*The Australian Journal of Agricultural and Resource Economics*, 50, pp. 1–15

# Modelling the supply of ecosystem services from agriculture: a minimum-data approach*

John M. Antle and Roberto O. Valdivia†

We argue that to support agriculture–environmental policy decision making, stakeholders need 'quantitative back-of-the-envelope' analysis that is timely and sufficiently accurate to make informed decisions. We apply this concept to the analysis of the supply of ecosystem services from agriculture. We present a spatially explicit production model and show how it can be used to derive the supply of ecosystem services in a region. This model shows that the supply of ecosystem services can be derived from the spatial distribution of opportunity cost of providing those services. We then show how this conceptual model can be used to develop a minimum-data (MD) approach to the analysis of the supply of ecosystem services from agriculture that can be implemented with the kinds of secondary data that are available in most parts of the world. We apply the MD approach to simulate the supply of carbon that could be sequestered in agricultural soils in the dryland grain-producing region of Montana. We find that the supply curve derived from the MD approach can approximate the supply curve obtained from a more elaborate model based on site-specific data, and can do so with sufficient accuracy for policy analysis.

**Key words:** carbon sequestration, ecosystem services, minimum data.

## 1. Introduction

Increasingly, the focus of agricultural policy is shifting from traditional subsidy and trade policies to conservation and environmental aspects of agriculture (i.e., to policies that provide farmers with incentives for increasing the supply of ecosystem services from agriculture). This shift in policy emphasis is explained partly by a growing public demand for these ecosystem services that are public goods, including wildlife habitat, visual amenities and open space, water quality protection, and greenhouse gas mitigation. This shift in policy focus also has been encouraged by the incorporation of agriculture into the General Agreement on Tariffs and Trade in the mid-1990s, and the ongoing negotiations to reduce industrialised-country agricultural subsidies in the Doha Round of negotiations led by the World Trade Organisation.

* This research was supported in part by the U.S. Agency for International Development Soil Management Collaborative Research Support Program and by Montana State University.

† John M. Antle (email: jantle@montana.edu) is a Professor and Roberto O. Valdivia is a Research Associate in the Department of Agricultural Economics and Economics, Montana State University, 312 Linfield Hall, PO Box 172920, Bozeman, MT, USA. This paper is based in part on the Distinguished Fellows Lecture Presented by J. Antle at the 2005 meetings of the Australian Agricultural and Resource Economics Society.


doi: 10.1111/j.1467-8489.2006.00315.x

Just and Antle (1990, p. 197) proposed a conceptual framework for analysis of agriculture–environment policy interactions: 'This framework integrates physical and economic models at a disaggregate level necessary to capture the heterogeneity of the physical environment and the economic behaviour of farmers'. Since then, a number of researchers have utilised site-specific data and models to implement analysis of agriculture–environment interactions and related policies, consistent with that conceptual framework (e.g., Pautsch *et al*. 2001; Antle *et al*. 2003; Feng *et al*. 2004; Wu *et al*. 2004; Lubowski *et al*. 2005). However, high-resolution biophysical and economic data with the geographic coverage needed for analysis of agriculture–environment interactions, such as the National Resources Inventory data in the United States, are exceptional, and provide limited economic information. In most cases, site-specific economic data are only available from special-purpose farm surveys, and the time and resources required to undertake special-purpose surveys precludes their use for most policy analysis.

In this paper, we develop a minimum-data (MD) approach to the analysis of the supply of ecosystem services from agriculture that can be implemented with data that are readily available in most parts of the world from existing secondary sources. Due to the need for timely analysis, often policy decisions must be based on qualitative analysis. The MD approach is motivated by the demand for timely, quantitative analysis of agriculture–environment interactions that can provide sufficiently accurate analysis of policy scenarios to support informed policy decision making. There is always a trade-off between timeliness and accuracy of information for policy decision making. Our experience suggests that policy decision makers need quantitative, 'back-of-the-envelope' analysis that is accurate within, say, an order of magnitude. The appropriate level of accuracy for policy analysis is arguably lower than standards acceptable for publication in scientific journals because *ex ante* policy modelling involves a large number of unquantifiable uncertainties, and because the policy process utilises many sources of information and must balance competing political interests.

In the first section of this paper, we present a spatially explicit production model and show how it can be used to derive the supply of ecosystem services in a region. This model shows that the supply of ecosystem services can be derived from the spatial distribution of opportunity cost of providing those services. In the second section of the paper, we show that under some plausible simplifying assumptions, the spatial distribution of opportunity cost can be approximated with a small number of parameters that can be estimated with the kinds of data that are generally available from secondary sources. We apply this MD approach to simulate the supply of carbon that could be sequestered in agricultural soils in the dryland grain-producing region of Montana. We find that the supply curve derived from the MD approach provides a good approximation to the supply curve obtained from an econometric-process simulation model based on site-specific data.

## 2. Modelling the supply of ecosystem services

The analysis of the supply of ecosystem services begins with the observation that farmers' land use and management decisions may affect biological and physical systems through a number of mechanisms. These effects may be limited to the land owned by the farmer, such as a change in soil productivity, or may have off-site effects, such as chemical run-off into surface waters. Absent policies that affect farmers' incentives, we assume farmers make land use and management decisions to maximise their perceived economic well-being. These decisions result in a supply of ecosystem services that is determined by farmers' economic incentives to supply market goods (crops and livestock), but does not take into account society's valuation of the ecosystem services. To increase the supply of ecosystem services beyond this private equilibrium, demanders of ecosystem services must provide farmers with incentives to change their management decisions in ways that increase those services. In most cases, ecosystem services are public goods, so some form of government intervention or assignment of property rights will be needed to create these incentives. For example, in the case of greenhouse gas mitigation through carbon sequestration, a government program could pay farmers to adopt practices that sequester carbon. Alternatively, a government regulation that caps greenhouse gas emissions could create a market for carbon emissions reduction credits, providing farmers with an incentive to sequester carbon.

We consider a simple model of a farmer's choice between two competing land uses, $a$ and $b$, in a geographic region. These land uses can be any kind of productive activity or a conserving use of the land, as long as a value can be associated with it. For example, if crop rotation is involved, then the appropriate value could be an annualised value of the rotational system. The land-use decision in each time period is based on the maximisation of expected value $v(p, s, z)$ where $p$ is a parameter (more generally a vector), interpreted here as an output price; $s$ indexes the site and $z = a, b$ indexes the activity at the site. This value function could be a linear or non-linear function of returns or any other objective function. The model can be generalised in various ways, for example by introducing risk aversion (interpreting $v$ as an expected utility function), or by interpreting $v$ as the expected present value of returns for analysis of investment problems. We assume for simplicity that adjustment costs associated with changing from one land use to another are zero (we discuss the implications of relaxing this assumption below). Under these assumptions, activity $a$ is chosen if $\omega(p, s) = v(p, s, a) - v(p, s, b) \geq 0$ and activity $b$ is chosen otherwise.

We assume that an ecosystem service of $e(s)$ units per hectare per time period is produced at each site $s$ when practice $b$ is in use, and that zero services are produced if practice $a$ is in use at the site. We also assume practice $a$ produces zero ecosystem services without loss of generality, since the analysis here is based on the difference between the two practices. Alternatively, it can be assumed that practice $a$ also produces a positive service flow, but practice *b* produces more ecosystem services than practice *a*. The parameter *e* is interpreted as the expected rate of ecosystem services obtained from changing land use, not as the realised supply of ecosystem services. Following Antle *et al*. (2003), we assume that because of costs associated with measuring and verifying changes in ecosystem services, contracts for provision of ecosystem services are based on an *ex ante* estimate of ecosystem services within relatively homogeneous agro-ecozones.

To derive the expected private equilibrium supply of ecosystem service in the region, we define the density function $\varphi(\omega)$ by ordering all land units according to the value of $\omega(p, s)$ for a given a value of $p$. The proportion of land units in activity $b$ is then given by

$$r(p) = \int_{-\infty}^{0} \varphi(\omega)\, d\omega,\ 0 \le r(p) \le 1, \tag{1}$$

where the dependence of $r$ on $p$ follows from the fact that $\omega(p, s)$ is a function of $p$. The expected private-equilibrium supply of ecosystem services per time period in the region with $H$ hectares of cropland is then given by

$$S(p) = r(p)\ H\ e. \tag{2}$$

Equation (2) shows that the private equilibrium, or baseline, supply of ecosystem services, $S(p)$, can be defined as a function of the parameter vector $p$, the number of hectares on which practice $b$ is used and the average rate of service production on those sites where practice $b$ is used.

Following Antle *et al*. (2003), to increase the supply of ecosystem services above the baseline quantity $S(p)$, we assume that a payment $p_e$ (\$/$e$) is offered to the land managers by a private or government entity for increasing the quantity of the ecosystem service. Without loss of generality, we assume that farmers are paid only for an *increase* in ecosystem services relative to a baseline, as might be the case for carbon sequestration when contracts reward only additional carbon sequestered in response to the offer of an environmental payment. Alternatively, it could be assumed that farmers who adopted practice $b$ before environmental payments were offered are paid for the ecosystem services they produce. This policy would increase the cost of producing any given quantity of ecosystem service by the amount $p_eS(p)$. Also note that an efficient incentive system would pay farmers a positive price for increasing ecosystem services and tax them for any actions that reduced them. This type of symmetric incentive could be modelled in essentially the same way as the positive incentive example presented here.

Note that the amount of ecosystem service supplied at each site is not known *ex ante*, and therefore payments must be based on expected ecosystem services. The estimated quantity of ecosystem services could be estimated on a site-specific basis if sufficiently good data were available. Alternatively, if the goal of the buyer of ecosystem services is to obtain a total quantity in the region and not to obtain a specific amount at a given site, payments could be based on an average rate of services for the region that could subsequently be verified on a regional basis through a statistically-based sampling and measurement scheme (e.g., for a discussion of this type of contract for soil carbon sequestration, see Antle *et al*. 2003). Here we make this latter assumption, and therefore the landowner receives a value of $v(p, s, a)$ for using practice $a$ and $v(p, s, b) + p_e\, e$ for using practice $b$, where $e$ is interpreted as the expected amount of services produced with practice $b$. Therefore, the farmer will choose activity $b$ if $\omega(p, s) - p_e\, e < 0$. The following three cases may occur at each site $s$:

*Case 1.* Activity $b$ is more profitable without payments ($\omega(p, s) < 0$), and therefore is also more profitable with payments for ecosystem services. At these sites, farmers adopt practice $b$ without an environmental payment.

*Case 2.* Activity $a$ is more profitable without a payment for ecosystem services, but activity $b$ is more profitable with environmental payments. This is the case where $\omega(p, s) > 0$, but $\omega(p, s) - p_e\, e < 0$. Thus, the decision maker will switch to activity $b$ if the payment per unit of additional ecosystem service is greater than the opportunity cost per unit of additional ecosystem service (i.e., if $p_e > \omega(p, s)/e$).

*Case 3.* Activity $a$ is more profitable whether or not environmental payments are received. Thus, $\omega(p, s) > 0$, $\omega(p, s) - p_e\, e > 0$ and $\omega(p, s)/e > p_e$. This last inequality means that the opportunity cost per unit of ecosystem service is always positive and greater than the payment per unit of ecosystem service, hence the farmer will not enter the land unit into a contract that would require switching from practice $a$ to practice $b$.

The site-specific land-use decisions can be linked to the regional supply of ecosystem services as shown in Figure 1. Using the spatial distribution of opportunity cost defined above, $\varphi(\omega)$, we can make a change of variable and define the spatial distribution of opportunity cost per unit of ecosystem service as $\phi(\omega/e) = \varphi(\omega)e$, as illustrated in the left side of Figure 1 (Freund 1962, p. 120). The area under the spatial distribution of opportunity cost on the interval $(-\infty, 0)$ equals $r(p)$ and represents those land units where farmers use practice $b$ without environmental payments (case 1). Thus, at the point where $p_e = \omega/e = 0$, the baseline supply of ecosystem services equals $S(p)$. Those land units corresponding to the range of opportunity cost between zero and $p_e$ will switch from activity $a$ to $b$ and thus increase the supply of ecosystem services to a quantity greater than $S(p)$ (case 2). Define this proportion of the land area as

$$r(p, p_e) = \int_0^{p_e} \phi(\omega/e)\, d(\omega/e). \tag{3}$$

The supply of ecosystem services at price $p_e > 0$ is equal to

$$S(p, p_e) = S(p) + r(p, p_e)\, H\, e. \tag{4}$$

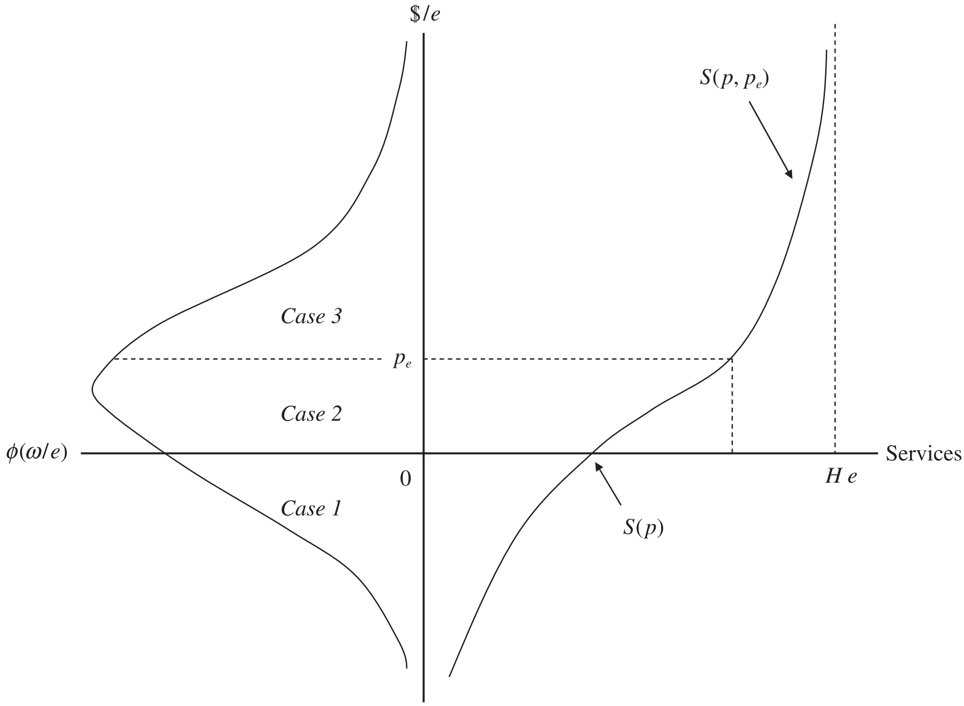


**Figure 1** Derivation of the supply of ecosystem services from the spatial distribution of opportunity cost per unit of ecosystem services.

Those land units where opportunity cost is greater than $p_e$ will remain in activity *a* (case 3). As $p_e$ increases, $r(p, p_e)$ increases and approaches $1 - r(p)$.

Equation (4) shows that the total quantity of ecosystem service is equal to the baseline quantity, $S(p)$, plus the additional quantity supplied, $r(p, p_e)\ H\ e$, due to the positive incentive. In some contexts, such as greenhouse gas mitigation, buyers of emission offsets get credit only for additional carbon sequestered above and beyond the baseline quantity, so the supply curve for carbon sequestration would be defined as $S(p, p_e) - S(p) = r(p, p_e)\ H\ e$.

Figure 1 shows that the shape of the supply curve is determined by the form of the spatial distribution of opportunity cost and the point at which the distribution crosses the origin. The supply curve approaches a vertical asymptote equal to the maximum amount of ecosystem service ($H\ e$) that can be produced when every land unit switches to activity *b*. The variance of the opportunity cost of changing practices plays a critical role in determining the shape of the supply curve of ecosystem services, as illustrated in Figure 2. When the variance is positive, the supply curve has a positive slope, with its concavity depending on the position of the distribution of opportunity cost in relation to the origin. As the variance increases, the supply curve rotates counterclockwise about the point A. As the variance decreases and approaches zero, the supply curve approaches the shape of a step function with the step

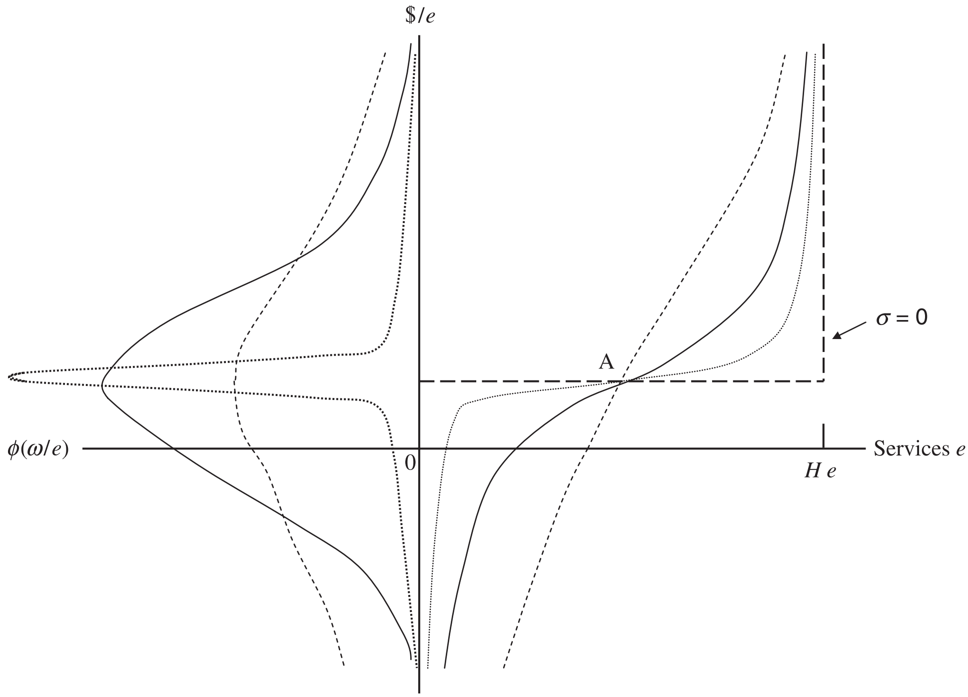


**Figure 2** Effect of the variance of the opportunity cost of changing practices on the supply curve for ecosystem services.

occurring at the value of $\omega/e$ where the mass of the distribution lies. This limiting case of a zero variance is equivalent to a representative farm model applied to all land units in the region. If $p_e < \omega/e$, then none of the land units enter contracts to supply additional ecosystem services above the baseline amount, so the supply curve would lie on the vertical axis. At the point where $p_e = \omega/e$, all land units where practice $b$ had not already been adopted would switch from $a$ to $b$ and so the supply function would increase to the maximum $H\ e$ at that point, and become vertical again for all higher prices.

If the rate of ecosystem services were known on a site-specific basis, then the spatial correlation between the opportunity cost and ecosystem services would need to be taken into account in deriving the spatial distribution of opportunity cost per unit of ecosystem service. For example, see Dorrman *et al.* (1990) who show that an approximation to the variance of a ratio of random variables can be made using their variances and covariance.

In addition to changes in expected returns, the opportunity cost of changing practices could include fixed costs of capital needed to utilise the new practice (defined here as an annualised value per hectare, $F$). Also, participation in government programs or ecosystem service markets could involve transaction costs ($TC$, defined here as an annual cost per hectare). The farmer now enters a contract if

$$v(p, b) + p_e e - F - TC > v(p, a)$$

or if

$$p_e > \omega(p, s)/e + (F + TC)/e. \quad (5)$$

Thus, fixed and transaction costs create a threshold value of $p_e$ below which it will not be profitable for farmers to switch practices. In the case where $e$ is a constant expected value in the region, this effect shifts the supply curve upwards (or leftwards) so that the baseline value $S(p)$ occurs at this positive threshold price $p_e = (F + TC)/e$ rather than at $p_e = 0$. In the case where $e$ varies spatially, the threshold price for the supply curve will equal $(F + TC)/e_{\max}$ where $e_{\max}$ is the maximum carbon rate in the region.

Finally, we observe that fixed costs of adoption and transaction costs are difficult to observe and are difficult to estimate *ex ante*. Similarly, in developing countries there may be a variety of factors, such as market imperfections, that imposed unobserved costs on farmers and thus constrain adoption of certain practices. The logic of this model suggests a simple method to estimate the population average of these adoption and transaction costs. To illustrate, suppose that data are available to estimate $\omega(p, s)/e$ (as discussed in the following section), and using that information the model predicts that the proportion of farmers adopting the practice should be $r(p, 0) > 0$. However, suppose that we observe that a smaller proportion of farmers, say $r_C < r(p, 0)$, are actually using the practice. We could infer that this difference is due to unobserved costs that shift the distribution of opportunity cost leftward and cause the model without these costs to overpredict adoption. Thus, the value of the fixed cost term sufficient to set $r_C = r(p, 0)$ could be interpreted as an estimate of the unobserved costs of adoption.

## 3. Minimum-data methods to simulate the supply of ecosystem services

The analysis in the preceding section shows that the economic information needed to simulate land-use decisions is the expected value for each competing land-use activity at each site. As noted in the introduction, one solution is to use site-specific data to estimate and then simulate an appropriate site-specific behavioural model, as has been performed in the literature cited there. Site-specific land-use and management decisions can be combined with outputs from biophysical simulation models to estimate the quantity of expected ecosystem service supplied on each land unit, $e$, for a given economic incentive (price) for that service. Repeating this analysis over a range of prices, the analyst can aggregate these quantities across land units and thus derive the regional supply curve for ecosystem services. A critical difficulty with this approach, as noted in the introduction, is that it requires sufficient site-specific data on inputs, outputs and prices of individual production activities to estimate behavioural equations. Such data are rarely available on the regional or national basis needed for policy analysis.

The alternative approach proposed here is to use secondary data (e.g., data available from governmental agencies), other available data (e.g., data on yield variability), expert judgement as necessary, and sensitivity analysis, to parameterise directly the spatial distribution of net returns to competing activities, and then simulate land allocation decisions using suitable decision rules (e.g., maximisation of expected returns or expected utility). Often secondary data are available for 'average' or 'representative' costs and returns for a geographical region such as a county or agro-ecozone. In the MD approach proposed here, secondary data are used to estimate mean expected net returns to each activity in each region. In addition, estimates of spatial variability in expected returns are needed. In the case study presented below, we assume that farmers in a region form similar output price expectations (e.g., based on future prices adjusted for transportation costs), and face similar factor prices and average costs of production. In addition, we assume that variable costs of production are approximately proportional to expected output. For example, Antle and Capalbo (2001) found that cost functions for wheat and barley production exhibited costs of production with approximately unitary output elasticities, and this is typical of field-scale or farm-scale econometric production studies. Therefore, we can plausibly assume that cost of production $c \approx \kappa\, y$, where $\kappa$ is a constant across space, and $y$ is yield. It then follows that for a given expected output price $p$, net returns are $v = py - c \approx (p - \kappa)y$, hence, it also follows that the coefficient of variation in net returns across land units in a region (at a point in time) can be estimated by the spatial coefficient of variation for $y$. Available data show that this approach provides an approximation that is well within an order of magnitude. For example, using field-level data and simulations from the study by Antle and Capalbo (2001), the spatial coefficients of variation for winter wheat and spring wheat yields in Montana are 79 and 91 per cent, whereas the coefficients of variation in simulated net returns are 90 and 86 per cent (Table 1). Data from a field-level survey in Peru show that coefficients of variation for potato, grain and legume crops are 135, 97 and 122 per cent, whereas simulated net returns had coefficients of variation of 151, 103 and 144 per cent. For Peru, crop simulation models were also used to estimate yields for each field in the survey data. These simulated yields had coefficients of variation 54 to 69 per cent lower than the observed yields. Table 1 also shows county-level yields for these crops for Montana and Peru. As expected, the aggregated data show lower variation than the field data. These data suggest that simulated, field-level net return variation can be approximated with field-level yield variation, but simulated variances and variances from aggregated data underestimate observed field-level yield variability by a factor of 1.5 to 3. Therefore, if researchers have only simulated or aggregated yield data available, estimates of variability will need to be increased, perhaps by a factor of two, although further research will need to be conducted to see if this rule of thumb is more generally valid.

As shown in the previous section, land-use decisions in a region are determined by the spatial distribution of the difference in expected value $\Delta v(p, s)$.

**Table 1** Coefficients of variation for actual and simulated yields and net returns for crops in Montana and Peru

| | Actual yield (field) | Actual yield (county) | Simulated yield | Simulated net returns |
|---|---|---|---|---|
| Montana | | | | |
| Winter wheat | 79 | 25 | n.a. | 90 |
| Spring wheat | 91 | 32 | n.a. | 86 |
| Peru | | | | |
| Potato and tubers | 135 | 43 | 53 | 151 |
| Grains | 97 | 67 | 30 | 103 |
| Legumes | 122 | 77 | 56 | 144 |

Source: Montana field data and simulation models reported in Antle and Capalbo (2001); Peru field data and simulation models described in Valdivia (2002). Montana county coefficients of variation (CV) are for 1999–2001 crops. Peru CV are for 2001–2003 crops grown in departments in the sierra with characteristics similar to the Cajamarca region where the field data were collected.

The expectation of this difference is simply $E(\Delta v) = E[v(a)] - E[v(b)]$, and the variance of the difference between returns to $a$ and $b$ is $\sigma_{a-b}^2 = \sigma_a^2 + \sigma_b^2 - 2\sigma_{ab}$ Although secondary data often can be used to estimate the variances $\sigma_a^2$ and $\sigma_b^2$ as discussed above, it may be more difficult to obtain data to estimate the covariance $\sigma_{ab}$. To gain some insight into how this covariance enters the analysis, suppose that $\sigma_a^2 \approx \sigma_b^2 = \sigma^2$, as is likely to be the case, for example, when the two practices involve growing the same crop with a different tillage practice or a different use of fallow in the rotation. Substituting $\sigma^2$ into the expression for $\sigma_{a-b}^2$ it follows that $\sigma_{a-b}^2 \approx 2\sigma^2(1 - \rho_{ab})$ where $\rho_{ab}$ is the correlation between returns for activities $a$ and $b$. Thus, as the correlation between the expected returns approaches 1, the variance of their difference approaches zero, and the supply curve approaches a step function, as illustrated in Figure 2. As the correlation approaches zero, the variance of their difference approaches $2\sigma^2$ and the supply curve takes on a positive slope. In many cases the spatial correlation between returns to alternative practices is likely to be relatively high – for example the returns to wheat grown with conventional tillage should be highly but not perfectly correlated with the returns to wheat grown with conservation tillage, as suggested by research showing that in some cases conservation tillage results in higher yield variability than conventional tillage (Uri 2000). In other cases, the correlation will be low and approach zero, for example in the case where one of the alternative land uses is participation in a government conservation program with a fixed return.

Given an estimate of the mean and variance of the difference in expected returns between the two competing activities, we can proceed to simulate land-use decisions in a region by parameterising the distribution of opportunity cost and using Equations (1) and (2) to derive the supply curve. This can be accomplished under the assumption of normally distributed returns by simply noting that the difference of two normally distributed random variables is itself normal. Alternatively, in cases where either normal or non-normal distributions are used (e.g., log-normal or gamma), the simulation may be implemented by repeatedly sampling from the distributions of net returns for each activity, and selecting the activity with the highest expected returns. This process is performed once for the baseline case (no payments for ecosystem services) and then for each payment level that is of interest. In the base case (no payment for ecosystem services), we would expect the land allocation to match the observed land allocation between the two practices (the point $S(p)$ in Figure 1). If more land is allocated by the model to the conservation practices than is observed, we assume that is because there are unobserved costs of adoption. We can therefore calibrate the model to the baseline by subtracting a cost term from the mean returns of the conservation practice that makes the model correctly predict the baseline land allocation. This cost term can be interpreted as an estimate of the mean adjustment cost for adopting the conservation practice.

## 4. An application to soil carbon supply

In this section we apply the MD approach to simulate the soil carbon supply curve for the dryland grain production system typical of the United States northern plains region. Farmers can receive payments per ton of carbon sequestered for switching from a crop–fallow rotation to continuous cropping or a conserving use (permanent grass cover). Our goal is to replicate that analysis using the MD approach and compare the results to those from the more detailed econometric-process simulation model.

Antle *et al.* (2003) used field-scale production survey data to estimate production models for winter wheat, spring wheat and barley crops. For each crop, a supply function, a machinery cost equation, and a cost function were estimated. These systems were used to parameterise an econometric-process simulation model to simulate farmers' participation contracts for soil carbon sequestration. The survey data were based on a representative sample of major land resource areas (MLRA) in the grain production region of Montana. The econometric models that were used to construct the econometric-process simulation model contain 124 parameters, including dummy variables for sub-MLRA effects. The analysis of soil carbon sequestration potential used the logic outlined above for the supply of ecosystem services. For computation of carbon rates, these MLRA were subdivided into low- and high-precipitation zones. The analysis was based on farmers receiving payments per metric ton of carbon sequestered.

To construct an MD simulation model for soil carbon sequestration, the survey data used by Antle *et al.* (2003) were used to estimate means and variances of net returns to each of the crops for each of the sub-MLRA zones. Logic suggests that the net returns to wheat produced under a crop–fallow rotation would be highly correlated to returns under continuous cropping, although yields under fallow are generally higher due to soil moisture stored during the fallow season. County-level data show a correlation generally

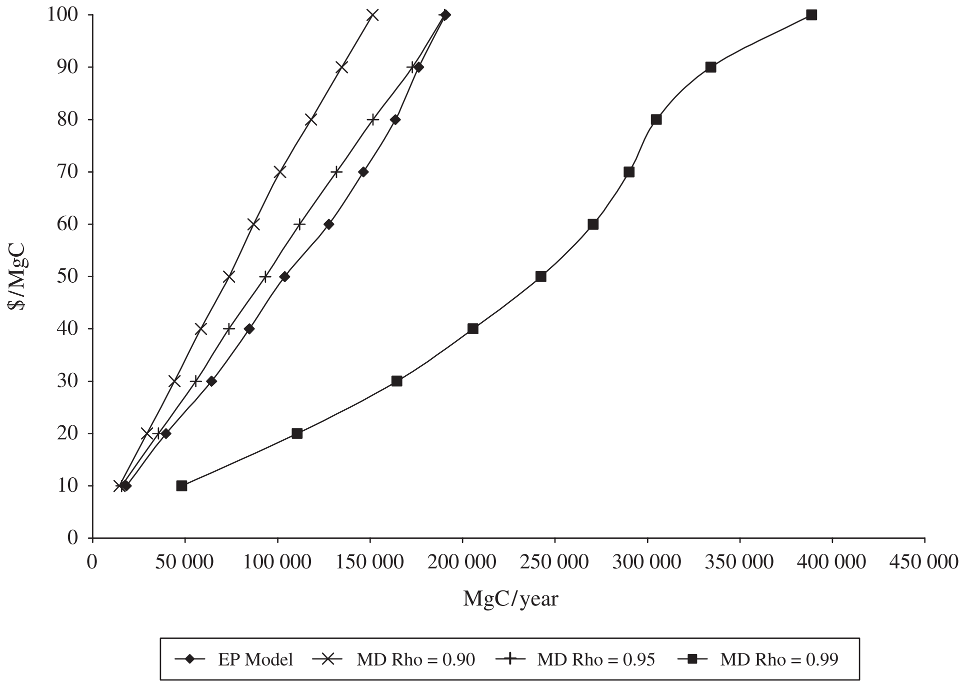


**Figure 3** Comparison of carbon supply curves from the Montana econometric-process model and from the minimum-data model with different correlations of crop returns.

between 0.8 and 0.9, with lower values in some years and values as high as 0.99 in other years. As with variances, we would expect field-level correlations to be higher than those calculated from county averages, so we conduct a sensitivity analysis for a plausible range of correlations.

Within each of the six sub-MLRA, a simulation was run by sampling net return distributions under the assumption of normality. Spring wheat is the most widely grown crop, so its variance was used, with different means for winter wheat, spring wheat and barley. The complete analysis involved four parameters (two means, a variance, and a correlation coefficient) for each of the six zones, for a total of 24 parameters. To change the system by reducing the use of fallow, there are no fixed costs of changing practices because the farmers use the same equipment. Changes in variable cost are accounted for in the estimates of the net return distributions.

Figure 3 presents the results of the MD analysis simulation aggregated across the six sub-MLRA zones, for values of the correlation coefficient equal to 0.9, 0.95 and 0.99, and the results from the original econometric-process model analysis. The MD supply curve for a 0.95 correlation is a very close approximation to the econometric-process model supply curve. Recall from Table 1, the spring wheat yield variance provided a close approximation to the variance of net returns, and is the most frequently grown crop. This fact presumably explains why the MD model is able to provide such a close

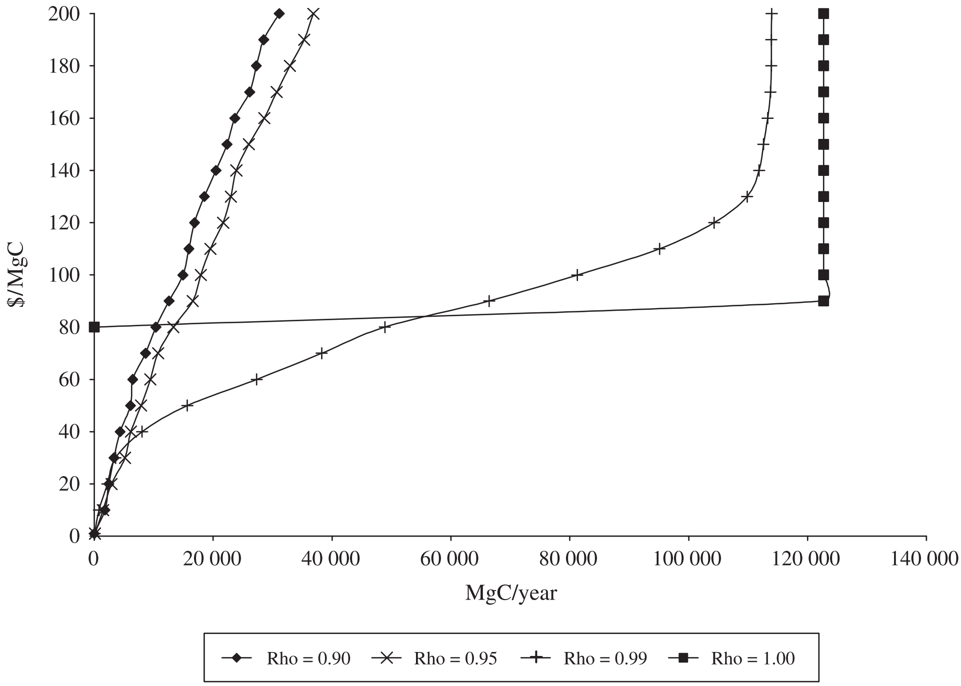


**Figure 4** Effect of changing the correlation between crop returns on the carbon supply curves in a Montana ecoregion.

approximation to the original econometric-process simulation model's supply curve. The key conclusion to be drawn from this example is that by setting the correlation coefficient within a plausible range of values, we can obtain a supply curve that is adequate for policy analysis purposes. Figure 3 shows that even though the MD supply curve changes substantially over the plausible range of correlations, nevertheless it is well within an order of magnitude, and much closer than that unless an extreme value of the correlation near unity is assumed.

Figure 4 shows results from one of the sub-MLRA zones for different correlation coefficients. This figure confirms that the supply curve follows the pattern illustrated in Figure 2 as the variance changes, approaching a step function as the correlation coefficient approaches a value of one and the variance of opportunity cost approaches a value of zero. It also shows that using a representative farm for the analysis – that is making the extreme assumption that the spatial variance is zero – produces a very poor approximation.

## 5. Conclusions

In this paper, we presented a spatially explicit production model and showed how it can be used to derive the supply of ecosystem services in a region. This model shows that the supply of ecosystem services can be derived from the spatial distribution of opportunity cost of providing those services. We then showed how this conceptual model can be used to develop an MD approach to the analysis of the supply of ecosystem services from agriculture that can be implemented with the kinds of secondary data that are available in most parts of the world. We applied the MD approach to simulate the supply of carbon that could be sequestered in agricultural soils in the dryland grain-producing region of Montana. We found that the supply curve derived from the MD approach can provide a reasonably close approximation to the supply curve obtained from a more elaborate econometric-process simulation model based on farm-level survey data, well within the order-of-magnitude range that we argue is needed for most policy analysis. Of course this does not necessarily mean that the MD approach as applied here will work in other settings. A challenge for future work is to test the MD approach with other systems, for example in an analysis of conservation tillage adoption, where fixed adoption costs are important. Our argument is that unless one can show that greater model complexity is needed, the virtues of simplicity in policy analysis are likely to outweigh greater model complexity, due to the need for timely quantitative policy analysis.

Based on our experience doing policy research, both in the United States and in developing countries, we see three principal uses of the MD approach. First, as suggested above, there is a demand for timely policy-relevant analysis to support policy decision making – even in countries such as the United States where relatively good data are available. Second, our experience in developing countries suggests that the MD approach is more appropriate than more elaborate modelling approaches because it can be learned and implemented at low cost. Our experience working with national research institutions in Latin America and Africa has shown that researchers in national research institutions often have training and data compatible with the MD approach, whereas implementing much more costly and elaborate modelling approaches is prohibitively expensive. Third, we have conducted training workshops on agricultural systems modelling in both Latin America and Africa and have found that the MD approach also can play a useful pedagogical role due to its simplicity and transparency.

## References

Antle, J.M. and Capalbo, S.M. (2001). Econometric-process models for integrated assessment of agricultural production systems, *American Journal of Agricultural Economics* 83, 389–401.

Antle, J.M., Capalbo, S.M., Mooney, S., Elliott, E.T. and Paustian, K.H. (2003). Spatial heterogeneity, contract design, and the efficiency of carbon sequestration policies for agriculture, *Journal of Environmental Economics and Management* 46, 231–250.

Dorfman, J.H., Kling C.L. and Sexton, R.J. (1990). Confidence intervals for elasticities and flexibilities: reevaluating the ratio of normals case, *American Journal of Agricultural Economics* 72, 1006–1017.

Feng, H.-L., Kling, C.L. and Gassman, P.W. (2004). Carbon sequestration, co-benefits, and conservation programs, Staff General Research Papers 12220, Department of Economics, Iowa State University. Available at URL: http://ideas.repec.org/e/pkl29.html.

Freund, J.E. (1962). *Mathematical Statistics*. Prentice Hall, Englewood Cliffs, NJ.

Just, R.E. and Antle, J.M. (1990). Interaction between agricultural and environmental policies: a conceptual framework, *American Economic Review* 80, 197–202.

Lubowski, R.N., Plantinga, A.J. and Stavins, R.N. (2005). Land-use change and carbon sinks: econometric estimation of the carbon sequestration supply function, Faculty Research Working Paper Series RP-2005-01, John F. Kennedy School of Government, Regulatory Policy Program, Harvard University, Cambridge, MA.

Pautsch, G.R., Kurkalova, L.A., Babcock, B.A. and Kling, C.L. (2001). The efficiency of sequestering carbon in agricultural soils, *Contemporary Economic Policy* 19, 123–134.

Uri, N.D. (2000). An evaluation of the economic benefits and costs of conservation tillage, *Environmental Geology* 39, 238–248.

Valdivia, R.O. (2002). The economics of terraces in the Peruvian Andes: an application of sensitivity analysis in an integrated assessment model. MS. Thesis, Department of Agricultural Economics and Economics, Montana State University, Bozeman, MT.

Wu, J., Adams, R.M., Kling, C.L. and Tanaka, K. (2004). From microlevel decisions to landscape changes: an assessment of agricultural conservation policies, *American Journal of Agricultural Economics* 86, 26–41.