Compensation mechanism, program scale, and the efficiency of voluntary carbon offset programs

Micah Cameron-Harp & Nathan P. Hendricks

mcameronharp@ksu.edu, nph@ksu.edu

Kansas State University, Department of Agricultural Economics

Author's note:

Abstract

In the present voluntary market for agricultural carbon credits, the fine scale nature of participation and evaluation can inhibit program administrators' efforts to predict producers' counterfactual behavior in the absence of a program and mitigate adverse selection. In addition, the uncertainty in estimates of soil carbon sequestration, the outcome used to issue carbon credits, increases rapidly as the scale of analysis decreases. Scaling up carbon farming programs to a jurisdiction level could address these issues by reducing uncertainty in both estimates, increasing the efficiency of achieving societal carbon sequestration goals. In such a jurisdictional approach to carbon farming programs, the program provides payments to jurisdictions based on the aggregate outcome of all producers within a jurisdiction. In this paper, we use a simulation approach to investigate the success of carbon farming programs with two different compensation mechanisms as a function of the program scale and design. When operating at the individual scale, we find programs using a per unit payment-for-sequestration approach are more efficient than those using a constant payment-per-practice mechanism only if the program administrator can estimate producers' carbon sequestration potential with sufficient accuracy. As the uncertainty in carbon sequestration rises, the payment-for-practice program becomes more efficient in terms of the average cost of an additional unit sequestered. Finally, increasing the scale of the program improves the efficiency of the payment-for-practice program to a greater degree such that it can be more efficient than the alternative, even if the error in estimating carbon sequestration is low.

1 Introduction

Adoption of conservation agriculture activities on U.S. agricultural land could sequester 234 Tg CO_2 Eq. per year, roughly 4% of the United States' annual emissions in 2020 (Sperow 2016; EPA 2020). Reductions in tillage intensity, which contribute the majority of this sequestration potential, are one of the more economical sources of sequestration costing roughly $14 to $55 / tonne CO_2 Eq. for the central United States (Antle et al. 2007). While the estimated cost is slightly lower for forest-based sequestration at between $8 and $25 / tonne CO_2 Eq. (Richards and Stavins 2005), public and private interest has increasingly focused on the carbon sequestration potential of agricultural land in recent year. For example, all four of the voluntary agriculture conservation programs receiving funds under the 2022 Inflation Reduction Act specify that projects reducing emissions or increasing sequestration of greenhouse gases must be prioritized when considering program applications (117th Congress 2022). In the private sector, there are now at least 11 voluntary agricultural carbon credit programs purchasing offsets generated by agricultural producers using conservation agriculture practices in the United States (Plastina and Wongpiyabovorn 2021).

Regardless of whether a program is public or private, voluntary programs suffer from an inherent information asymmetry between the policy maker and potential participants. Ideally, the policy maker would only enroll participants whose behavior is additional, meaning they would not have altered their emissions or sequestration in the baseline, counterfactual scenario where the policy is unavailable. In an analysis of U.S. voluntary agriculture conservation programs, Claassen, Duquette, and Smith (2018) find that just 47% of payments for conservation tillage practices produced additional adoption. Similarly, Sawadgo and Plastina (2021) conclude that 54% of the expenditures on cost-share agreements in their study of cover-crop adoption in Iowa were additional.

Programs producing offsets, like the 11 agricultural carbon credit programs analyzed in Plastina and Wongpiyabovorn (2021), combat this problem of non-additionality by comparing producers' behavior to an assigned baseline, or benchmark, based on their production history for example. However, even with extensive information on their past behavior, the baseline assigned by the policy maker is still an estimate of producers' counterfactual. The variety of approaches for setting baselines among current voluntary agricultural carbon credit programs, and the magnitude of error each entails, are the result of each entity weighing the cost of uncovering additional information on potential participants against the benefits of reducing spurious expenditures (Montero 2000).

In addition to determining how baselines will be set, administrators of agricultural carbon offset programs must also decide how they will compensate producers for the carbon sequestration or avoided emissions resulting from their production practices. One option is to compensate each producer on a per-acre basis for its use of practices above an assigned baseline based on the policy maker's prediction of counterfactual practice use in acres. Alternatively, producers could be compensated for results they produce on a per-unit of sequestration basis. In theory, a payment-for-sequestration policy is more efficient than the alternative payment-for-practice option because producers' per-unit compensation can be equated with the marginal social benefit of carbon sequestration (Sandmo 1975). So, even for difficult-to-measure ecosystem services, like biodiversity, results-based payment for ecosystem services programs are generally considered to be more cost effective than action-based programs (Hanley et al. 2012). But, while Wuepper and Huber (2021) find this to be true in an empirical study comparing action-based and results-based payments for biodiversity outcomes in Switzerland, they mention that their results are limited by the fact that they do not observe the outcome variable, changes in biodiversity, directly. This limitation highlights the principal difficulty of results-based programs, accurately measuring outcomes.

This issue of measuring outcomes is especially pertinent for agricultural carbon offset programs due to the large uncertainties in carbon sequestration outcomes. Regional uncertainties in such estimates can range between $\pm 114\%$ to 118%, but site scale uncertainties can be between $\pm 674\%$ and $\pm 739\%$ (Ogle et al., 2010). One solution to this issue, originally proposed to mitigate error in estimates of producers' baselines for offset programs, is to increase the scale of the program and take advantage of the inverse relationship between the scale of analysis and uncertainty. Using a model in which agents are offered a fixed conservation incentive akin to a payment-for-practice program, van Benthem and Kerr (2013) demonstrate that increasing the scale at which baselines are set and objectives measured improves the efficiency, offset quality, and average cost for a carbon offset program. In their model, increasing the program's scale involves aggregating the predicted counterfactual behavior across groups of agents, and then compensating the group for its performance relative to this aggregate baseline.

In this paper, we build on the approach of van Benthem and Kerr (2013) and investigate how the choice of compensation mechanism, payment-for-practice or payment-for-sequestration, affects the benefits of increasing the program's baseline scale when producers' have heterogenous carbon sequestration outcomes. We begin by determining whether an individual-level program (i) paying for practice(s) adopted on a per-acre basis or (ii) paying for the carbon offset, or sequestered, on a per-unit basis generates greater efficiency gains when the policy maker estimates producers' baseline behavior with error. Then, we introduce error in the policy maker's estimates of producers' carbon sequestration and compare the programs' performance. Due to the significance both types of error have for determining the relative performance of the compensation mechanisms, we test the impact of program scale for various combinations of carbon sequestration and baseline estimation error.

The two most similar works concerning agricultural carbon sequestration are the empirical work of Antle et al. (2003) and theoretical investigation by van Benthem and Kerr (2013). Antle et al. (2013) include heterogeneous outcomes and producers' costs of adoption in their model using site-specific estimates of carbon sequestration and producer characteristics. They find a per-acre payment-for-practice program could be as much as five times more costly than a per-ton payment-for-sequestration program depending on the acceptable level of uncertainty in carbon sequestration estimates. Unlike in our model, the two programs in Antle et al. (2013) do not measure performance relative to a counterfactual scenario. So, there is no role for error in the policy maker's estimates of producers' behavior in the absence of the program. In addition, the carbon sequestration estimates in Antle et al. (2013) are formulated at a regional scale based on the USDA's Major Land Resource Areas and the effect of program scale on these estimates is not investigated. As mentioned earlier, the work of van Benthem and Kerr (2013) focuses primarily on the effect of program scale on the efficiency of a carbon offset program that compares producers' behavior to a predicted counterfactual scenario. As the carbon sequestration potential of all agents in van Benthem and Kerr's (2013) model is assumed to be identical, they do not consider how the choice of compensation mechanism or uncertainty in carbon sequestration outcomes modifies the effect of increasing program scale.

Our work addresses the limitations of these two previous works by uniting the heterogenous sequestration outcomes and competing programs designs of Antle et al. (2003) with the baseline prediction and program scale features of van Benthem and Kerr (2013). In doing so, we make three contributions to our understanding of carbon offset policies. First, if the policy maker knows the sequestration potential of producers with certainty, we find the payment-for-sequestration program is more efficient regardless of the policy maker's error in estimating baselines because the per-unit compensation mechanism selects for individuals with greater carbon sequestration potential. Second, when we introduce uncertainty in the policy maker's estimates of sequestration outcomes, the payment-for-practice program is more efficient under certain conditions. For example, when the policy maker's error in estimating sequestration is high relative to its error estimating producer's baseline adoption behavior, the average cost of additional sequestration is lower for the payment-for-practice program. Finally, increasing the program scale reduces the average cost of additional sequestration for both compensation mechanisms, but the current level of adoption in the population determines which program benefits to a greater degree. In areas with low levels of adoption, scaling up the program benefits the payment-for-practice program to a greater extent, and the opposite holds true for the payment-for-sequestration program.

2 Model of an individual-scale voluntary carbon credit program

The initial formulation of our model modifies the jurisdiction-level model from van Benthem and Kerr (2013) for a hypothetical agricultural context where adoption of a practice produces a positive marginal carbon externality by increasing carbon sequestration. To demonstrate how the model functions, envision an example case where the agents in our model are producers considering switching from conventional tillage to no-till for the next growing season. Next growing season, even in the absence of any program, some of these hypothetical producers may be planning to adopt no-till next growing season purely because of its agronomic benefits. We refer to this counterfactual scenario as the baseline scenario in our model, and we construct this baseline scenario for a population of N producers with identically sized parcels such that parcel and producer are synonymous. The population of producers is divided evenly amongst a total of N_j jurisdictions, so the number of producers in each jurisdiction is $N_{i,j} = N/N_j$. In our no-till example, these jurisdictions could be the countries that producers are located within.

The number of producers using the practice in the baseline scenario is determined by each individual producer's net return to adopting the practice, a random variable defined as $r_{i,j}$ in our model. The subscript i indexes individual producers within jurisdictions indicated by the subscript j. We assume producers are profit maximizing, so the baseline behavior for producers with net returns greater than zero is to use the practice, indicated by the binary variable $BL_{i,j}$ equaling one:

$$BL_{i,j} = \begin{cases} 1 & if\ r_{i,j} > 0 \\ 0 & if\ r_{i,j} \leq 0 \end{cases}. \quad \text{(Eq. 1)}$$

As such, the number of producers in jurisdiction j adopting the practice in this baseline scenario is: $BL_j = \sum_{i=1}^{N_{i,j}} BL_{i,j}$. Defining the distribution of the net return random variable as f_r, then the expected level of baseline adoption for any given jurisdiction is:

$$E[BL_j] = E\left[\sum_{i=1}^{N_{i,j}} BL_{i,j}\right] = N_{i,j} * E[BL_{i,j}] = N_{i,j} * \int_0^\infty f_r\,(r)dr.$$

(Eq. 2)

The other distinguishing characteristic of agents in our model is another random variable which indicates their individual carbon sequestration potential, defined as $c_{i,j}$. More specifically, $c_{i,j}$ represents the marginal change in carbon sequestration due to producer i in jurisdiction j adopting the conservation agriculture practice. We allow $c_{i,j}$ to vary between producers to mimic the variation in carbon sequestration that occurs when producers with different soil types and production histories adopt conservation agriculture practices. For example, switching from conventional tillage to no-till is estimated to sequester twice as much carbon in the Corn Belt in comparison to states in the arid southwest (Swan et al. 2020). Given this definition of $c_{i,j}$, we define the carbon sequestration of producer i in jurisdiction j that occurs because of its adoption behavior in the baseline scenario as:

$$BLC_{i,j} = \begin{cases} c_{i,j} & if\ r_{i,j} > 0 \\ 0 & if\ r_{i,j} \leq 0 \end{cases}. \quad \text{(Eq. 3)}$$

Consequently, the total sequestration that occurs because of producers' baseline adoption behavior across all producers in jurisdiction j can be expressed as: $BLC_j = \sum_{i=1}^{N_{i,j}} BL_{i,j} * c_{i,j}$. Letting f_c be the distribution governing the random variable $c_{i,j}$, and assuming independence between f_c and f_r, we can express the expected value of $BLC_{i,j}$ for any given producer i in a jurisdiction j as: $E\left[BLC_{i,j}\right] = \int_{-\infty}^{\infty} c_{i,j} * f_c(c_{i,j}) * \int_0^{\infty} f_r\,(r)dr$. As such, the total expected baseline sequestration for jurisdiction j is therefore:

$$E\left[BLC_j\right] = N_{i,j} * \int_{-\infty}^{\infty} c_{i,j} * f_c(c_{i,j})dc * \int_0^{\infty} f_r\,(r)dr. \quad \text{(Eq. 4)}$$

In the next section, we describe how the expected jurisdictional baseline adoption, $E\left[BL_j\right]$, and carbon sequestration outcomes, $E\left[BLC_j\right]$, in Equations 2 and 4 are approximated by the policy maker. Then, we describe how these approximations are used to determine jurisdictions' eligibility and outcomes for two carbon offset programs of differing design.

3 Policy Design

The objective of the policy maker in our model is to offset carbon emissions by providing incentives for the adoption of conservation agriculture practices at the jurisdiction level. Both voluntary offset program designs we will consider involve the policy maker predicting producers' returns to adopting the conservation agriculture practice and, therefore, their baseline behavior. The policy maker uses these predictions to assign individual baselines which are then aggregated across all producers in a jurisdiction to create jurisdictional assigned baselines. In doing so, the policy maker aims to reduce participation by non-additional producers. To demonstrate how this approach addresses non-additionality, consider the simplified case when jurisdictions contain a single producer, meaning $N_{i,j} = 1$. For either program design, the policy maker aims to only enroll and compensate producers who would not have adopted the practice in the baseline scenario, those with $BL_j = BL_{i,j} = 1$ given $N_{i,j} = 1$. So, the policy maker generates predictions of producers' counterfactual behavior and assigns baselines written as $\widehat{BL}_{i,j}$ and $\widehat{BL}_j$ for individuals and jurisdictions respectively. Then, the policy maker only offers compensation to producers with assigned baselines indicating they would not have adopted the practice in the counterfactual scenario, those with $\widehat{BL}_j = \widehat{BL}_{i,j} = 0$. If the policy maker can perfectly predict producers baseline behavior, then it would only offer compensation to individuals would not adopt have adopted the practice without an incentive and avoid wasted expenditures on non-additional adoption. While the compensation mechanisms defined later for the general case when $N_{i,j} \geq 1$ are slightly more complicated, the general approach of using predicted behavior to minimize non-additionality is the same for both offset programs.

However, because the policy maker cannot observe the counterfactual scenario where the voluntary programs do not exist, its predictions and the resulting assigned baselines are made with some error. We assume the policy maker generates the assigned baselines using estimates of individuals' net returns, defined as $\hat{r}_{i,j}$, which include an individual specific error term, $\varepsilon^r_{i,j}$, such that: $\hat{r}_{i,j} = r_{i,j} + \varepsilon^r_{i,j}$. The baseline assigned by the policy maker for producer i in jurisdiction j is therefore:

$$\widehat{BL}_{i,j} = \begin{cases} 1 \ \ if\ \hat{r}_{i,j} > 0 \\ 0 \ \ if\ \hat{r}_{i,j} \leq 0 \end{cases}. \qquad \text{(Eq. 5)}$$

The error $\varepsilon_{i,j}^{r}$ could originate because the policy maker, for instance, uses the average change in no-till adoption for jurisdiction j over the past few years to predict the likelihood of individual producers using no-till next year in jurisdiction j. Even if this prediction is correct on average, the policy maker will likely make mistakes for individual producers in the jurisdiction because of their unique production characteristics. In a similar fashion, we assume the error in the policy maker's estimate of net returns, $\varepsilon_{i,r}$, is symmetric around zero, so the policy maker's estimates for individuals' returns to adoption are unbiased. We also assume the error in predicting a producer's net returns, $\varepsilon_{i,r}$, is independent of the producer's true return to adoption drawn from f_r, and that it has standard deviation σ_{ε}^{r}. Letting f_{ε}^{r} be the distribution governing the random variable $\varepsilon_{i,r}$, then we can express the expected value for a producer's assigned baseline as: $E\left[\widehat{BL}_{i,j}\right] = \int_0^{\infty} f_r(r) f_{\varepsilon}^{r}(\varepsilon) dr$. Like for the true baselines, we assume the policy maker adds up the individual assigned baselines for all individuals in each jurisdiction to construct the assigned baseline for an entire jurisdiction: $\widehat{BL}_j = \sum_{i=1}^{N_{i,j}} \widehat{BL}_{i,j}$. The expected value of this jurisdictional assigned baseline, $\widehat{BL}_j$, is simply the product of the expected value for the individual baseline and the number of producers in each jurisdiction:

$$E\left[\widehat{BL}_j\right] = N_{i,j} \int_0^{\infty} f_r(r) f_{\varepsilon}^{r}(\varepsilon) dr. \qquad \text{(Eq. 6)}$$

Notice, that the expected value for true jurisdictional baseline in Equation 2 differs from the expected value of the assigned jurisdictional baseline in Equation 6. As shown in van Benthem and Kerr, even though the policy maker's estimates of producers' returns are unbiased, it's estimates of baselines will be biased unless f_r is symmetric around zero.

Proposition 1: If $E[\varepsilon\,] = 0 \,\&\, \sigma_{\varepsilon}^{r} > 0$, then $E\left[\widehat{BL}_j\right] > E\left[BL_j\right]$ when $E[f_r] < 0$ and $E\left[\widehat{BL}_j\right] > E\left[BL_j\right]$ when $E[f_r] > 0$.

Proposition 1 states that the policy makers assigned jurisdictional baselines will overestimate baseline adoption of the practice if adopting the practice is unprofitable on average among the population of producers, and they will underestimate baseline adoption when the opposite is the case. Despite the bias present in these assigned jurisdictional baselines, we do not modify them because doing so would imply that the policy maker has knowledge of either the true distribution of returns, its prediction error, or both.

Given its objective of encouraging addition carbon sequestration, the policy maker also estimates the carbon sequestration associated with its assigned baselines concerning producers' adoption behavior. These assigned carbon sequestration baselines are generated in a very similar manner based on the policy maker's estimates of producers' carbon sequestration potential. Like the estimates of net returns, the policy maker estimates producers carbon sequestration with some error, $\varepsilon_{i,j}^{c}$, such that: $\hat{c}_{i,j} = c_{i,j} + \varepsilon_{i,j}^{c}$. As such, we define the assigned carbon sequestration baseline for a representative producer as:

$$\widehat{BLC}_{i,j} = \begin{cases} \hat{c}_i \ if \ \hat{r}_i > 0 \\ 0 \ \ if \ \hat{r}_i \leq 0 \end{cases} . \qquad \text{(Eq. 7)}$$

We make similar assumptions concerning $\varepsilon_{i,j}^{c}$ as for the net return error random variable. Namely, we assume the error in the policy maker's estimate of each producer's carbon sequestration potential, $\varepsilon_{i,c}$, is symmetric around zero, independent of f_r and f_ε^r, and has standard deviation σ_c. As $\widehat{BLC}_{i,j} = \widehat{BL}_{i,j} * \hat{c}_{i,j}$, and given the independence between each random variable, the expected value of $\widehat{BLC}_{i,j}$ is as follows:

$$E\left[\widehat{BLC}_{i,j}\right] = \int_0^\infty f_r(r) f_\varepsilon^r(\varepsilon) dr * \int_{-\infty}^\infty c_{i,j} * f_c(c) dc. \qquad \text{(Eq. 8)}$$

Next, we introduce a payment-for-practice program, in which the policy maker offers a payment equal to p_p to producers who it predicts would not adopt the practice in the absence of the program. In the no-till example, this program would pay producers p_p for adopting no-till in the coming growing season only if the policy maker predicts they were going to use conventional tillage in the absence of a program. The policy maker in our model predicts producers' baseline use of the conservation agriculture practice based on its estimates of their net returns, defined as $\hat{r}_i$ for producer i.

We can now illustrate two inefficiency categories created by the policy maker's net return observation error, inefficient participants and non-participants, by examining a representative producer's adoption decision when this payment-for-practice program is available. Given its assigned baseline, a producer will enroll in the program if the net return to participating is positive, and its assigned baseline indicates they would not have used the practice in the absence of the program: $r_i + p_p > 0$ and $\widehat{BL}_i = 0$. However, the producer will only be efficiently enrolled in the program if its true baseline behavior was to not adopt the practice, meaning $r_i + p_p > 0$, $\widehat{BL}_i = 0$, and $BL_i = 0$.

Inefficient participants are the non-additional producers who are enrolled even though they would have adopted the practice if the program was unavailable, $r_i + p_p > 0$, $\widehat{BL}_i = 0$, and $BL_i = 1$. For the example concerning the adoption of no-till, these would be producers who were planning to adopt no-till next growing season, but the policy maker mistakenly predicts that they were going to use conventional tillage. As a result, the policy maker lets them participate and they are paid for adopting a practice they had already planned to use. In our model, these non-additional, inefficient participants can enroll because the policy maker underestimates their net return to using the practice.

Conversely, the inefficient nonparticipants occur because the policy maker overestimates some producers' net return to adoption, so the resulting assigned baselines prevent producers from efficiently participating: $r_i + p_p > 0$, $\widehat{BL}_i = 1$, and $BL_i = 0$. These would be producers who were going to use conventional tillage next season, but the policy maker predicts that they're going to use no-till in the example case. As a result, these producers are not offered the payment p_p, even though it would have caused them to switch from conventional tillage to no-till next season.

4 Heterogeneity in the marginal carbon externality from adoption

4.1 Payment-for-practice program

As described earlier, the payment-for-practice program offers the fixed level of compensation p_p to producers who adopt the practice when their assigned baseline indicates they would not have otherwise. Let the binary variable E_i^p indicate the enrollment status of producer i, having a value of one if producer i enrolls, according to:

$$E_i^p = \begin{cases} 1 \; if \; r_i + p_p > 0 \; and \; \widehat{BL}_i = 0 \\ 0 \; if \; r_i + p_p \leq 0 \; or \; \widehat{BL}_i = 1 \end{cases}. \qquad \text{(Eq. 5)}$$

In Equation 5, the superscript p is used to indicate variables associated with the payment-for-practice program. Notice that E_i^p can only take a value of 1, indicating the producer is participating, if the policy maker predicts that the producer would not use the practice in the baseline scenario ($\widehat{BL}_i = 0$). A producer may not be participating, $E_i^p = 0$ in Equation 5, for one of two reasons. First, it could be because the incentive p_p is too small to make adoption profitable, meaning $p_p < |r_i|$. Or a producer may not be participating because the policy maker estimates that it would adopt the practice in the baseline scenario $\widehat{BL}_i = 1$, so it is not eligible for the payment.

For simplicity, we assume perfect compliance by producers meaning all producers who accept the payment and participate will use the practice without fail.

Given the variation in producers carbon sequestration outcomes, we also define the sequestration that occurs based on producers' behavior when the payment-for-practice program is available as:

$$c_i^p = \begin{cases} c_i \ \ if \ r_i + E_i^p * p_p > 0 \\ \ \ 0 \ \ if \ r_i + p_p \leq 0 \end{cases}. \qquad \text{(Eq. 6)}$$

In Equation 6, there are three cases in which a producer sequesters carbon when the payment-for-practice program is available, such that $c_i^p = c_i$. First, there are the efficient enrollees, as defined above, who would not have adopted the practice in the baseline so $BLC_i = 0$. Second, there are inefficient participants who would have adopted the practice in the baseline scenario, meaning $BLC_i = c_i$. Last, notice in Equation 6 that a producer may sequester c_i even if they are not enrolled, meaning $E_i^p = 0$. This type of producer has positive net returns to adopting the practice, so they adopt the practice even though their correctly assigned baseline of $\widehat{BL}_i = 1$ prevents them from receiving the incentive p_p. When calculating the additional sequestration that the payment-for-practice program produces, the second and third cases where $c_i^p = c_i$ necessitate subtracting producers' true baseline carbon sequestration values, BLC_i, from the sequestration that occurs when the program is available, c_i^p. We define the additional sequestration that occurs because of the payment-for-practice program as AS^P, and it is calculated using the following equation:

$$AS^P = \sum_i^N c_i^p - BLC_i. \qquad \text{(Eq. 7)}$$

For our no-till example, Equation 7 is the mathematical equivalent of only counting the sequestration that occurs because of individuals who made the switch from conventional tillage to no-till because of the incentive provided to them.

The policy maker, however, cannot calculate the true amount of carbon sequestration occurring due to the policy using Equation 7, because it does not know producers' true returns to adopting the practice and, consequently, their true baseline sequestration values, BLC_i. Instead, the policy maker estimates the amount of additional sequestration generated by its payment-for-practice program using its estimates of producers' baseline sequestration according to:

$$\widehat{AS}^P = \sum_i^N c_i^p - \widehat{BLC}_i. \quad \text{(Eq. 8)}$$

Notice that even though the policy maker knows the carbon sequestration potential of producers with certainty, the policy maker's estimate of carbon sequestration resulting from the payment-for-practice program in Equation 8 may be significantly different than the true additional sequestration in Equation 7 depending on how poorly the policy maker estimates producers net returns.

If the policy maker's estimate differs from the true additional sequestration caused by the program, meaning $AS^P \neq \widehat{AS}^P$, this difference is due to inefficient participants who receive an incorrect baseline of zero carbon sequestration from the policy maker. The fraction of carbon sequestration estimated by the policy maker that is based on the behavior of these inefficient participants (IP^p) is expressed as:

$$IP^p = 1 - \frac{\sum_i^N (c_i^p - BLC_i)}{\widehat{AS}^P}. \quad \text{(Eq. 9)}$$

We will also refer to the quantity of estimated sequestration contributed by inefficient participants (IP^p) as the portion of offsets from inefficient enrollees. This definition reflects the policy maker's objective of incentivizing carbon sequestration, or purchasing carbon offsets, such that Equation 9 represents the portion of offsets purchased by the policy maker which are non-additional. For the example involving the adoption of no-till, Equation 9 would indicate the fraction of the carbon sequestration that the policy maker thinks they have incentivized but is truly from individuals who were planning to adopt no-till regardless. As such, it indicates the portion of the outcome estimated by the policy maker which would have occurred in the absence of any payments.

The other source of inefficiency described earlier, inefficient nonparticipants, occurs because the policy maker overestimates some producers' net returns and assigns them a positive baseline carbon sequestration value when the true value is zero. These would be individuals who were not planning on adopting no-till next season in our example, but the policy maker does not offer them an incentive to adopt because it incorrectly predicts that they were going to use no-till in the baseline scenario. We define the efficiency loss (EL^p) as the ratio of the carbon that is not sequestered because inefficient nonparticipants receive incorrect assigned baselines to the total carbon that would be efficiently sequestered if the policy maker observed producers' net returns without error:

$$EL^p = \frac{\sum_i^N (c_i - BLC_i \mid r_i + p_p > 0 \ \& \ \widehat{BL}_i = 1)}{\sum_i^N (c_i - BLC_i \mid r_i + p_p > 0 \ \& \ BL_i = 0)}. \qquad \text{(Eq. 10)}$$

The final metric we construct to measure the program's success is the average cost per unit of additional sequestration (AC^p), which is calculated by dividing the total transfers from the policy maker to producers by the additional sequestration the program results in:

$$AC^p = \frac{p_p * \sum_i^N E_i^p}{AS^P}. \qquad \text{(Eq. 11)}$$

4.2 Payment-for-sequestration program

The second distinguishing feature of our model is an alternative payment-for-sequestration program which offers to compensate each producer p_c per unit of carbon it sequesters above its assigned baseline carbon sequestration, $\widehat{BLC}_i$. The programs are offered in two separate states of the word, so that either the payment-for-practice or payment-for-sequestration program is available in any given scenario. As with the payment-for-practice program, we assume perfect compliance, so any producer enrolling in the payment-for-sequestration program uses the conservation agriculture practice. Additionally, we continue to assume the policy maker knows the carbon sequestration of individual producers with certainty. Given these assumptions and the per unit of sequestration incentive p_c, the binary variable E_i^c is defined as follows to indicate whether producer i enrolls in the payment-for-sequestration program:

$$E_i^c = \begin{cases} 1 \ if \ r_i + p_c * \left(c_i - \widehat{BLC}_i\right) > 0 \ and \ \widehat{BL}_i = 0 \\ 0 \ if \ r_i + p_c * \left(c_i - \widehat{BLC}_i\right) \leq 0 \ or \ \widehat{BL}_i = 1 \end{cases}. \qquad \text{(Eq. 12)}$$

Note, in equation 12, that the payment-for-sequestration program shares the enrollment criteria of the payment-for-practice program requiring producers to have an assigned baseline of not using the practice ($\widehat{BL}_i = 0$). Unlike the payment-for-practice program, the compensation an individual producer receives varies as a function of its carbon sequestration potential, c_i. We define the carbon sequestration for producers when the payment-for-sequestration program is available as:

$$c_i^c = \begin{cases} c_i \ \ if \ r_i + p_c * (c_i - \widehat{BLC}_i) > 0 \\ 0 \ \ if \ r_i + p_c * (c_i - \widehat{BLC}_i) \leq 0 \end{cases}. \qquad \text{(Eq. 13)}$$

Given these outcomes, we can define the same four metrics used to evaluate the payment-for-practice program. The additional carbon sequestration which occurs (AS^C) due to the payment-for-sequestration program is:

$$AS^C = \sum_i^N c_i^c - BLC_i. \qquad \text{(Eq. 14)}$$

But the sequestration estimated by the policy maker for the payment-for-sequestration program is:

$$\widehat{AS}^c = \sum_i^N c_i^c - \widehat{BLC}_i. \qquad \text{(Eq. 15)}$$

The fraction of sequestration estimated by the policy maker which is due to adoption by inefficient participants (IP^c) and the efficiency loss which occurs due to the incorrectly assigned baselines for inefficient nonparticipants (EL^c) are:

$$IP^c = 1 - \frac{\sum_i^N (c_i^c - BLC_i)}{\widehat{AS}^c} \qquad \text{(Eq. 16)}$$

and

$$EL^c = \frac{\sum_i^N (c_i - BLC_i \mid r_i + p_c * (c_i - BLC_i) > 0 \;\&\; E_i^c = 0)}{\sum_i^N (c_i - BLC_i \mid r_i + p_c * (c_i - BLC_i) > 0 \;\&\; BL_i = 0)}\,. \qquad \text{(Eq. 17)}$$

Finally, the last metric used to evaluate the success of the program is the average cost of an additional unit of sequestration:

$$AC^c = \frac{p_c * \sum_i^N E_i^c * (c_i - \widehat{BLC}_i)}{AS^c}\,. \qquad \text{(Eq. 18)}$$

4.3 Simulation of program performance as a function of net return observation error

To test the performance of the two program types when the policy maker knows producers' carbon sequestration with certainty, we create a simulation of our conceptual model and examine how the four metrics defined for each compensation mechanism vary as a function of the policy maker's error in observing producers' net returns, $\varepsilon_{i,r}$. The simulation is parameterized by assigning values for the two incentives, $p_p = 0.5$ and $p_c = 1$, and making three distributional assumptions for r_i, c_i, and $\varepsilon_{i,r}$.

First, we assume producers' net returns to adopting the practice in the absence of any program are distributed normally such that: $f_r(r_i) \sim N(-0.5, 1^2)$. On average, this means that roughly 31% of producers would adopt the practice without any incentive. Second, we assume producer's carbon sequestration due to practice adoption, c_i, is drawn from a standard uniform distribution, $f_f(c_i) \sim U[0,1]$. Given the incentives, $p_p = 0.5$ and $p_c = 1$, the expected compensation is equal between the two program types as: $E[f_f(c_i)] = 0.5$. Last, the policy maker's error in estimating producers' net returns, $\varepsilon_{i,r}$, is assumed to be a normally distributed random variable with mean zero and standard deviation, σ_r.

For both payment mechanisms, the simulation is run for 31 values of σ_r running from zero to 1.5 in increments of 0.05, with 800 replicates for each value of σ_r. In every model run, there are a total of 10,000 producers, or plots, simulated. The average efficiency loss (EL), average cost of a unit of additional sequestration (AC), percent of offsets from inefficient enrollees (IP), and additional sequestration (AS) per plot across all model runs for the two program types are displayed in Figure 1.

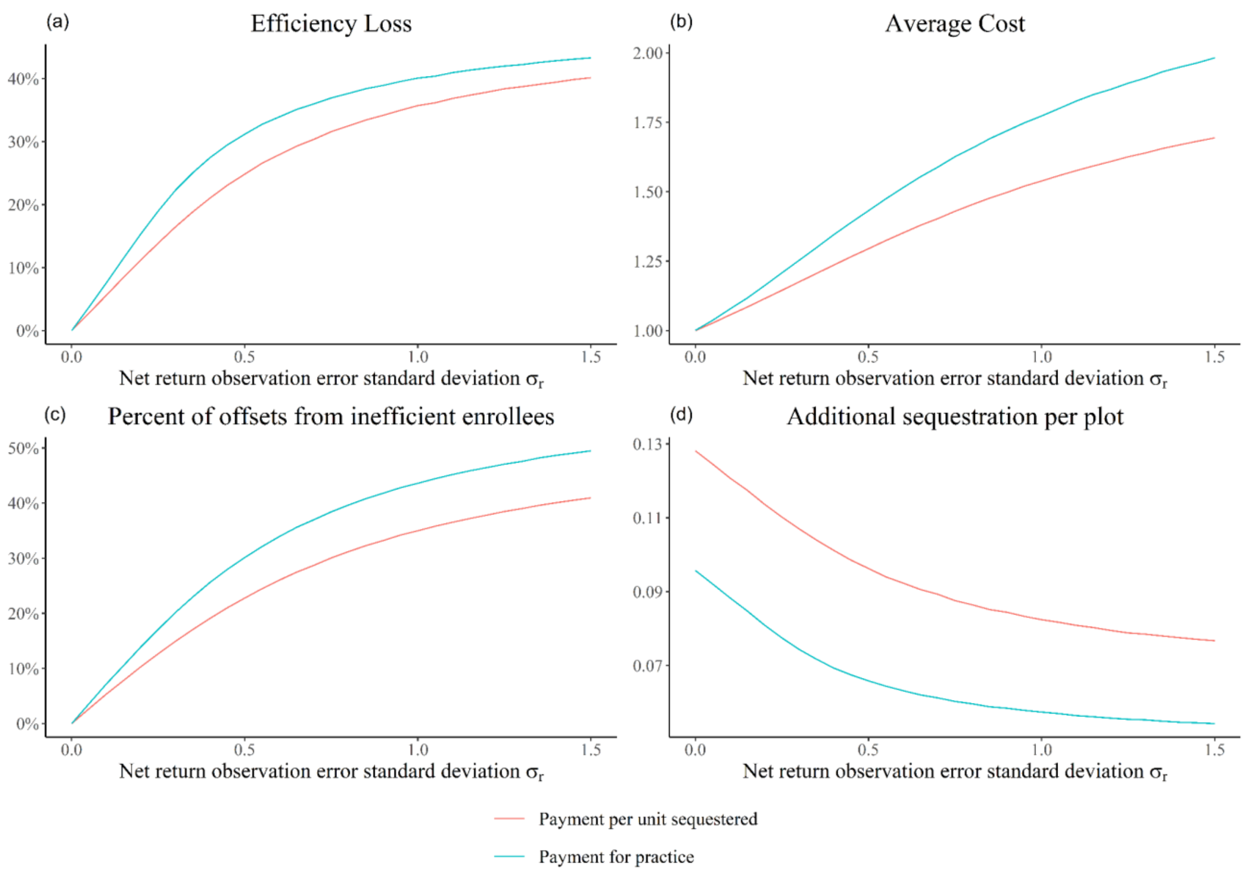

Figure 1. Effect of error in policy maker's estimates of producers' net returns to practice adoption. The x-axis indicates the standard deviation of the policy maker's error.

In Figure 1, the payment-for-sequestration program outperforms the payment-for-practice program in every metric for all values of $\sigma_r > 0$ when the policy maker knows producers' carbon sequestration with certainty. The difference between the two compensation mechanisms' performance is strictly increasing in the standard deviation of the net return observation error for two of the metrics examined, the average cost of an additional unit of sequestration in panel (b) and the percent of estimated sequestration from inefficient enrollees in panel (c) of Figure 1. The efficiency loss in panel (a) of Figure 1 initially increases more slowly for the payment-for-sequestration program as σ_r grows, but it then accelerates at intermediate levels of σ_r such that it approaches the efficiency loss of the payment-for-practice program at extreme values of σ_r. In contrast to the behavior of the other metrics, the difference in additional sequestration per plot between the two program types, displayed in panel (d) of Figure 1, remains constant as σ_r increases.

The reason for the improved performance of the payment-for-sequestration program and the behavior of the four metrics just described, is the built-in selection criteria of results-based compensation mechanisms. To demonstrate how this selection occurs, consider the potential carbon sequestration values for an example producer x. Letting c_x indicate the carbon producer x would sequester if it adopts the practice, then c_x can take any value on the interval $c_x \in [0,1]$ given the assumed standard uniform distribution for c_i. Now, assume producer x would efficiently enroll in the payment-for-practice program if given the opportunity. Producer x's carbon sequestration, c_x could still lie anywhere on the closed interval between zero and one, meaning $c_x \in [0,1]$, because eligibility and profitability are solely determined by producer x's returns to adoption, r_x, for the payment-for-practice program. However, if we assume producer x would efficiently enroll in the payment-for-sequestration program, the potential values of c_x are constrained to $c_x \in$

$(|r_x|/p_c, 1]$ due to the producer's profit maximizing behavior. For producer x will only efficiently enroll if $r_x + p_c * c_x > 0$. So, when the policy maker knows producers' carbon sequestration with certainty, the payment-for-sequestration program effectively filters participants such that the average carbon sequestered by a participant will always be greater than that of the payment-for-practice program.

5 Including observation error for carbon sequestration

5.1 Payment-for-practice program with carbon observation error

Producers' enrollment decision is identical for the payment-for-practice program when the policy maker observes their carbon sequestration with error because compensation is based on producers' use of the practice relative to their assigned practice use baselines. This means equations 5, 6, and 7 produce the enrollment status (E_i^{p,σ_c}), true carbon sequestration outcomes for producers (c_i^{p,σ_c}), and total additional sequestration (AS^{p,σ_c}) when the payment-for-practice program is available despite the policy maker's error in estimating carbon sequestration. The efficiency loss (EL^{p,σ_c}) and average cost of an additional unit sequestered (AC^{p,σ_c}) are calculated using equations 10 and 11 for the same reason.

In contrast, the outcomes reported by the policy maker for the payment-for-practice program change significantly. First, we define $\hat{c}_i^{p,\sigma_c}$ as the carbon sequestration that is predicted due to producers' use of the conservation agriculture practice when the payment-for-practice program is available:

$$\hat{c}_i^{p,\sigma_c} = \begin{cases} \hat{c}_i \ \ if \ r_i + E_i^p * p_p > 0 \\ 0 \ \ if \ r_i + E_i^p * p_p \leq 0 \end{cases}. \qquad \text{(Eq. 20)}$$

As the policy maker no longer knows producers' sequestration potential with certainty, the policy maker's estimate for the carbon sequestered due to the program is now a function of the predicted carbon sequestration outcomes $\hat{c}_i^{p,\sigma_c}$ for enrolled producers:

$$\widehat{AS}^{p,\sigma_c} = \sum_i^N \hat{c}_i^{p,\sigma_c} * E_i^{p,\sigma_c}. \qquad \text{(Eq. 21)}$$

Consequently, the fraction of sequestration expected by the policy maker based on the sequestration of inefficient participants (IP^p, σ_c) becomes:

$$IP^{p,\sigma_c} = \frac{\sum_i^N (\hat{c}_i^{p,\sigma_c} * E_i^{p,\sigma_c} \mid BL_i = 1)}{\widehat{AS}^{p,\sigma_c}}. \qquad \text{(Eq. 22)}$$

5.2 Payment-for-sequestration with carbon observation error

For the payment-for-sequestration program, there are changes to both the producers' enrollment decisions and the outcomes reported by the program administrator. Each producer's enrollment decision is altered because its potential compensation is calculated using the policy maker's estimates of carbon sequestration:

$$E_i^{c,\sigma_c} = \begin{cases} (1 - \widehat{BL}_i) \; if \; r_i + p_c * (\hat{c}_i - \widehat{BLC}_i^{\sigma_c}) > 0 \; and \; (\hat{c}_i - \widehat{BLC}_i^{\sigma_c}) > 0 \\ 0 \; if \; r_i + p_c * (\hat{c}_i - \widehat{BLC}_i^{\sigma_c}) \leq 0 \; or \; (\hat{c}_i - \widehat{BLC}_i^{\sigma_c}) \leq 0 \end{cases}. \qquad \text{(Eq. 23)}$$

In addition to using the estimated carbon sequestration potential for producers, the enrollment decision in equation 23 differs from the case without carbon sequestration error due to the restrictions on the predicted carbon sequestration above the assigned baseline, $(\hat{c}_i - \widehat{BLC}_i^{\sigma_c})$. These restrictions account for the possibility that the estimated carbon sequestration associated with using a practice is negative due to the policy maker's error in estimating a producer's carbon sequestration potential. If an eligible producer's return to adopting a practice is positive but the predicted carbon sequestration is negative, we assume the producer would use the conservation agriculture practice but would not enroll in the program as: $r_i > r_i + p_c * (\hat{c}_i - \widehat{BLC}_i^{\sigma_c})$. Using equation 23, we construct the true and estimated carbon sequestration by producers when the payment-for-sequestration program is available:

$$c_i^{c,\sigma_c} = \begin{cases} c_i \ \ if \ r_i + p_c * \hat{c}_i * E_i^{c,\sigma_c} > 0 \\ 0 \ \ if \ r_i + p_c * \hat{c}_i * E_i^{c,\sigma_c} \leq 0 \end{cases}; \quad \text{(Eq. 24)}$$

$$\hat{c}_i^{c,\sigma_c} = \begin{cases} \hat{c}_i \ \ if \ r_i + p_c * \hat{c}_i * E_i^{c,\sigma_c} > 0 \\ 0 \ \ if \ r_i + p_c * \hat{c}_i * E_i^{c,\sigma_c} \leq 0 \end{cases}. \quad \text{(Eq. 25)}$$

The true additional sequestration which occurs (AS^{c,σ_c}) and the sequestration estimated by the policy maker ($\widehat{AS}^{c,\sigma_c}$) for the payment-for-sequestration program are therefore:

$$AS^{c,\sigma_c} = \sum_i^N c_i^{c,\sigma_c} - BLC_i \, ; \quad \text{(Eq. 26)}$$

$$\widehat{AS}^{c,\sigma_c} = \sum_i^N \hat{c}_i^{c,\sigma_c} * E_i^c \, . \quad \text{(Eq. 27)}$$

Constructing the average cost per unit of additional sequestration for the payment-for-sequestration program with error in estimating carbon sequestration involves minor changes to equation 18:

$$AC^{c,\sigma_c} = \frac{p_c * \sum_i^N E_i^{c,\sigma_c} * (\hat{c}_i^{c,\sigma_c} - \widehat{BLC}_i)}{AS^{c,\sigma_c}} \, . \quad \text{(Eq. 28)}$$

To define the fraction of offsets due to inefficient participation (IP^{c,σ_c}) and the efficiency loss (EL^{c,σ_c}) in this context, we must first consider how including carbon sequestration error affects each of the two categories of inefficient producers. It is still possible to have non-additional enrollment by producers if the policy maker incorrectly assigns a baseline of not using the practice when the opposite is true.

But unlike when the policy maker knows the carbon sequestration potential of producers with certainty, there is also another group of inefficient participants who enroll and alter their behavior relative to their baselines because the policy maker overestimates the carbon they will sequester. These participants are inefficient because the profit-maximizing choice for them, in the absence of error estimating carbon sequestration, would be to refrain from using the practice. Despite this, the sequestration occurring due to this second type of inefficient participants is included in AS^{c,σ_c} because it truly is additional relative to these producers' baselines. With error in carbon sequestration estimates, the fraction of the sequestration estimated by the policy maker due to inefficient participants combines the non-additional enrollees as well as the producers inefficiently participating due to overestimated values for their carbon sequestration potential:

$$IP^{c,\sigma_c} = \frac{\sum_i^N \left(\hat{c}_i^{c,\sigma_c} * E_i^{c,\sigma_c} \mid BL_i = 1\right)}{\widehat{AS}^{c,\sigma_c}} + \frac{\sum_i^N \left(\hat{c}_i^{c,\sigma_c} * E_i^{c,\sigma_c} \mid BL_i = 0 \ \& \ r_i + p_c * c_i \leq 0\right)}{\widehat{AS}^{c,\sigma_c}}. \quad \text{(Eq. 29)}$$

The first term in Equation 29 is the fraction offsets issued based on the predicted sequestration of non-additional enrollees whose baselines were incorrectly estimated, and the second term is the fraction of offsets issued based on the estimated carbon sequestration of producers who were assigned and overly generous carbon sequestration value.

The inclusion of error in carbon sequestration estimates necessitates a similar adjustment to calculate efficiency loss (EL^{c,σ_c}) as well. There is still a possibility of producers inefficiently not participating in the program because of an incorrectly assigned baseline that indicates they would have used the practice in the absence of the program when the opposite is true. But, even if the policy maker correctly estimates returns to adoption, it is possible that a producer inefficiently refrains from enrolling enroll because the policy maker underestimates its carbon sequestration potential. The efficiency loss for the payment-for-sequestration program when there is error in the policy maker's estimates of producers' carbon sequestration potential (EL^{c,σ_c}) includes the carbon sequestration missed out on from both types of inefficient non-participants:

$$EL^{c,\sigma_c} = \frac{\sum_i^N (c_i - BLC_i \mid r_i + p_c * (c_i - BLC_i) > 0 \ \& \ E_i^{c,\sigma_c} = 0)}{\sum_i^N (c_i - BLC_i \mid r_i + p_c * (c_i - BLC_i) > 0 \ \& \ BL_i = 0)} \qquad \text{(Eq. 30)}$$

5.3 Simulation of program performance as a function of net return and carbon sequestration observation error

As earlier, we now test the performance of the two program types when the policy maker estimates producers' sequestration potential with error. We use the four metrics defined above for the case involving carbon observation error and examine how they vary as a function of the policy maker's errors in observing producers' net returns, $\varepsilon_{i,r}$, and carbon sequestration potential, $\varepsilon_{i,c}$. We use the same values as in the earlier simulation for the two incentives, $p_p = 0.5$ and $p_c = 1$, and the same distributional assumptions for r_i, c_i, and $\varepsilon_{i,r}$: $f_r(r_i) \sim N(-0.5, 1^2)$, $f_f(c_i) \sim U[0,1]$, and $\varepsilon_{i,r} \sim N(0, \sigma_r)$. For the policy maker's estimates of producers' carbon sequestration potential $\varepsilon_{i,c}$, we assume their error in observing c_i is drawn from a normal distribution with mean zero and standard deviation σ_c.

In figure 2, we depict the average efficiency loss (EL^{σ_c}), average cost of a unit of additional sequestration (AC^{σ_c}), percent of offsets from inefficient enrollees (IP^{σ_c}), and additional sequestration (AS^{σ_c}) per plot as a function of the standard deviation in the net return observation error (σ_r) for three values of the standard deviation in carbon sequestration error: $\sigma_c = 0, 0.2, \& \ 0.4$. We do not display the values of σ_c for the payment-for-practice program results because they are identical across all values of σ_c. This is because the policy maker's estimates of carbon sequestration are unbiased given $\varepsilon_{i,c} \sim N(0, \sigma_c)$.

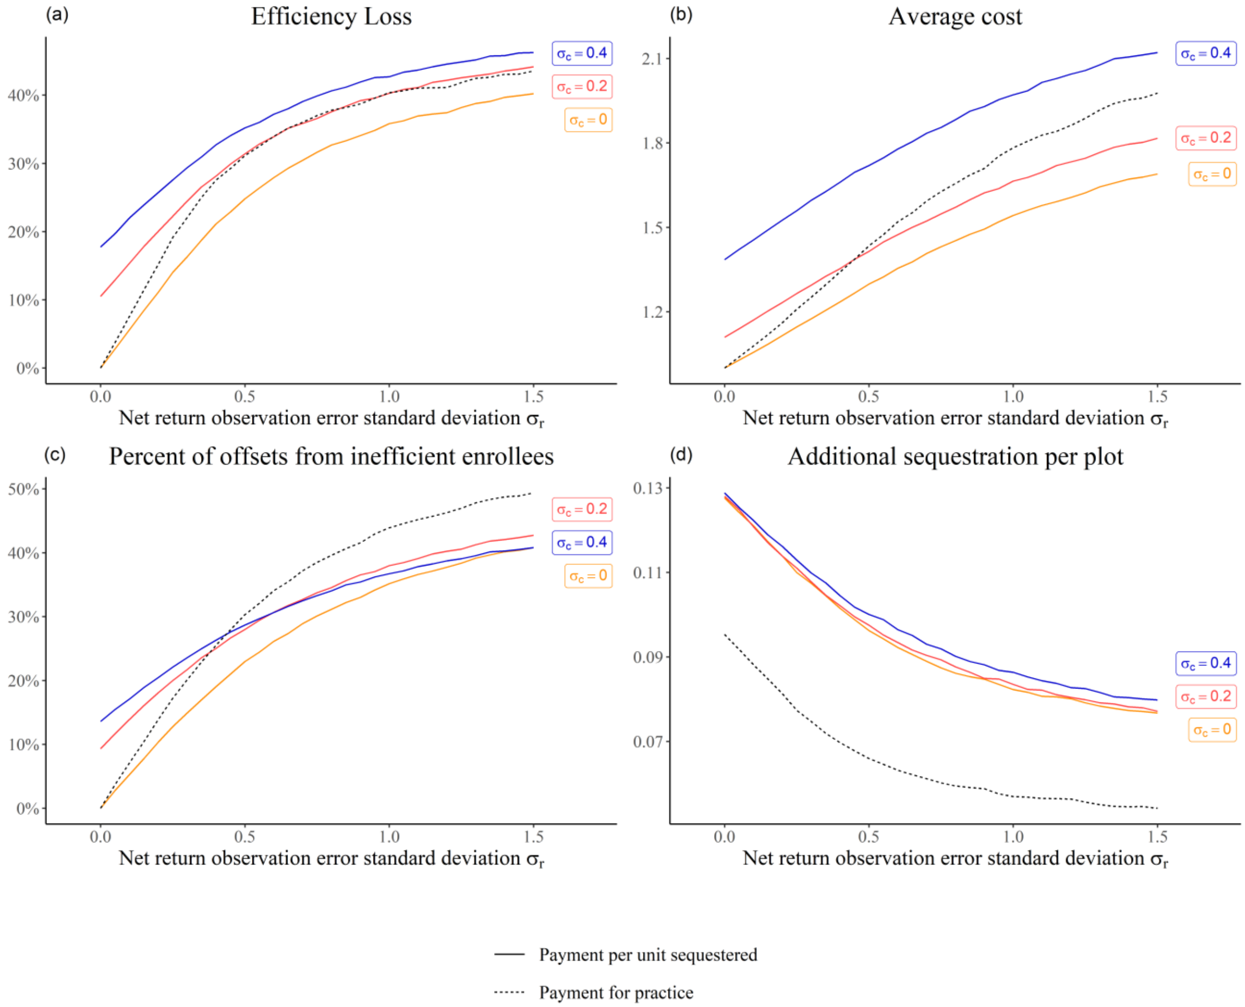

Figure 1. Effect of error in policy maker's estimates of producers' net returns to, and carbon sequestration produced by, practice adoption. The x-axis indicates the standard deviation of the policy maker's error, $\boldsymbol{\sigma_r}$. Individual lines are labelled with their respective values for the standard deviation of the policy maker's error in estimating carbon sequestration, $\boldsymbol{\sigma_c}$.

As in Figure 1, the payment-for-sequestration program results in panel (d) of Figure 2 depict greater additional sequestration for all values of net return observation error (σ_r), and this holds true for all values of carbon sequestration observation error (σ_r). For the rest of the metrics, however, the payment-for-practice program sometimes outperforms the payment-for-sequestration program. For example, the average cost of an additional unit of sequestration in panel (b) and the percent of offsets from inefficient enrollees in panel (c) of Figure 2 are both lower for the payment-for-practice program when the net return observation error (σ_r) is relatively low in comparison to the carbon sequestration observation error (σ_c). Conversely, when the net return observation error is relatively high, the payment-for-sequestration program outperforms the payment-for-practice program in terms of its average cost and percent of offsets from inefficient enrollees. This does not hold true, however, at very high levels of error in the policy maker's estimates of carbon sequestration. For the highest value of carbon sequestration observation error, $\sigma_c = 0.4$, the payment-for-practice program's efficiency loss and average cost are lower for all values of net return observation error, σ_r.

From the results in Figure 2, the payment-for-sequestration program's efficiency is clearly reduced by the presence of error in the policy maker's estimates of carbon sequestration. When there is any error in carbon sequestration estimates ($\sigma_c > 0$), the payment-for-sequestration program suffers large efficiency losses even in the absence of net return observation error ($\sigma_r = 0$). These efficiency losses represent the carbon which is not sequestered because producers received carbon sequestration estimates below their real sequestration potential. However, as shown in panel (d) of Figure 2, the greater additional sequestration for the payment-for-sequestration program indicates the carbon represented by these efficiency losses was somehow recouped. The curves for the payment-for-sequestration program in panel (c) of Figure 2 depict the new source of additional carbon: producers who would not have participated in the absence of carbon sequestration error but who enroll in the program because the policy maker overestimates their carbon sequestration. The carbon sequestered by these individuals is truly additional, but the policy maker is overpaying these individuals due to its overly generous sequestration estimates. The cumulative result of these effects is the increasing average cost of the payment-for-sequestration program.

6 Effect of increasing baseline scale on voluntary carbon credit program performance

To examine the impact of scale on the efficiency of the payment-for-practice and payment-for-sequestration programs, we first redefine the program enrollment decision and outcomes at a jurisdictional level. We consider a setting identical to that of the previous section involving error in policy maker's estimates of producers' net returns to practice adoption and the sequestration due to each producer's use of the practice. However, instead of evaluating outcomes across all individuals indexed by i, we now consider the aggregate outcomes for all producers in identically sized jurisdictions indexed by j. We hold the total number of producers (N) constant and define the number of jurisdictions as N_j, so the number of producers in each jurisdiction is $N_{i,j} = N/\, N_j$.

Within each jurisdiction, the individual producers are indexed by i running from $i = 1$ to $i = N_{i,j}$. The characteristics of individual producer's, $r_{i,j}$ and $c_{i,j}$, share the same features as their counterparts in the individual-scale model, r_i and c_i. The baseline adoption behavior for the conservation agriculture practice in jurisdiction j is the sum of the individual baselines for all producers in jurisdiction j such that the true baseline use of the practice in jurisdiction j is:

$$BL_j = \sum_{i=1}^{N_{i,j}} BL_{i,j}\,;\; where\; BL_{i,j} = \begin{cases} 1 \;\; if\; r_{i,j} > 0 \\ 0 \;\; if\; r_{i,j} \le 0 \end{cases}. \qquad \text{(Eq. 31)}$$

The carbon that would be sequestered in jurisdiction j based on its baseline practice use is:

$$BLC_j = \sum_{i=1}^{N_{i,j}} BLC_{i,j}\,;\ where\ BLC_{i,j} = \begin{cases} c_{i,j} & if\ r_{i,j} > 0 \\ 0 & if\ r_{i,j} \leq 0 \end{cases}. \qquad \text{(Eq. 32)}$$

Like in the individual-scale case, the policy maker assigns a baseline level of practice use for each jurisdiction based on its estimates of producers' returns to adopting the practice, $\hat{r}_{i,j} = r_{i,j} + \varepsilon_{i,j,r}$. We continue to assume the error in the policy maker's estimate of net returns, $\varepsilon_{i,j,r}$, is symmetric around zero, independent of f_r, and has standard deviation σ_r. The policy maker constructs its prediction of jurisdiction j's baseline by adding up the assigned baselines of all the constituent producers:

$$\widehat{BL}_j = \sum_{i=1}^{N_{i,j}} \widehat{BL}_{i,j}\,;\ where\ \widehat{BL}_{i,j} = \begin{cases} 1 & if\ \hat{r}_{i,j} > 0 \\ 0 & if\ \hat{r}_{i,j} \leq 0 \end{cases}. \qquad \text{(Eq. 33)}$$

In similar fashion, the policy maker continues to estimate producers' carbon sequestration with some error such that: $\hat{c}_{i,j} = c_{i,j} + \varepsilon_{i,j,c}$. We likewise assume the error in the policy maker's estimate of net returns, $\varepsilon_{i,j,c}$, is symmetric around zero, independent of f_r and $\varepsilon_{i,j,r}$, and has standard deviation σ_c. As such, the baseline carbon sequestration that the policy maker assigns for jurisdiction j is:

$$\widehat{BLC}_j = \sum_{i=1}^{N_{i,j}} \widehat{BLC}_{i,j}\,;\ where\ \widehat{BLC}_{i,j} = \begin{cases} \hat{c}_{i,j} & if\ \hat{r}_{i,j} > 0 \\ 0 & if\ \hat{r}_{i,j} \leq 0 \end{cases}. \qquad \text{(Eq. 34)}$$

6.1 Jurisdiction scale payment-for-practice program

For the jurisdictional payment-for-practice program, each jurisdiction is offered a fixed compensation p_p times the number of its constituent producers who adopt the conservation agriculture practice above the assigned jurisdictional baseline. As an example, consider a program in which countries are compensated on a per-acre basis for reforestation efforts. The jurisdiction, a country in this example, would be assigned a baseline that reflects the expected forest acreage across its entire geographic extent in the absence of any reforestation initiatives.

To determine its compensation, the policy maker would subtract this country-level baseline from the total forested acreage observed after the country institutes its payment-for-practice reforestation program and multiply the result by p_p. Instead of measuring whether individual producers in each country reforest their particular plots of land, this jurisdictional approach focuses on the aggregate outcomes across all members of jurisdictions.

Note that if each jurisdiction contains a single producer ($N_{i,j} = 1$), this is identical to the individual-scale payment-for-practice program examined above. But, when the jurisdiction size is greater than one, it is necessary to define how each jurisdiction decides whether to participate or not. To do so, similar to van Benthem and Kerr (2013), we assume that jurisdictions know the returns to adoption with certainty and that they are informed of their assigned jurisdictional baseline for practice use ($\widehat{BL}_j$) by the policy maker. In essence, the assumption that jurisdictions know their constituents returns to adoption with certainty is akin to stating each jurisdiction can perfectly predict the total cost of adoption among its constituents for any given level of adoption.

To demonstrate the jurisdictions' enrollment decision-making process, consider a jurisdiction j which is offered the payment-for-practice incentive by the policy maker p_p. Jurisdiction j computes the potential benefit to enrollment by totaling how many of its constituent producers would adopt the practice given incentive p_p. We define $N_j^{p_p}$ as this total number of producers in jurisdiction j who would adopt given the incentive p_p, and note that this sum includes producers who would have adopted without an incentive. Then, jurisdiction j's potential return to enrollment is: $p_p * \left(N_j^{p_p} - \widehat{BL}_j\right)$. To determine the cost of enrolling, each jurisdiction then totals the individual cost of adoption across all its constituents who would alter their behavior when offered p_p, yielding: $\sum_{i|r_{i,j}\in(-p_p,0]}^{N_{i,j}} r_{i,j}$. A jurisdiction will enroll in the payment-for-practice program, reflected by the indicator variable E_j^p equaling one, if the potential benefits from enrollment are greater than the total cost of adoption born by its constituents:

$$E_j^p = \begin{cases} 1 \; if \; p_p * \left(N_j^{p_p} - \widehat{BL}_j\right) > \sum_{i|r_{i,j} \in (-p_p,0]}^{N_{i,j}} r_{i,j} \\ 0 \; if \; p_p * \left(N_j^{p_p} - \widehat{BL}_j\right) \leq \sum_{i|r_{i,j} \in (-p_p,0]}^{N_{i,j}} r_{i,j} \end{cases} . \qquad \text{(Eq. 35)}$$

In equation 35, notice that the jurisdiction does not make its enrollment decision based on the expenditures it would make internally for each of its constituents who use the practice, so we are not modeling how the jurisdiction disperses payments amongst producers. Instead, we assume that all $N_j^{p_p}$ producers in jurisdiction j will adopt the conservation agriculture practice if the jurisdiction enrolls in the payment-for-practice program, meaning $E_j^p = 1$.

To determine the efficiency of the jurisdiction-scale payment-for-practice program in the presence of error in both the estimates of net returns and carbon sequestration potential, we focus our analysis on the average cost of an additional unit of sequestration $(AC^{p,j})$. The j component in the superscript of $AC^{p,j}$ indicates this term refers to the average cost of the payment-for-practice program at a jurisdictional scale. To calculate the average cost, we add up the cost of the program to the policy maker across all jurisdictions to produce:

$$TC^{p,j} = \sum_{j=1}^{N_j} E_j^p * p_p * \left(N_j^{p_p} - \widehat{BLC}_j\right) . \qquad \text{(Eq. 36)}$$

Then, we find the total additional sequestration that occurs because of jurisdictions' enrollment decisions using:

$$AS^{p,j} = \sum_{j=1}^{N_j} E_j^p * \sum_{i|r_{i,j} \in (-p_p,0]}^{N_{i,j}} c_{i,j} . \qquad \text{(Eq. 37)}$$

The average cost per unit of additional sequestration is then:

$$AC^{p,j} = \frac{\sum_{j=1}^{N_j} E_j^p * p_p * \left(N_j^{p_p} - \widehat{BLC}_j\right)}{\sum_{j=1}^{N_j} E_j^p * \sum_{i|r_{i,j} \in (-p_p,0]}^{N_{i,j}} c_{i,j}} . \qquad \text{(Eq. 38)}$$

6.2 Jurisdiction scale payment-for-sequestration program

When the jurisdictional payment-for-sequestration program is available, each jurisdiction is offered compensation p_c times the amount of carbon the policy maker estimates is sequestered above the assigned jurisdictional baseline $(\widehat{BLC}_j)$. Like for the jurisdictional payment-for-practice program, we assume the jurisdiction knows the net returns of its producers with certainty and determines its returns to participation based on its constituent producers' potential returns given the payment-for-sequestration incentive it receives from the policy maker, p_c. As was the case for the jurisdictional payment-for-practice program, if jurisdictions are comprised of one producer $(N_{i,j} = 1)$, the jurisdiction-scale and individual-scale payment-for-sequestration programs are identical. However, unlike producers' returns to practice adoption, we do not assume that the jurisdiction knows its constituents' true carbon sequestration potential $(c_{i,j})$ with certainty. Instead, we assume the jurisdiction relies on the policy maker's estimates of its constituent producers' carbon sequestration potential $(\hat{c}_{i,j})$ when calculating individual producers' returns in the jurisdictional payment-for-sequestration program.

When $N_{i,j} > 1$ for the jurisdictional payment-for-sequestration program, jurisdictions compute their return to enrollment by totaling the estimated carbon sequestration of all constituent producers who would adopt the practice when offered $p_c * \hat{c}_{i,j}$. Let $\hat{C}_j^{p_c}$ be this total estimated carbon sequestration of producers in jurisdiction j who would participate given the incentive $p_c * \hat{c}_{i,j}$, defined as:

$$\hat{C}_j^{p_c} = \sum_{i|r_{i,j}+p_c*\hat{c}_{i,j} > 0 \ \& \ \hat{c}_{i,j} > 0}^{N_{i,j}} \hat{c}_{i,j} \ . \qquad \text{(Eq. 39)}$$

Jurisdiction j's potential return to enrollment is then $p_c * \left(\hat{C}_j^{p_c} - \widehat{BLC}_j\right)$, and the total cost of participating summed across all individuals in a jurisdiction who would alter their behavior is

$\sum_{i|r_{i,j}\in(-p_c*\hat{c}_{i,j},0]\ \&\ \hat{c}_{i,j}>0}^{N_{i,j}} r_{i,j}$. Taken together, these two quantities determine if jurisdiction j will enroll in the jurisdictional payment-for-sequestration program, reflected by the indicator variable E_j^c equaling one, according to:

$$E_j^c = \begin{cases} 1\ if\ p_c * \left(\hat{C}_j^{p_c} - \widehat{BLC}_j\right) > \sum_{i|r_{i,j}\in(-p_c*\hat{c}_{i,j},0]\ \&\ \hat{c}_{i,j}>0}^{N_{i,j}} r_{i,j} \\ 0\ if\ \ p_c * \left(\hat{C}_j^{p_c} - \widehat{BLC}_j\right) \leq \sum_{i|r_{i,j}\in(-p_c*\hat{c}_{i,j},0]\ \&\ \hat{c}_{i,j}>0}^{N_{i,j}} r_{i,j} \end{cases}. \qquad \text{(Eq. 40)}$$

So, when a jurisdiction chooses to participate, meaning $E_j^c = 1$, the jurisdiction's total estimated carbon sequestration across its constituent producers equals $\hat{C}_j^{p_c}$. Like for the payment-for-practice jurisdictional program, we focus our analysis on the average cost of an additional unit of sequestration $(AC^{c,j})$. The total cost of the program born by the policy maker across all jurisdictions is defined as:

$$TC^{c,j} = \sum_{j=1}^{N_j} E_j^c * p_c * \left(\hat{C}_j^{p_c} - \widehat{BLC}_j\right). \qquad \text{(Eq. 41)}$$

While the total additional sequestration due to the jurisdictional payment-for-sequestration program is:

$$AS^{c,j} = \sum_{j=1}^{N_j} E_j^p * \sum_{i|r_{i,j}\in(-p_c*\hat{c}_{i,j},0]\ \&\ \hat{c}_{i,j}>0}^{N_{i,j}} c_{i,j}\ . \qquad \text{(Eq. 42)}$$

The average cost per unit of additional sequestration is then:

$$AC^{c,j} = \frac{\sum_{j=1}^{N_j} E_j^c * p_c * \left(\hat{C}_j^{p_c} - \widehat{BLC}_j\right)}{\sum_{j=1}^{N_j} E_j^p * \sum_{i|r_{i,j}\in(-p_c*\hat{c}_{i,j},0]\ \&\ \hat{c}_{i,j}>0}^{N_{i,j}} c_{i,j}}. \qquad \text{(Eq. 43)}$$

6.3 Numerical simulation of jurisdictional voluntary carbon credit programs

To test the performance of the jurisdictional versions of the payment-for-practice and payment-for-sequestration programs, we narrow our focus to the impact of jurisdiction size $(N_{i,j})$ on the average cost of an additional unit of sequestration $(AC^{p,j}, AC^{c,j})$. To account for the impact of error in the policy maker's estimates of net returns and carbon sequestration, we test this relationship at several values for the standard deviations in returns to adoption (σ_r) and carbon sequestration (σ_c). We make the same distributional assumptions as in the earlier simulations such that $f_r(r_{i,j}) \sim N(-0.5, 1^2)$, $f_f(c_{i,j}) \sim U[0,1]$, $\varepsilon_{i,j,r} \sim N(0, \sigma_r)$, and $\varepsilon_{i,j,c} \sim N(0, \sigma_c)$. We also use the same values for the two incentives, $p_p = 0.5$ and $p_c = 1$, and hold the total number of producers in the simulation constant across runs at $N = 10{,}000$.

In figure 3, we depict the average cost of an additional unit of sequestration ($AC^{p,j}, AC^{c,j}$) as function of the jurisdiction size ($N_{i,j}$) for nine different scenarios. The nine scenarios depicted are defined by the combination of three values for the standard deviation in policy maker's estimates of returns to adoption ($\sigma_r = 0.1, 0.3, 0.5$) and carbon sequestration ($\sigma_c = 0.1, 0.2, 0.3$). The plotted lines depict the mean average cost of an additional unit of sequestration ($AC^{p,j}, AC^{c,j}$) across the 50 simulation runs performed for each combination of jurisdiction size ($N_{i,j}$), standard deviation in return observation error (σ_r), and standard deviation in carbon sequestration error (σ_c).

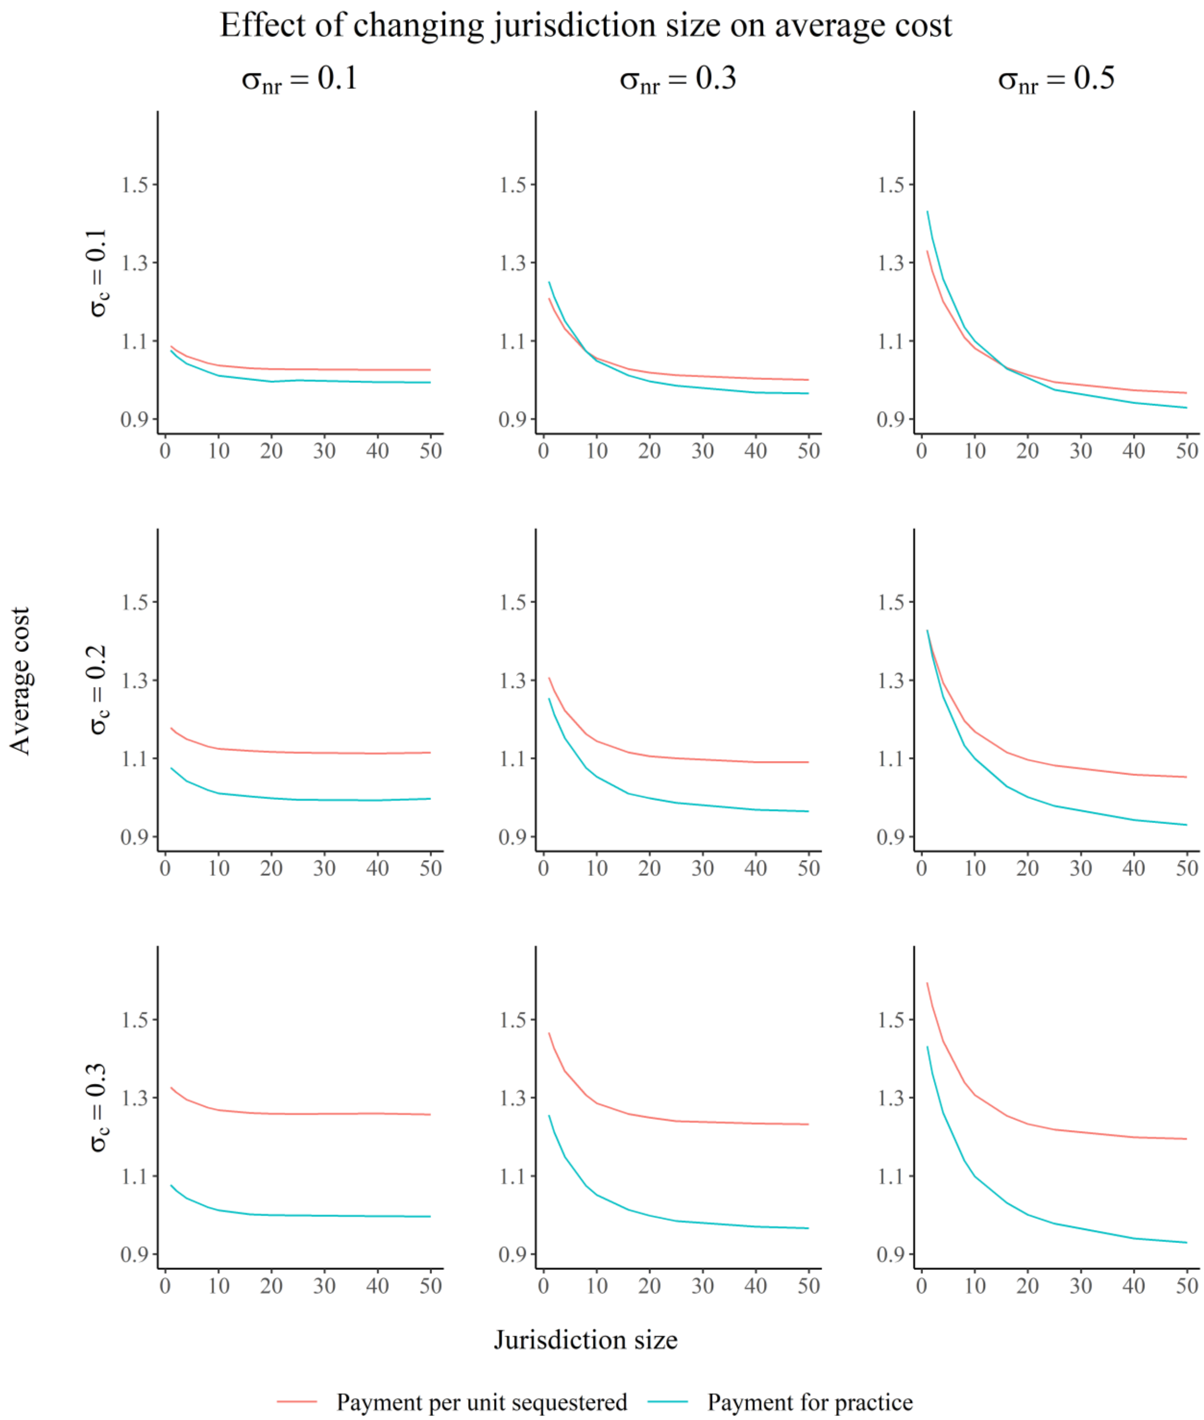

Figure 1. Effect of jurisdiction size on the average cost of an additional unit of sequestration by compensation mechanism.

The top row of graphs in figure 3 depict the effect of increasing the size of jurisdictions when the error in policy maker's estimates is lowest ($\sigma_c = 0.1$). As expected from the earlier results displayed in Figure 2, the average cost of an additional unit of sequestration is lower for the payment-for-sequestration program when jurisdictions are small in size and there is little error in carbon sequestration estimates relative to the error in estimates of net returns. But, across all three values for σ_r, the payment-for-practice program eventually becomes the more cost-effective program as the size of jurisdictions increases when $\sigma_c = 0.1$. One potential reason for this is the selection for greater carbon contributions involved in the payment-for-sequestration payment mechanism. In comparing the graphs in the top rows of Figure 3, it becomes clear that error in the policy maker's estimates of net returns (σ_r) mediates the relationship between scale and average cost of the payment-for-practice program. For example, holding the error for carbon sequestration constant at $\sigma_c = 0.1$, an increase in σ_r from 0.3 to 0.5 causes the jurisdiction size at which the payment-for-practice program outperforms the payment-for-sequestration program to increase.

For the intermediate or high levels of error in estimates of carbon sequestration ($\sigma_c = 0.2, 0.3$), shown in the second and third rows of graphs in Figure 3, the greater impact of increasing jurisdiction size on the payment-for-practice program causes it to outperform the payment-for-sequestration program in almost every circumstance. The exception to this statement is the case with the intermediate level of error in carbon sequestration estimates ($\sigma_c = 0.2$), very small jurisdiction sizes, and the highest level of error in estimates of producers' net returns, ($\sigma_r = 0.5$). One reason for the markedly superior performance of the payment-for-practice program is our assumption that the jurisdiction uses the policy maker's estimates of carbon sequestration ($\hat{c}_{i,j}$) to calculate the returns to participation for its constituents. This assumption implies that the jurisdiction does not have additional information on the rates of sequestration amongst its producers in comparison to the policy maker.

Overall, the effect of increasing program scale mitigates the inefficiencies caused by the policy maker's error in estimating producers' net returns to a greater degree. As the payment-for-practice program is only affected by the inefficiencies related to net return observation error, its performance improves to a greater degree as the scale of the program increases. The insensitivity of the payment-for-practice program's efficiency to the policy maker's error in estimating carbon sequestration is readily apparent when viewing a single column of Figure 3. In contrast, the payment-for-sequestration program, being affected by the inefficiencies born of carbon sequestration error, benefits from the effect of increasing the program's scale to a lesser degree.

7 Conclusion

In this work we first demonstrate that the superior performance of a payment-for-sequestration, or results-based, carbon offset program relies on the policy maker's ability to accurately predict the sequestration outcomes of individual producers. The payment-for-sequestration program excels because it selects for producers whose actions produce larger results, a greater increase in carbon sequestration due to adoption of a conservation agriculture practice in this context. Furthermore, the selection caused by the payment-for-practice compensation mechanism causes the results-based policy to outperform the actions-based offset policy even when the policy maker predicts producers' counterfactual behavior with significant error. However, the unequivocal, superior performance of the payment-for-sequestration program only occurs when the policy maker knows the results of producers' actions, or the sequestration each producer would contribute, with certainty.

When we introduce error in the policy makers estimates of the carbon sequestration resulting from producers using the practice, our results indicate that the payment-for-practice program outperforms the payment-for-sequestration program when the uncertainty in carbon sequestration values is high relative to the uncertainty in estimates of producers' returns to adoption. This result builds on the conclusions of White and Hanley (2016) who find that an input-based contract, analogous to the payment-for-practice program, is more cost effective than an output-based contract, similar to our payment-for-sequestration program, because it has a lower cost approach to adverse selection. With our model, we demonstrate that the cost of either approach to adverse selection, meaning the compensation mechanism chosen, will depend on the policy maker's ability to accurately estimate crucial producer characteristics. For example, the payment-for-sequestration mechanism selecting for producers with greater sequestration potential becomes increasingly costly as the error in carbon sequestration estimates rises. This increasing cost of mitigating adverse selection is due to the growth in costly contributions of inefficient participants who are assigned inflated carbon sequestration baselines.

Finally, in examining the impact of a jurisdictional approach to setting baselines and evaluating outcomes, we find that the positive relationship between program scale and efficiency established in van Benthem and Kerr (2013) applies to both the payment-for-practice and payment-for-sequestration programs. In addition, we find that increasing the program scale improves the cost-efficiency of the payment-for-practice program to a greater degree, such that it becomes the more cost-effective jurisdictional program in the majority of scenarios. Future work will clarify whether this result is due to the assumptions made in our work concerning the independence of model parameters or the distributions used in the numerical simulations.

References

117th Congress. 2022. *Inflation Reduction Act of 2022*.

Antle, John, Susan Capalbo, Siân Mooney, Edward Elliott, and Keith Paustian. 2003. "Spatial Heterogeneity, Contract Design, and the Efficiency of Carbon Sequestration Policies for Agriculture." *Journal of Environmental Economics and Management* 46 (2): 231–50.

Antle, John M., Susan M. Capalbo, Keith Paustian, and Md Kamar Ali. 2007. "Estimating the Economic Potential for Agricultural Soil Carbon Sequestration in the Central United States Using an Aggregate Econometric-Process Simulation Model." *Climatic Change* 80 (1–2): 145–71.

Benthem, Arthur van, and Suzi Kerr. 2013. "Scale and Transfers in International Emissions Offset Programs." *Journal of Public Economics* 107: 31–46.

Claassen, Roger, Eric N. Duquette, and David J. Smith. 2018. "Additionality in U.S. Agricultural Conservation Programs." *Land Economics* 94 (1): 19–35.

EPA. 2020. "Inventory of U.S. Greenhouse Gas Emissions and Sinks: 1990-2020."

Hanley, Nick, Simanti Banerjee, Gareth D. Lennox, and Paul R. Armsworth. 2012. "How Should We Incentivize Private Landowners to 'produce' More Biodiversity?" *Oxford Review of Economic Policy* 28 (1): 93–113.

Montero, Juan Pablo. 2000. "Optimal Design of a Phase-in Emissions Trading Program." *Journal of Public Economics* 75 (2000): 273–91.

Ogle, Stephen M., F. Jay Breidt, Mark Easter, Steve Williams, Kendrick Killian, and Keith Paustian. 2010. "Scale and Uncertainty in Modeled Soil Organic Carbon Stock Changes for US Croplands Using a Process-Based Model." *Global Change Biology* 16 (2): 810–22.

Plastina, Alejandro, and Oranuch Wongpiyabovorn. 2021. "How to Grow and Sell Carbon Credits in US Agriculture." *Ag Decision Maker. Iowa State University Extension and Outreach*.

Richards, Kenneth R, and Robert N Stavins. 2005. "The Cost of U.S. Forest-Based Carbon Sequestration." Arlington, Virginia.

Sandmo, Agnar. 1975. "Optimal Taxation in the Presence of Externalities." *The Swedish Journal of Economics* 77 (1): 86–98.

Sawadgo, Wendiam, and Alejandro Plastina. 2021. "Do Cost-Share Programs Increase Cover Crop Use? Empirical Evidence from Iowa." *Renewable Agriculture and Food Systems* 1 (2016): 1–9.

Sperow, Mark. 2016. "Estimating Carbon Sequestration Potential on U.S. Agricultural Topsoils." *Soil and Tillage Research* 155: 390–400.

Swan, Amy, Mark Easter, Adam Chambers, Kevin Brown, Stephen Williams, Jeff Creque, John Wick, and Keith Paustian. 2020. "COMET-Planner Carbon and Greenhouse Gas Evaluation for NRCS Conservation Practice Planning."

White, Ben, and Nick Hanley. 2016. "Should We Pay for Ecosystem Service Outputs, Inputs or Both?" *Environmental and Resource Economics* 63 (4): 765–87.

Wuepper, David, and Robert Huber. 2021. “Comparing Effectiveness and Return on Investment of Action- and Results-Based Agri-Environmental Payments in Switzerland.” *American Journal of Agricultural Economics*, no. January: 1585–1604.