration opportunities across regions. Still, the use of field-level data allows site-specific factors—or spatial heterogeneity—to be considered in the design and evaluation of incentives to encourage farmers to adopt the desired land uses and production practices.

Neither Antle et al. nor Pautsch et al. explicitly addresses C-stock equilibrium, leakage, or permanence. The studies omit the C-stock equilibrium issue because the timeframes of the scenarios are less than the period needed for most of the affected lands to reach new carbon equilibrium levels. Leakage is omitted because all land-management changes reflect changes to carbon-sequestering uses or practices. Regarding permanence, both studies implicitly assume that farmers continue the new land uses or production practices when the payments end.

Antle et al. assess the potential costs of paying farmers in Montana's dryland grain region to sequester carbon by switching from crop-fallow to continuous cropping rotations and by shifting cropland into permanent grass. The study analyzes each activity separately. In the scenarios, farmers commit to maintaining the new land use or production practice for a period of 20 years. Antle et al. develop econometric production models for wheat, barley, and permanent grass in both continuous crop and crop-fallow rotations using 1995 data for 425 farms and 1,200 fields. These models are then incorporated into a producer-decision framework that allocates land among the different rotations in response to payments to switch additional land to continuous cropping or permanent grasses. The study estimates changes in soil carbon by running the land-

---

[3] These dollar and sequestration values were obtained from the graphs in Plantinga et al. (pp. 820-21). We convert the results—presented in 1995 dollars and short tons—to 1997 dollars and metric tons.

use changes obtained in the model simulations through the Century crop-ecosystem model.

Three findings in Antle et al. are especially noteworthy. First, eliminating fallow periods appears to be a cost-effective method of carbon sequestration while expanding grasslands does not—at least in eastern Montana. Over a 20-year simulation, sequestering about 7 MMT of carbon by expanding grasslands requires an annual rental payment of $51.77 per acre and costs the Government $3.15 billion (not discounted). By contrast, sequestering the same amount of carbon by increasing continuous cropping requires an annual payment of $4.14 per acre and costs the Government $206.34 million (again, not discounted).

Second, a comprehensive GHG-mitigation strategy can include activities that sequester relatively small amounts of carbon but at a very low cost. At the highest payment level, switching from crop-fallow rotations to continuous cropping sequesters less than 19 MMT of carbon over the 20-year simulation. While this amount represents only a fraction of the sequestration potential of afforestation, two-thirds, or about 12 MMT, could be captured for an annual payment of $8 per acre.

Finally, for a given activity, the cost of sequestering carbon can vary significantly within a region due to site-specific biophysical and economic characteristics. Across locations in the study area, the marginal cost of sequestering carbon ranges from $51.15 to $511.51 per mt in the permanent grassland simulation and from $12.28 to $143.22 per mt in the continuous cropping simulation. While it is difficult to account for differences among fields and farms in national agricultural sector models, policymakers should note the potential for cost savings if sequestration incentives are designed with flexibility to take advantage of heterogeneity among locations in each region.

Pautsch et al. examine the cost of paying farmers in Iowa to expand use of no-till systems under alternative payment and program eligibility structures. Payment schemes include a uniform per acre payment and a variable per acre payment based on the amount of carbon sequestered. Program eligibility options include paying all users of no-till systems and limiting payments to new adopters.

Using field-level soil and weather data, county-level crop yield data, and State-level price and cost data, Pautsch et al. develop an econometric model in which the probability of adopting no-till is a function of net returns to conventional tillage, local soil characteristics, and regional temperature and precipitation variables. Production possibilities include 14 rotations consisting of mixes of corn, soybeans, wheat, sorghum, and hay. As in Antle et al., the changes in area under no-till are fed into a biophysical model to estimate changes in soil carbon.

Two findings in Pautsch et al. highlight the importance of payment structure and eligibility criteria in determining the costs associated with incentives to increase the quantity of carbon stored in agricultural soils. First, the average cost per metric ton of carbon sequestered is lower when per acre payments to adopt no-till are based on carbon sequestered rather than set at a fixed level. Parks and Hardie also support this finding, but to a lesser degree; in Pautsch et al., the average cost per unit sequestered is about four times lower under the price-discriminating structure than under the fixed payment structure. Second, cost savings are considerable when eligibility is limited to new adopters. In this analysis, paying all farmers who typically use no-till doubles the average cost per unit of carbon sequestered relative to paying only new adopters.

McCarl and Schneider expand the focus beyond a single carbon-sequestering activity, developing a comprehensive GHG-mitigation strategy that pays farmers to change land uses, crop mixes, crop-management practices, and livestock-management practices in ways that increase carbon sequestration, decrease GHG emissions (i.e., CO_2, CH_4, or N_20), or increase biofuel crop production. Their ASMGHG model—a market and spatial equilibrium mathematical programming framework—depicts production and consumption in 63 U.S. regions for 22 traditional crop commodities, 3 biofuel crops, 29 livestock commodities, and more than 60 processed agricultural products. In responding to relative changes in input prices, farmers can adjust tillage, fertilization, irrigation, manure treatment, and feed mixes. The study uses the biophysical Environmental Policy Integrated Climate (EPIC) model to calculate changes in carbon sequestration associated with changes in crop management activities. Thirty-year simulations of FASOM (see Alig et al., 1999) provide data on land shifting from agriculture to forestry and the associated quantities of carbon sequestered. Changes in emissions associated with livestock management are based on EPA data.

McCarl and Schneider simulate ASMGHG with carbon valued at $0.0 (i.e., the baseline), $9.60,

$48.10, $96.20, and $480.80 per MT.[4] These values are treated as a subsidy when farmers switch to activities that reduce GHG emissions and a charge when farmers change to activities that increase emissions. By design, the framework pays farmers only for the net GHG mitigation that results from their response to the full set of incentives. Additionally, the link between the agricultural and forestry sectors allows ASMGHG to reflect the forest sector's response to the afforestation of agricultural lands—at least in the long run. Hence, the framework explicitly accounts for carbon leakage through forest-sector activities. While McCarl and Schneider do not explicitly address the permanence and C-stock equilibrium issues, they adapt the framework to address these issues in McCarl et al. (2003).

For carbon valued at $9.60, $48.10, $96.20, and $480.80 per MT, net carbon sequestration is, 51.80, 146.40, 238.50, and 395.50 MMT, respectively, while net GHG mitigation is equivalent to 53.90, 154.10, 255.70, and 425.90 MMT of carbon, respectively. McCarl and Schneider conclude that paying farmers to sequester carbon, reduce GHG emissions, and increase biofuel crop production would positively affect farmers and negatively affect U.S. consumers. For carbon valued at $9.60, $48.10, $96.20, and $480.80 per MT, farm welfare increases $ 0.40, $4.30, $13.40, and $76.90 billion, respectively, while U.S. consumer welfare decreases $0.40, $5.20, $18.50, and $104.60 billion, respectively. More generally, at low carbon prices soil carbon sequestration, afforestation, and CH_4 / N_2O emissions reduction dominate GHG-mitigation activities. At high carbon prices, the dominant mitigation activities are afforestation and biofuel production. Regardless of the value assigned to carbon, CH_4 and N_2O emissions reduction activities make a relatively small contribution to GHG mitigation.

More recently, Jones et al. (2002) and McCarl et al. (2003) identified approaches for treating permanence and C-stock equilibrium in static models that do not implicitly assume sequestration is permanent.[5] As described in chapter 2 of this report, this approach employs a series of annual payments based on a pay-as-you-store principle. In essence, a discount—based on the length of the contract period for carbon sequestration and the choice of discount rate—is applied to the value of a unit of emissions reduction to determine the value of a unit of carbon sequestration during the contract period. Our analysis also employs this approach to valuing sequestered carbon.

Conclusions

Differences in scope (including geographic region, sector, activity, and GHG coverage), methodology, and underlying assumptions make it difficult to directly compare the results of previous studies of the economic potential to sequester carbon in agriculture. Recognizing these limitations, however, the literature does provide some insights regarding carbon-sequestering land uses and production practices that may be competitive at different carbon prices and the levels of carbon that might be sequestered at those prices. Additionally, the studies reveal how researchers have addressed permanence, C-stock equilibrium, and leakage.

Taken collectively, the studies reviewed here suggest that the farm sector's economic potential to sequester additional carbon is significantly less than amounts deemed technically possible in soil science-based assessments. For example, for the United States as a whole, the studies cited in table 2.2 estimate the technical potential for sequestering carbon at 35-107 MMT per year for expanding conservation tillage and 6-18 MMT per year for eliminating summer fallow and changing rotations. In contrast, for expanding no-till in Iowa, Pautsch et al. estimate the cost of sequestering 1 MMT of carbon at about $200 per MT. On the national scale, McCarl and Schneider estimate the maximum economic potential of about 70 MMT per year for all farm-sector soil management activities at a carbon price of $500 per MT.

The most noteworthy divergence between soil science and economic assessments concerns conversions of cropland to permanent grasses. From a national perspective, Eve et al. (2000) estimate the technical potential of converting cropland to grassland at between 26 and 54 MMT of carbon per year. In contrast, economic assessments by Antle et al. and McCarl and Schneider find that sequestering carbon via this land use change would not be competitive with other carbon-sequestering activities.

---

[4] Money values in McCarl and Schneider are in 2000 dollars. We have converted their values to 1997 dollars. To obtain the values reported in McCarl and Schneider, multiply the values reported here by 1.04.

[5] The Jones et al. citation refers to a presentation at the Forestry and Agriculture Greenhouse Gas Modeling Forum in Shepherdstown, WV (October 9-11, 2002). This presentation and the paper by McCarl et al. are not included in appendix 1 because of similarities to, respectively, the present analysis and McCarl and Schneider.

The published studies indicate that the most cost-effective mix of carbon-sequestering activities will depend on the level of carbon payment offered—or equivalently, the target quantity of total sequestration to be achieved by the program. Across studies, changes in production practices—such as expanding no-till and shifting to carbon sequestering rotations—dominate farm sector responses at very low payment levels. Afforestation becomes the dominant sequestration activity at a carbon payment between $20 and $100 per MT, depending on the specific features of the context modeled. McCarl and Schneider, the one study that considers multiple GHGs, finds that CH_4 and N_2O emissions reduction activities become feasible as carbon payments approach $50 per MT. Above payments of $100 per mt, additional emissions reductions from these activities become very limited at any price. McCarl and Schneider also find that production of biofuel crops starts to become economically attractive at a payment of about $50 per mt and joins afforestation as a dominant GHG-mitigation activity at payments above $75 per mt.

The treatment of permanence, C-stock equilibrium, and leakage in previous studies is noteworthy for several reasons. First, the finite time period for annual increments to the carbon stock in agricultural soils is generally sidestepped, without comment. In the predominant static frameworks, the sequestration programs generally end before affected soils or forests reach their new carbon equilibrium levels. The longrun dynamic models employed for forestry explicitly address the issue because they typically incorporate a carbon-stock accumulation function by age of timber.

Treatment of permanence varies in previous studies. Studies that employ dynamic frameworks generally track changes in net carbon storage over time; however, some of the forestry modeling has not employed negative charges for emissions from terrestrial storage. In dynamic models, permanence can be incorporated in sequestration incentives by paying landowners when they switch to land uses or production practices that sequester carbon and charging them when they switch to uses and practices that emit carbon. In static models, permanence can be accounted for by using the rental price for a specified commitment period, derived from the full asset value employed for permanent sequestration. In our analysis, which employs a static model, we follow this approach to account for the temporary nature of carbon sequestration in the design of sequestration incentives.

The literature also varies in the treatment of leakage. By design, models that are limited to a single land use or production practice cannot account for leakage related to changes in other land uses or production practices. Similarly, single-sector models cannot account for leakage related to changes in land uses or production practices in other economic sectors. Past farm sector studies (Parks and Hardie, Antle et al., and Pautsch et al.) have generally focused on a single carbon-sequestration activity, and the carbon accounting has not included activities outside of the incentive program to check for leakage. With such limited coverage of activities, the potential for leakage from other activities in agriculture or forestry is great. In contrast, the forestry literature, which uses dynamic longrun models in which agriculture and forestry are linked via competition for land (Alig et al., Stavins, and Plantinga et al.), has examine the issue. These studies generally find that forest sector responses to farm sector carbon sequestration incentives can be an important source of leakage—particularly at higher carbon payment levels.

Our framework is limited to the agricultural sector, but it does account for the net effect on sequestered carbon after farmers have adjusted for land-use changes among crop, grass, or forest land and for a variety of crop land-management practices. It does not, however, capture forest-sector leakage, so we acknowledge the potential for upward bias in our sequestration results relative to a full accounting of the agriculture and forestry sectors. With respect to past studies, we note that the magnitude of forest-sector leakage is critically linked to assumptions previous researchers have employed regarding the longrun (50-100 years) time paths of income growth, population growth, technological change, price expectations, and consumer preferences. Different assumptions about the longrun time paths of these variables can have very different implications for the longrun demand for land in the two sectors and thus very different implications concerning potential leakage.

Finally, researchers to date have formulated hypothetical incentives—often based on those in USDA's CRP—which are then incorporated into economic models and simulated for various exogenously specified carbon payments. The result is to trace out scenario-specific supply curves for sequestered carbon. Our analysis will also trace out a farm-sector supply curve for sequestered carbon.