

***A Simulating Annealing Approach to Non-Market
Environmental Benefit Aggregation***

Stephen Hynes¹, Nick Hanley², Cathal O'Donoghue³

Working Paper No. 166

December 2010

Discipline of Economics
National University of Ireland, Galway

<http://www.economics.nuigalway.ie>

A Simulating Annealing Approach to Non-Market Environmental Benefit Aggregation.

Stephen Hynes¹, Nick Hanley², Cathal O'Donoghue³

1. Socio-Economic Marine Research Unit, J.E. Cairnes School of Business and Economics, National University of Ireland, Galway, Ireland, Email: stephen.hynes@nuigalway.ie

2. Department of Economics, University of Stirling, Scotland,

3. Environmental Modelling Unit, RERC, Teagasc, Ireland

Abstract

This paper considers the use of a “combinatorial optimization” technique in the aggregation of environmental benefit values. Combinatorial optimization is used to statistically match population census data to a Contingent Valuation survey. The matched survey and census information is then used to produce regional and national total WTP figures. These figures are then compared to figures derived using more standard approaches to calculating aggregate environment benefit values. The choice of aggregation approach is shown to have a major impact upon estimates of total benefits at a regional level, especially when the target population displays considerable heterogeneity across space.

Keywords: Combinatorial Optimization, Contingent Valuation, spatial microsimulation, non-market benefit aggregation

JEL: Q51, Q57

This paper was written as part of the Rural Stimulus Funded project, financed by The Irish Department of Agriculture, Fisheries and Food. We thank two referees for their comments on an earlier version of this paper.

1. Introduction

The benefits and costs of environmental policy can be expected to vary significantly over space, which causes problems for the estimation of policy or program net benefits (Nijkamp, 2002). This is particularly true for the aggregation of welfare estimates derived from non-market environmental valuation techniques. If one attempts to transfer a nationally-representative random sample's average welfare estimates to a particular region of interest, it is difficult to fully account for the heterogeneous characteristics in that region's population, so that the analyst may therefore under- or over-estimate the aggregate regional impacts from an environmental policy change. To address this problem, we employ the principle of synthetic estimation (Williamson et al., 1998), using a method known as "combinatorial optimization" to take into account spatial heterogeneity across target populations in the aggregation process at different spatial levels.

Consider, as an example, a common Contingent Valuation (CV) survey situation. Assume that we have a random sample of 1500 urban households in Ireland each of which has been asked their willingness to pay (WTP) for an improvement in the quality of their local drinking water supply due to the implementation of new filtering technology. This sample of households is representative of the entire national population, but is not representative of each small census area jurisdiction in the country. Now suppose we are particularly interested in the aggregate WTP of residents in one jurisdiction, Galway City, for a cleaner water supply. We know from census records that households in Galway City display higher incomes and higher education levels than the average urban dweller in the rest of the country. This means that using the average WTP in our sample to aggregate up to the target population may under-estimate the aggregate WTP of Galway city for a clean water supply.

Using a combinatorial optimization methodology we can instead select individuals from the national CV survey in order to define a synthetic dataset which is representative of people living in Galway city as described in the census “small area statistics” for this jurisdiction. On completion of the statistical matching process we have a synthetic *population* of individuals, derived from the CV sample, representing the target Galway population by income and education levels. We can then use these synthetic households’ WTP values to get a truer aggregate estimate for WTP in Galway city than would be possible if we multiplied the original sample’s average WTP by the population in the jurisdiction of interest (since we are accounting for the fact that Galway city displays different characteristics from the average urban area in Ireland).

The case study used in this paper considers a CV study that asks Irish farmers their WTP to conserve an endangered farmland bird, the Corncrake (*Crex crex*). Corncrakes depend on the maintenance of suitable farm habitats for nesting success, and have been in rapid decline in Ireland due to changes in farming methods over the last 25 years. The CV survey and Census of Agriculture data are statistically matched to produce small area population environmental benefit micro data estimates for the year 2006 using a combinatorial optimization technique. The focus of attention in this paper is on the aggregation methodology employed, which, in theory, could be used in conjunction with any stated preference survey. For this reason we do not spend time discussing the WTP study itself, but instead concentrate on how the combinatorial optimization method can be used to generate regional and national aggregated WTP estimates¹.

The combinatorial optimization method used in this paper could provide three benefits for non-market valuation. Firstly, the technique allows us to efficiently and accurately expand a sample's individual welfare estimates to any particular spatial scale of interest. Secondly, the methodology allows us to take into account measurable spatial heterogeneity of the target population in the aggregation process. Finally, given time and money constraints it may be impossible to conduct a comprehensive enough valuation survey to obtain an adequate sample from every small area jurisdiction we are interested in. The combinatorial optimization method, however, facilitates the estimation of WTP in diverse jurisdictions by re-weighting responses from a national survey.

In the next section, we discuss the aggregation of environmental benefit values. Section 3 then briefly reviews the design of the WTP survey used as a case study, and discusses the datasets used in the combinatorial optimization process. In section 4 we discuss the combinatorial optimization approach used to aggregate the WTP values for corncrake conservation at varying levels of regional jurisdiction. Model results and the aggregated WTP estimates are presented in section 5. Finally, section 6 concludes with some recommendations for further research.

2. The Aggregation of Environmental Benefit Values

Aggregating environmental benefit values is the process whereby sample mean values of WTP or other welfare measures are converted to a total value figure for a population (Hanley et al., 2003). Bateman et al. (2006) and Smith (1993) point out that because the methods for measuring non-market benefit values are based on analyses of individual behaviour, a problem arises in knowing how changes in a resource will affect aggregate values, since aggregation will depend on both the

benefits per person and the population of beneficiaries (the extent of the market, and the characteristics of members of that market).

A number of issues regarding the aggregation of environmental value estimates can be resolved using the combinatorial optimization approach developed in this paper. At a basic level, Loomis (1987) states that the problem of generalizing results from a sample to the population relates to low response rates and small sample sizes. By statistically matching our sample of farmers and their associated WTP values with associated farm characteristics obtained from a comprehensive census, we can generate a synthetic dataset with individual WTP values for an entire population. The combinatorial optimization approach also alleviates the problem highlighted by Morrison (2000) in relation to how representative the sample of respondents is of the actual socioeconomic and demographic characteristics of the population in question.

Another concern relates to spatial representativeness in the aggregation process (Bateman et al., 2006). This issue is also relevant for our study. For example, different regions of Ireland are represented by different types of farmers. The western seaboard is predominately represented by relatively small, extensively-operated, livestock farmers, whilst the south east of the country is populated by larger, more intensive dairy and tillage (arable) farm holdings. In any aggregation process, it is vital that these spatial differences in farm size and type be taken into account, especially if we wish to examine regional variations in the total benefits of the corncrake conservation program which forms the empirical focus of this paper.

Another issue that needs consideration when dealing with environmental benefit aggregation is aggregation error. Aggregation errors arise when estimates from a

sample are aggregated to represent the total population value. These errors are inversely related to the degree of correspondence between the sample and the population (Rosenberger and Stanley, 2006). Calculating the extent of this error when aggregating up to a total population is a difficult prospect, but ensuring the correspondence of socio-economic characteristics between the sample and population (as is done in the combinatorial optimization approach developed here) should increase the accuracy of the aggregation. In the discussion at the end of the paper we also propose a follow-up field study that would facilitate a test of aggregation error where our estimates of total WTP at the Electoral Division (ED) level of geographical area are compared to the actual total WTP of farmers in the corresponding EDs, as revealed by the interviewing of all farmers in each of the EDs.

The present paper conducts what is, to our knowledge, the first systematic aggregation of contingent valuation data which accounts for heterogeneity in the target population using combinatorial optimization techniques. Our results show that the choice of aggregation approach can have a major impact upon estimates of total benefits at a regional level, especially when the target population displays a large amount of heterogeneity across space.

3. Data and WTP Format

In this section we describe the data used and the format of the willingness to pay questions. *The National Farm Survey (NFS)* is collected as part of the Farm Accountancy Data Network (FADN) of the European Union. The aim of this network is to gather accounting data from farms in all member states of the EU for the determination of incomes and business analysis of agricultural holdings (FADN, 2005). Table 1 provides summary statistics of a number of key variables in the NFS

sample. In the 2006 NFS, contingent valuation (CV) was employed by the authors to estimate the value to the Irish farmer of conserving the corncrake, a rare farmland bird species. Questions were asked in terms of farmers' willingness to pay towards the restoration of the corncrake in the Irish countryside. The Payment Card Method (Cameron and Huppert, 1989) was the elicitation format used in the survey. As with any of the response formats in a CV study, the use of this method has advantages and disadvantages. A review of these is beyond the scope of the current paper but the interested reader should see Boyle (2003) for further discussion – details of the CV study are given in Hynes and Hanley (2009). A total of 1,117 surveys were collected. The total number of usable responses was 928.

- Table 1 here

The other dataset used in this paper is the Census of Agriculture. The objective of the census is to identify every operational farm in the country and collect data on agricultural activities undertaken on them (CSO, 2002). The census classifies farms by physical size, type and geographical location. Of the 3,440 EDs in the country, 2,850 contain farms; the average number of farms in each of these EDs being 53 (min 10, max 320). There is substantial evidence within the census and from other sources (Lafferty et al., 1999; Matthews, 2000 and Hennessy, 2004) of a marked regional variation in farm structures and farm income across Ireland. Given this regional variation and the assumption that more intensive farm enterprises may be less willing to pay for corncrake conservation (as a sustainable reintroduction of the bird into parts of its former range could mean a relatively large change in their farm management practices compared to less-intensive farmers), it is important that any

regional aggregation of farmers' WTP for corncrake conservation take into account these substantial differences in farm structures across the country.

4. Methodology

The spatial scale at which sample data sets are released (including those from environmental valuation studies) is often national. These data may not be regionally representative. In such situations one could improve the relation between the sample and the regional population of interest by adjusting the weights of the cases in the sample so that the adjusted weights on specified characteristics agree with the corresponding totals for the regional population. This operation is known as sample balancing, and it offers a number of potential methods to solve the problem of micro-data estimates of WTP not being available at the desired spatial scale. Approaches include balancing factor methods in spatial interaction models (Wilson, 2000), Iterative Proportional Fitting (IPF) techniques (Birkin, 1987; Clarke, 1996; Williamson et al., 1998) and combinatorial optimization techniques (Wu and Wang, 1998; Ballas et al., 2005).

Spatial interaction models are commonly used to model the trade-off between spatial convenience (for example, visiting a retail outlet close by) and the attractiveness of particular outlets (measured by proxies such as size, brand and quality of the service). Balancing factors ensure that the total flow of individuals from a residential zone to all available service centers does not exceed an imposed number of individuals that represents the total demand for the particular service in a given area (Morrissey et al., 2008). IPF on the other hand is a mathematical scaling procedure that ensures that a two-dimensional table of data is adjusted so that its row and column totals are equivalent to row and column totals from some alternative source. IPF acts as a

weighting system whereby the original table values are gradually adjusted through repeated calculations to fit (usually census) row and column constraints (Norman, 1999). The resultant table of data is a joint probability distribution of maximum likelihood estimates obtained when the probabilities are convergent within an acceptable defined limit (Birkin, 1987).

Combinatorial optimization methods on the other hand re-weight an existing micro data sample to fit Small Area Population (SAP) statistics. We employ a combinatorial optimization technique known as simulated annealing in our WTP aggregation exercise. Simulated annealing is used to assign integer weights to each national observation to create a synthetic population for each region. When an observation from the national sample is not represented in region q , according to the matching variables, a weight of zero is assigned. The combinatorial optimization approach has been increasingly adopted to study the spatial impacts of social and economic policies (Ballas et al., 2005). However, this paper is the first application of the methodology to the valuation of public goods.

The use of simulated annealing as an optimization method was first suggested in the early 1980s when Kirkpatrick et al. (1983) discovered an analogy between minimizing the cost function of a combinatorial optimization problem and the slow cooling of a solid until it reaches a low energy state. Since then, simulated annealing has been employed as an optimization technique to solve a variety of combinatorial optimization problems (Mertz, 1991; Wu and Wang, 1998 and Ballas et al., 2005). Goffe et al. (1992) first introduced the method to econometric problems. It has also been used for applications in trading and finance marketing (Ingber, 1996). The main advantage of simulated annealing over the other methods mentioned above is that it

uses survey data on ‘real’ persons to generate small-area population data. The main disadvantage of simulated annealing is long processing time, though this is becoming less of an issue with improvements in computer processing capabilities.

- Table 2 here

In the context of the research presented here we wish to generate synthetic individual farm micro data estimates of individual WTP for each of the 3,440 EDs in Ireland. This implies a combinatorial optimization problem where we try to find the set of NFS farms that can best reproduce the Census of Agriculture SAP tables of farm size, farm system and soil type. These SAP tables simply indicate the number of farms by different size, system and soil type in each ED. The categories within each of these SAP tables are shown in table 2 for a sample of EDs. The SAP tables of farm size, farm system and soil type were chosen as linking variables between the two data sources (farm survey and census) for three main reasons. Firstly, they are believed to be the best descriptors of the regional heterogeneity in the farm population across Ireland, and it is hoped will lead to the most realistic synthetic population of farms in each ED when used as constraints. Secondly, these variables are believed *a priori* to be useful in predicting farmers’ WTP for corncrake conservation, or for being the main explanatory variables that will determine other variables in our valuation function that are themselves key explanatory variables for WTP. Farm system and soil type, for example, would be key variables in determining Rural Environment Protection Scheme (REPS) participation, the quantity of land under crops and pasture (the corncrakes’ main habitat) and the level of organic nitrogen production on a farm, all of which might be expected to effect WTP for corncrake conservation. Finally, in the statistical matching of census to sample, the researcher is limited by the SAP

tables in the census that are made available at the desired spatial level. In our case, the variables in other Census of Agriculture tables available at the ED level would not be considered as reliable for the creation of regional synthetic farm populations using the combinatorial optimization approach, nor for their predictive power in terms of WTP. Neither would the variables in other Census of Agriculture tables (such as number of tractors or type of farm workers per ED) explain much of the variation in the explanatory variables which we may wish to use in the valuation function.

To formalize our combinatorial optimization problem consider a pair (R, Err) , where R is the finite set of farm configurations (set of NFS farm records representing the number of farms in an ED by size, system and soil type) and Err is an error function ($Err : R \rightarrow \mathbf{R}$), which assigns a real number to each farm configuration. Err is defined such that the lower the value, the better the corresponding configuration of NFS farms represents the census SAPs tables. The problem then is to find the configuration of farms for which Err takes its minimum value subject to a computation time constraint and Err being less than some upper bound, i.e. an optimum configuration i_o satisfying:

$$Err_{opt} = Err(i_o) = \min_{i \in R} Err(i),$$

where Err_{opt} denotes the minimum error between the actual census tables of size, system and soil type and the simulated tables constructed using the configuration of NFS farms. In order to solve this combinatorial optimization problem, a simulating annealing (SA) approximation algorithm was employed, which yields an approximate solution in an acceptable amount of computation time (Wu and Wang, 1998). SA is

used to locate a good approximation to the global optimum of a given function in a large search space using randomization techniques. The SA algorithm used in this paper was adopted from the algorithm employed by Ballas and Clarke (2000), where the authors generated a synthetic urban population in Leeds, UK to analysis urban planning issues². The mathematical model underlying the algorithm is described fully in Laarhoven and Aarts (1987, chapter 2).

The process selects a set of farms from the 928 records of the NFS that best fits the census SAP tables for every ED in the country. We initially choose a random configuration i of NFS farms to represent the SAP tables of farms for a single ED. Given configuration i , another configuration j can be obtained by randomly selecting a number of records in configuration i and replacing them with ones chosen at random from the universe of NFS records. The number of records to be replaced at each step is defined as T . In the first step, T equals half the number of farms in the ED. Also, in this first step we set the number of iterations at 5000 (i.e. T farm records are swapped 5000 times in the first step). We then re-tabulate the census SAP tables for the selected set of farms i and calculate the total absolute error or difference from the known small area constraints. The total absolute error between the new tabulation and the actual census tabulation for configuration i is given by $Err(i)$.

$$Err(i) = ErrorSum/CellSum, \text{ where } ErrorSum = \sum_t \sum_r \sum_c |O_t[r][c] - E_t[r][c]|$$

and $CellSum = \sum_t \sum_r \sum_c |O_t[r][c]|$. O_t is the actual census tabulation for each SAP table t (ranging from 1 to T) and E_t is the estimated tabulation using configuration i (derived from the NFS sample) for each SAP table. Also, r is the row number in each table, ranging from 1- R_t and c is the column number in each table,

ranging from $1 - C_t$. $Err(i)$ is chi-squared distributed with degrees of freedom equal to $\sum_t [R_t][C_t] - 1$. Letting $\Delta Err_{ij} = Err(j) - Err(i)$, where once again $Err(*)$ is the error between the tabulation of each configuration and the census tables, then the probability that configuration j will be the next configuration of farms in a predefined sequence of configurations is given by 1, if $\Delta Err_{ij} < 0$ and by $\exp(-\Delta Err_{ij}/T)$ if $\Delta Err_{ij} > 0$. The acceptance of a new configuration when $\Delta Err_{ij} > 0$ is decided by drawing random numbers from a uniform distribution on $[0, 1]$ and comparing these with $\exp(-\Delta Err_{ij}/T)$. If $\exp(-\Delta Err_{ij}/T)$ is greater than the drawn random number the new configuration is accepted even though $Err(*)$ is larger than for the previous configuration. These uphill movements prevent the process getting trapped in local minima.

This process continues, with T being lowered at each step. At each step, T is reduced by a set percentage written in to the Java program (in our case 10%). The number of iterations at each step is inversely proportional to T , so that as the number of farms per swap is reduced, the number of iterations is increased. In our program, we increased the iterations by an increment of 1500 at each step³. As T is lowered fewer uphill moves are accepted because the value of $\exp(-\Delta Err_{ij}/T)$ is a positive function of T . Eventually, the number of farms per swap is reduced to 1. The process is complete when either the maximum number of iterations has been hit or the total absolute error falls within the desired setting⁴.

- Figure 1 here

To demonstrate the above process, figure 1 presents a graph of the three types of error in the combinatorial optimization process for a sample set of iterations. The line of diamonds represents the error from the current configuration of farms. The line containing the triangles shows the best error. The best error refers to the lowest total absolute error found for any configuration of farms since the optimization process began. For example, if in iteration 1 the total absolute error between the simulated census tables and the actual census tables was found to be 67 and in iteration 2 it was found to be 52 then the configuration of farms associated with iteration 2 will be considered the configuration with the best error to date. If, on the other hand, the error between the simulated census tables and the actual census tables was found to be 85 in iteration 2 then the configuration of farms associated with iteration 1 will continue to be the configuration with the best error. The dashed line of squares shows the accepted error, which equals the current error if the swap is accepted and the best error if it is not accepted. Points A and B are iterations in which the farm swaps were accepted even though their error exceeded the best error. In the iteration after point B, the best error falls in line with the accepted error showing that sometimes a sub-optimal swap or an uphill movement can lead to a decrease in the overall error.

The combinatorial optimization process is complete when the selected configuration of farms from the NFS can reproduce the census SAP tables with less than 5% of a difference between the original SAP tables and those generated from the NFS selection. Once this point is reached the program stores the simulated configuration of NFS farm records for that ED and repeats the process to find the configuration of NFS farms that best fits the census SAP tables for the next ED and so on⁵. Matching the NFS and the SAP data creates synthetic demographic, socio-economic and farm level variables, such as age, fertilizer usage, livestock units, etc, and most importantly

from our research perspective, predicted WTP values *for each farmer in the population*. The combinatorial optimization process conducted for this research produces 145,057 individual farm records with their associated maximum WTP variable. This is clearly a big expansion of the initial sample of $n = 928$.

5. Results

Chi-square distributed statistics are used to assess how well the simulated SAP results compare to the actual census tables. The synthetic micro data estimates produced by the spatial microsimulation model are also validated by re-aggregating the model results up to levels at which observed data sets exist (Irish Central Statistics Office (CSO) figures) and then comparing the estimated distributions with the observed. Table 1 presents a comparison of summary statistics for both the NFS and our micro-simulated population. Both the chi-square distributed statistics and the comparison of re-aggregated results (county level) to CSO figures indicate a high level of “goodness of fit” for the simulated population. The validation of the optimization results using the chi-square distributed statistics and the re-aggregation comparisons are not reproduced here but are discussed fully in Hynes et al. (2009).

The main goal of the combinatorial optimization exercise carried out in this paper was to calculate the aggregate value of the proposed corncrake conservation project to the Irish farming community at different levels of spatial aggregation. In order to compare our simulated estimate of aggregate WTP to other more traditional approaches of aggregation used in the literature we calculate the aggregate environmental value of the corncrake conservation program in three alternative ways. These are:

1. **WTP_{NFS}**: Aggregation using a CV generalized Tobit model (the CV function approach) where the estimated average value of WTP in the NFS sample is multiplied by the number of farms in the country (if interested in national aggregation) or the number of farms at the desired level of aggregation (county, ED, province, etc) $\overline{(nWTP_{NFS})}$.
2. **WTP_{NATIONAL SIM}**: The coefficients from the NFS generalized Tobit are applied to the simulated farm population and the resulting estimated values of WTP in the synthetic population for each farm i are summed to calculate total WTP at the desired level of aggregation $(\sum_{i=1}^n \hat{WTP}_{SIMi})$ and
3. **WTP_{REGIONAL SIM}**: Aggregation using a CV generalized Tobit model estimated using the synthetic population data at the required spatial level where the resulting estimated values of WTP in the synthetic population for each farm i are summed to calculate total WTP in the population $(\sum_{i=1}^n \hat{WTP}_{REGIONAL_SIMi})$. Approach (1) can be viewed as a conventional method of aggregation in CV. Approaches (2) and (3) rely on micro-simulation.

The parametric CV regression calculated using the 2006 NFS sample (weighted using the individual farm population weights provided in the NFS) and a description of the associated censorship points are presented in table 3. As previously mentioned, the elicitation format chosen in this study was the Payment Card Method, where each farmer was shown a payment card listing various Euro amounts and asked to indicate the maximum amount they were willing to pay. Following Cameron and Huppert (1989), the response is interpreted not as an exact statement of willingness to pay but rather as an indication that the WTP lies somewhere between the chosen value and

the next largest value on the payment card. Since WTP is a censored distribution in this case, the appropriate foundation from which to develop the estimation procedure is to use a censored regression or Tobit model (Greene, 2000). However, for an accurate treatment of the WTP variable we need to adapt the estimation procedure in order to account for the mixture of point, interval and censored observations. Therefore, a generalized Tobit model was employed⁶.

- Table 3 here

A generalized Tobit subsumes the Tobit or censored regression model by employing a log-likelihood function adjusted to make provision for point, left-censored, right-censored and interval data. For farmers $j \in C$, we observe WTP_j , i.e. point data and for farmers $j \in L$, WTP_j are left censored. In our case we have no left censored observations as it is assumed that WTP cannot be less than zero, i.e. our survey design rules out people registering negative WTP. Where applicable, zero is therefore treated as a point rather than a left censored observation⁷. Farmers $j \in R$ are right censored; we know only that the unobserved WTP_j is greater than or equal to WTP_{Rj} . Finally farmers $j \in I$ are intervals; we know only that the unobserved WTP_j is in the interval $[WTP_{1j}, WTP_{2j}]$. Thus, regardless of the types of observations, the estimation method is able to account for them simultaneously.

It can be seen from the model results that WTP increases significantly with the income generated on the farm. The “REPS farm” variable indicates that farmers participating in the REPS⁸ are willing to pay (significantly) higher amounts than those farmers not participating in the scheme. Given the environmental education

component involved in the uptake of this scheme and the fact that farmers participating in an agri-environmental scheme are more likely to favor a biodiversity conservation program, this finding is not surprising. The Organic Nitrogen Production per hectare variable is an indicator of the intensiveness of the farming enterprise. Farms with higher rates of organic nitrogen per hectare are found to be willing to pay significantly less for a corncrake conservation program, possibly indicating a higher opportunity cost of conservation actions. As can be seen from the second column in table 3, the parametric CV regression estimated using the entire simulated farm population produces coefficients of the same sign and of a similar magnitude as the NFS regression (except for the ‘Total crops and pasture as fraction of Farm Size’ variable which is -8.24 compared to -4.97 in the NFS model).

- Table 4 here

Table 4 presents average and total WTP values estimated using the three alternative valuation approaches. The NFS generalized Tobit model produces a similar value for the average WTP per farm to the value generated from applying the model coefficients to the simulated farm data but a significantly (at the 95% level) higher value to that generated from re-running the model on the simulated data (although in absolute terms it still only differs by approximately €1, €10.65 compared to €9.58). In relation to the aggregation of the WTP values it can be seen from table 4 that at the national level of aggregation⁹ the first two approaches produce estimates that are similar in magnitude, but once again **WTP_{REGIONAL SIM}** produces an aggregate figure that is significantly lower. The larger negative coefficient on the ‘Total crops and pasture as fraction of Farm Size’ variable in the generalized Tobit estimated using the entire simulated farm population would appear to be driving this differential. It would

appear that the weight given to land under crops and pasture in the simulated data is different to that which is actually the case in the weighted NFS sample.

- Figure 2 here

As can be seen from figures 2 and 3, the aggregate value of the proposed corncrake conservation project to the Irish farming community can be examined at a number of different levels of spatial aggregation using our simulated population results. These include ED, county and regional level. It is evident from the maps (figures 2 and 3), produced with our simulated farm population estimates of WTP (aggregation approach 2), that farmers in EDs found in the west, south west and border areas of the country seem to be willing to pay higher amounts on average to have the corncrake restored into the Irish countryside. This is an interesting finding given that the remaining singing male population of corncrakes in Ireland is largely restricted to four areas (as shown in figure 2): Co. Fermanagh (which is on the border on the Northern Ireland side), Donegal, West Connacht, and the Shannon Callows (Schäffer and Green, 2001).

The positive spatial correlation between the WTP values in our simulated population of farmers and the areas where the corncrake can still be found highlights what Bateman (2006) refers to as an ‘ownership’ dimension to aggregate benefit values. Of course, as previously mentioned, these areas are also associated with the smaller, less intensive dry stock systems of farming where costs of conserving the corncrake are relatively low, and it may be this fact that is driving the observed spatial distribution of average WTP per ED across the country.

- **Figure 3 here**

Similar to the water quality example discussed in the introduction and as demonstrated in figures 2 and 3, the real strength of the combinatorial optimization approach is the fact that we can examine the aggregate WTP of particular regions within the country while taking into account the regional variation in farming activity. To this end we choose to examine two extensive farming counties and two intensive farming counties as defined in the Census Atlas of Agriculture in the Republic of Ireland (Lafferty et al., 1999)¹⁰. The results of this regional aggregation using the alternative aggregation approaches are displayed and compared in table 5. In terms of the regional aggregation of the estimated WTP values, $WTP_{NATIONAL\ SIM}$ produces regional aggregate estimates that are significantly lower (at the 95% level) for the *intensive* farming counties compared to WTP_{NFS} . On the other hand, for the *extensive* farm counties, $WTP_{NATIONAL\ SIM}$ produces regional aggregate estimates that are significantly higher compared to WTP_{NFS} . This is an interesting result considering that nationally average $WTP_{NATIONAL\ SIM}$ is less than average WTP_{NFS} and yet we still find that for the extensive farming counties analyzed aggregate $WTP_{NATIONAL\ SIM}$ is greater than aggregate WTP_{NFS} .

- **Table 5 here**

In these aggregation approaches we have made the assumption that the preferences of farmers are stable across space¹¹. However, one of the main advantages of deriving the synthetic regional data is that it enables different preferences (the estimation of different parameter values in the WTP equation) across regions. Instead of applying the NFS model coefficients on the synthetic data we can derive the regional WTP

estimates by running separate WTP regressions based on the synthetic regional data sets. Tests can also be conducted to compare these WTP regressions across regions. When we estimate these regional models we find (table 5) that the pattern we saw previously is even stronger.

Compared to the traditional aggregation approach, WTP_{NFS} , the total WTP estimates for the counties estimated using the separate county CV functions ($WTP_{REGIONAL\ SIM}$) are significantly higher for extensive farming counties and lower for intensive counties. A likelihood ratio (LR) test was also performed to test if the estimated regional models were significantly different from the national level model. The test statistic indicates that at any reasonable level of significance we can reject the hypothesis that the national level model applies to any of the 4 regions. The results demonstrate that by not recognizing the differences in farming activity across Irish rural space through the use of the simulated population, we would be overestimating the regional aggregation of WTP for intensively farmed areas and underestimating the regional aggregation of WTP for extensively farmed areas using the standard benefit transfer approach. These biased estimates could lead to an inefficient allocation of resources across these regions.

6. Discussion and Conclusions

As discussed in section 2, there are numerous examples in the literature where reliance upon sample means will fail to yield an accurate measure of aggregate WTP. As an alternative, we propose a new approach based upon combinatorial optimization which takes into account the impact of variation in the characteristics of the relevant population over which benefits are to be aggregated, and which allows for the calculation of benefits at a wide variety of spatial scales. The comparison between

aggregate estimates in table 4 (for national total WTP) and table 5 (for county total WTP) demonstrated that the combinatorial optimization approach provides similar estimates of aggregate environmental value as the simple sample mean WTP aggregation approach at the national level when transferring the sample model coefficients across the simulated population, but resulted in regional values which were significantly different when assessing total WTP values at the county level. We also saw that the combinatorial optimization approach leads to considerable gains in measuring the variation in aggregate WTP across types of individual (here, between different types of farms). The results ultimately demonstrate that researchers failing to take account of the spatial heterogeneity in their study population may be introducing biases in their attempts to estimate the spatial distribution of aggregate environmental benefit values.

For our data we found that the **WTP_{REGIONAL SIM}** approach produced significantly different national estimates of average and aggregate WTP compared to the other 2 approaches. The different weighting given to the *land under crops and pasture* variable in the simulated data resulted in a much larger value for this coefficient in the parametric CV regression estimated using the entire simulated farm population. This suggests that it may be more appropriate, in our case, to use the regional WTP estimates that result from the transfer of the coefficients from the NFS model across the simulated regional population rather than those got from re-estimating on the simulated regional population. Alternatively, it could be possible to calibrate the *land under crops and pasture* variable in the simulated data to match known county level totals. Ideally, using just the combinatorial optimization constraining table variables (in our case, size, system and soil type) as the only explanatory variables in estimating the CV functions should produce the most comparable regional WTP

estimates across all approaches. However, given census data limitations this option may not be available nor produce the best fitting model.

Looking beyond the detail of the example reported here, we would speculate that there are a number of benefits of using combinatorial optimization methods to create synthetic micro data. Firstly, using the spatial microsimulation modeling framework allows us to efficiently and accurately aggregate a sample's individual welfare estimates to a particular region (or spatial scale) of interest. Secondly, the methodology allows us to take into account the spatial heterogeneity of the target population in the aggregation process, to the extent that this can be linked to observables. Thirdly, when adequate funds or time are not available to conduct a large-scale non-market valuation survey across diverse regions, the combinatorial optimization method could be cost effectively employed to produce regional valuation estimates.

Finally, it should also be noted that the creation of the synthetic population of farmers with their accompanying WTP values is not a technique that is unique to the datasets in this study. In theory, it is an approach that could be replicated and used with sample data sets in other CV studies and with revealed preference techniques such as the recreation travel cost method and the hedonic price valuation technique. Also, a number of SA algorithms are now available to download free from the internet¹² and as Hynes et al. (2009) point out, once the matching datasets are structured in a manner that can be read by the programming language being employed, the synthetic data can be produced without to much "reinventing of the wheel" on the part of the researcher.

It should also be pointed out that for logistical reasons, non-market valuation surveys are often conducted within a smaller region and researchers may try to draw inferences about a population outside their sampling frame. The method developed in this paper could be used to improve aggregation estimation in this other direction (“broadening out” rather than “focusing down”). In principle, one could use WTP data which had been collected in a particular region and match it to census data to generate a simulated national dataset which includes simulated values for WTP. This matching could be done using variables such as age, sex or ethnicity, which vary within the regional survey. The combinatorial optimization algorithm would reweight the regional sample to produce a nationally representative population and a nationally representative estimate of aggregate environmental benefit.

A revealed preference example of where the combinatorial optimization may be highly applicable is in aggregating estimated benefits in recreational travel cost studies. It can be difficult to calculate the aggregate non-market value of a recreational site due to lack of information on the total population that might use that site on an annual basis (Hynes and Hanley, 2006). By using distance decay functions in a spatial micro-simulation framework this obstacle could be overcome. The approach could also be used in a choice experiment setting. For example, choice experiment data could be used to produce a spatial micro-simulation model to estimate the economic value of improvements in the ecology of water bodies which is required under the European Union’s Water Framework Directive. This would echo work by Lewis and Plantinga (2007), where the authors use a GIS-based approach to integrate an econometric model of land use with simulations that predict the spatial pattern of land-use change, in order to examine how the costs of reducing habitat fragmentation varied with landscape conditions.

One further area for future research in terms of using the combinatorial optimization techniques in public good valuation studies would be a “ground truthing” exercise to examine whether the WTP estimates in the micro-simulated population are statistically equivalent to what one would find if the WTP questionnaire was to be conducted in each ED. To this end it would be worthwhile to pick a number of EDs around the country and survey the farmers in them to see how close the actual WTP values of the farmers in the chosen EDs are to the estimates for those farmers in the corresponding simulated ED populations. Given that our simulated population is constrained to statistically match the Census of Agriculture tables, and the fact that our aggregation method meets the criteria discussed in the literature for the production of more reliable aggregate welfare estimates, (Bateman et al., 2006; Rosenberger and Stanley, 2006)¹³ we would expect the aggregation error to be relatively small.

It has been claimed that environmental policy has to be region-specific in the light of distributional issues and site and population heterogeneity (van Pelt, 1993). The results of this paper support that viewpoint. Similarly, Nijkamp (2002) contended that progress at the interface of regional and environmental economics is contingent on the availability of proper spatial information systems and models. We believe that we have offered a new perspective for analyzing this linkage between space and the environment, where land use and heterogeneous spatial behavior are shown to be closely connected to alternative regional aggregate environmental benefit values.

References

- Ballas, D. and G. Clarke. 2000. "GIS and Microsimulation for Local Labour Market Analysis." *Computers, Environment and Urban Systems* 24: 305-330.
- Ballas, D., G. Clarke and D. Dorling. 2005. *SimBritain: A Spatial Microsimulation Approach to Population Dynamics*, Rowntree Publishers, York.
- Bateman, I., B. Day, S. Georgiou, and I. Lake. 2006. "The Aggregation of Environmental Benefit Values: Welfare Measures, Distance Decay and Total WTP." *Ecological Economics*, 60: 450–460.
- Birkin, M. 1987. *Iterative Proportional Fitting (IPF): Theory, Method and Example*. Computer Manual 26, School of Geography, University of Leeds, Leeds.
- Boyle, K. 2003. "Contingent Valuation in Practice", in Champ, P., Boyle, K., and Brown, T. *A Primer on Nonmarket Valuation*. Kluwer Academic Publishers.
- Cameron, T., and D. Huppert. 1989. "OLS versus ML Estimation of Non-Market Resource Values with Payment Card Interval Data." *Journal of Environmental Economics and Management*, 17: 230-246.
- Central Statistics Office (CSO). 2002. *The Census of Agriculture - Main Results*. Central Statistics Office, Cork.
- Clarke, G. 1996. *Microsimulation for Urban and Regional Policy Analysis*. Pion, London.
- Greene, W. 2003. *Econometric Analysis*: Fifth Edition, Prentice Hall, New Jersey.
- Goffe, W., G. Ferrier and J. Rogers. 1992. "Global Optimization of Statistical Functions: Preliminary Results." In Amman, H., Belsley, D. and Pau, L. *Computational Economics and Econometrics*, 19–32.
- Hanley, N., F. Schläpfer and J. Spurgeon. 2003. "Aggregating the Benefits of Environmental Improvements: Distance-Decay Functions for Use and Non-Use Values." *Journal of Environmental Management* 68: 297–304.

- Hennessy, T. 2004. *Projecting Farm Numbers*. 2015 Agri-Vision Report. Irish Department of Agriculture and Food, Dublin.
- Hynes, S. and Hanley, N. 2006. "Preservation versus Development on Irish Rivers: Whitewater Kayaking and Hydro Power in Ireland." *Land Use Policy* 23: 170 - 180.
- Hynes, S and N. Hanley. 2009. "The "Crex crex" Lament: Estimating Landowners Willingness to Pay for Corncrake Conservation on Irish Farmland". *Biological Conservation* 142: 180-188.
- Hynes, S., K. Morrissey, and C. O'Donoghue. 2009. "Building a Static Farm Level Spatial Microsimulation Model for Rural Development and Agricultural Policy Analysis in Ireland." *International Journal of Agricultural Resources, Governance and Ecology* 8: 282 - 299.
- Ingber, L. 1996. "Statistical Mechanics of Nonlinear Non-Equilibrium Financial Markets: Applications to Optimized Trading." *Mathematical Computer Modelling*, 23: 101–121.
- Kirkpatrick, S., C. Gelatt, and M. Vecchi, 1983. "Optimization by Simulated Annealing." *Science*, 220: 671–680.
- Laarhoven, van, P. and Aarts, E. 1987. *Simulating Annealing: Theory and Applications*. Kluwer Academic Publishers, London.
- Lafferty, S., P. Commins, and J. Walsh. 1999. *A Census Atlas of Agriculture in the Republic of Ireland*. Dublin: Teagasc Publication.
- Lewis, D. and A. Plantinga. 2007. "Policies for Habitat Fragmentation: Combining Econometrics with GIS-Based Landscape Simulations." *Land Economics* 83 (2): 109-127.

- Loomis, J. 1987. Expanding Contingent Value Sample Estimates to Aggregate Benefit Estimates: Current Practices and Proposed Solutions. *Land Economics* 63: 396-402.
- Matthews, A. 2000. *Farm Incomes: Myths and Reality*, Cork University Press, Cork.
- Mertz, J. 1991. "Microsimulation - A survey of principles, developments and applications." *International Journal of Forecasting* 7: 77-104.
- Morrison, M. 2000. "Aggregation Biases in Stated Preference Studies." *Australian Economic Papers*, 39 (2), 215–230.
- Morrissey, K., G. Clarke, D. Ballas, S. Hynes and O'Donoghue, C. 2008. "Analysing Access to GP Services in Rural Ireland using Microsimulation Analysis". *Area* (40), 354-364.
- National Parks & Wildlife Service (NPWS), Department of the Environment, Heritage and Local Government Ireland, 2005. *All-Ireland Species Action Plans for the Irish Hare, the Corncrake, the Pollan and Irish Lady's Tresses* (www.npws.ie).
- Nijkamp, P. 2002. "Space in Environmental Economics" in J. van den Bergh (ed.), *Handbook of Environmental and Resource Economics*, Edward Elgar Publishers, Cheltenham, UK.
- Norman, P. 1999. "Putting Iterative Proportional Fitting (IPF) on the researcher's desk". Working Paper 99/03. School of Geography, University of Leeds, Leeds.
- Pelt, van, M., 1993. *Ecological Sustainability and Project Appraisal*. Aldershot, UK: Avebury.
- Rosenberger, R. and T. Stanley. 2006. "Measurement, Generalization and Publication: Sources of Error in Benefit Transfers and their Management." *Ecological Economics* 12 (2):372-378.

- Schäffer, N. and R. Green, 2001. "The Global Status of the Corncrake." *RSPB Conservation Review* 13: 18-24.
- Smith, V. 1993. "Nonmarket Valuation of Environmental Resources: An Interpretive Appraisal." *Land Economics* 69 (1), 1-26.
- Williamson, P., M. Birkin, and P. Rees. 1998. "The Estimation of Population Microdata by Using Data from Small Area Statistics and Samples of Anonymised Records." *Environment and Planning A* 30: 785-816.
- Wilson, A. 2000. *Complex Spatial Systems: the modelling foundations of Urban and Regional Analysis*, Prentice Hall, London.
- Wu L. and Y. Wang. 1998. "An Introduction to Simulated Annealing Algorithms for Computation of Economic Equilibrium." *Computational Economics* 12: 151-169.

Tables

TABLE 1. SUMMARY STATISTICS OF THE NFS AND MICROSIMULATED FARM POPULATIONS

	National Farm Survey Sample 928 Observations		Microsimulated Farm Population 145,057 Observations	
<i>Variable</i>	<i>Mean</i>	<i>Standard Deviation</i>	<i>Mean</i>	<i>Standard Deviation</i>
Size of Farm (acre)	90.54	78.04	76.42	64.91
Crop Pasture (acre)	82.05	68.89	72.53	61.13
Gross margin (€)	38653.09	39785.29	35039.79	37645.17
Farm income (€)	22387.14	24402.69	20026.95	22417.42
Gross output (€)	55315.96	58365.37	50421.83	54912.12
REPS payment (€)	2420.78	3388.42	1892.79	2959.51
Age (years)	54.30	12.72	54.34	12.83
Max stated WTP (€)	7.18	13.39	6.79	12.19

TABLE 2. CENSUS SAP TABLES FOR A SAMPLE OF EDs

<i>Size</i>	<10	10-20	20-30	30-50	50-100	>=100	Total Farms
Oranmore	11	11	11	22	11	0	66
Kinvara	10	20	10	20	10	0	70
<i>System</i>	Tillage	Dairying	Beef	Sheep	Mixed Livestock	Crops and Livestock	Total Farms
Oranmore	0	0	9	19	28	10	66
Kinvara	10	0	10	20	20	10	70
<i>Soil Class</i>	Class 1	Class 2	Class 3	Class 4	Class 5	Class 6	Total Farms
Oranmore	0	55	1	9	1	0	66
Kinvara	0	40	22	8	0	0	70

The soil classes refer to the range of potential uses that the dominant soil type on a farm can be put to, where 1 indicates that soils are of wide use range that have no limitations which cannot be overcome by normal management practices and 6 refers to soils where the agricultural potential is virtually non-existent. Farm size is measured in hectares.

TABLE 3. GENERALIZED TOBIT OF WTP FOR CORNCRAKE CONSERVATION

Variable	NFS Model	Simulated Farm Model
Size of Farm (acres)	-0.043 (-0.02)**	-0.045 (-0.002)***
Family Farm Income (€/1000)	0.075 (0.03)***	0.040 (0.002)***
Age of Farm Operator	0.027 (0.05)	0.104 (0.003)***
Organic Nitrogen Production (kg/hectare)	-0.036 (-0.02)**	-0.020 (-0.001)***
REPS farm ^a	2.072 (1.26)*	3.231 (0.07)***
Total crops and pasture as fraction of Farm Size	-4.97 (-4.85)	-8.244 (-0.32)***
Constant	14.73(5.31)***	13.07 (0.35)***
Log of the estimated standard error	2.722 (0.033)***	2.56 (0.002)***
Sample Size	928	145057
Log likelihood	-274263	-381357
Likelihood Ratio χ^2 (6) test	18.55	6126
Left Censored Observations	0	0
Right Censored Observations	4	532
Uncensored Observations	538	58950
Interval Observations	386	85575

Robust standard error in parentheses.* significant at 10%; ** significant at 5%; *** significant at 1%. (a) REPS farm indicates that the farmer participates in the Rural Environment Protection scheme in the reference year, 2006.

TABLE 4. WTP ESTIMATES FOR THE 4 ALTERNATIVE ESTIMATION METHODS (€)

Method of Analysis	Average WTP Per Farm (€)	Total environmental value of a corncrake conservation program (€)
Payment Card generalized Tobit for NFS sample	10.65 (10.45, 10.82)	1,544,857
NFS Payment Card generalized Tobit applied to simulated farm population	10.40 (10.38, 10.42)	1,508,592
Payment Card generalized Tobit estimated using simulated farm population	9.58 (9.56, 9.59)	1,388,195

95% confidence Intervals in brackets

TABLE 5. TOTAL WTP ESTIMATES PER COUNTY FOR THE 4 ALTERNATIVE AGGREGATION METHODS (€)

County	WTP _{NFS}	WTP _{NATIONAL SIM}	WTP _{REGIONAL SIM}
<i>Extensive Counties</i>			
Galway	141,293	145,010	151,238
Donegal	89,939	91,996	98814
<i>Intensive Counties</i>			
Tipperary S	40,289	37,170	33,616
Cork	162,572	149,369	135,506

Note that n , the relevant aggregation figure, in this case represents the total number of farms in each county.

Figure 1. Sample of Errors by Iteration


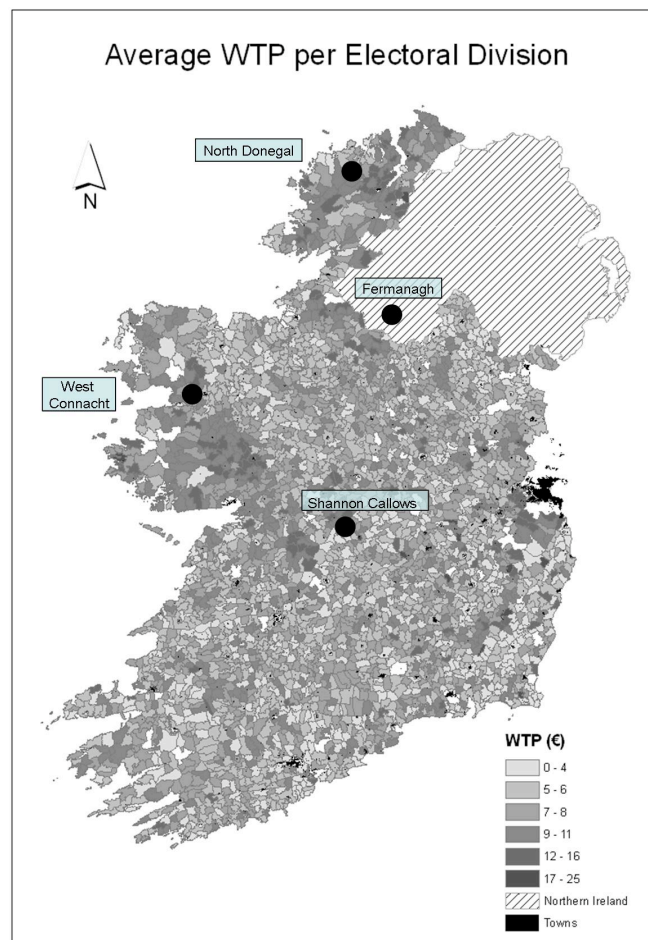
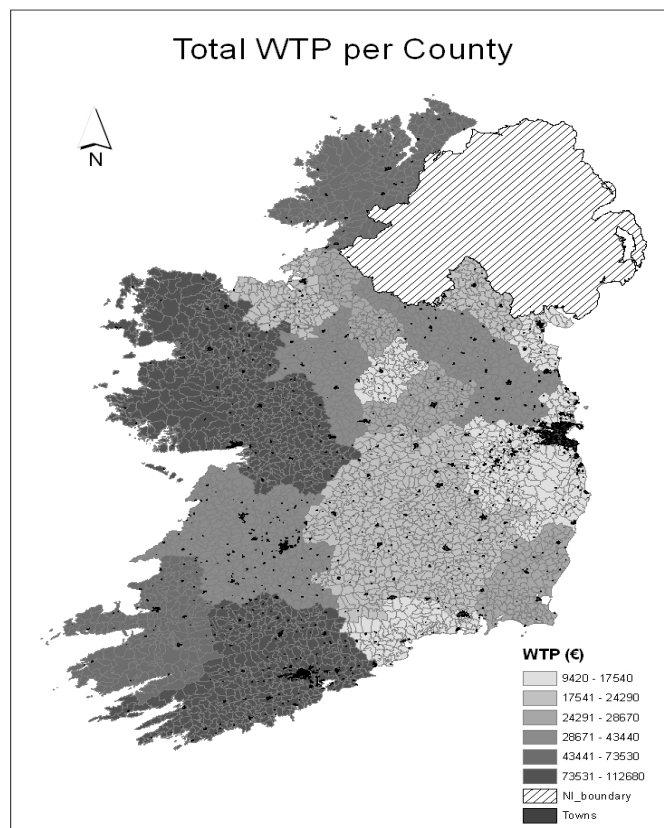


Figure 2. Average WTP for a Corncrake Conservation Program at the ED Level of Spatial Aggregation



The four labeled dots indicate the remaining breeding grounds of the corncrake in Ireland. Parcels that are blank indicate no farm activity in those Electoral Divisions. Of the 3,440 EDs in the country, 2,850 contain farms.

Figure 3. Total WTP for a Corncrake Conservation Program at the County Level of Spatial Aggregation



Footnotes

¹ The WTP study employed in this paper is fully discussed in Hynes and Hanley (2009).

² We implement the SA algorithm in Java, an object-oriented programming language, which has been accepted as the most suitable type of programming language for spatial microsimulation modelling (Ballas and Clarke, 2000; Wu and Wang, 1998).

³ The number of iterations in the first step, the iteration increment increase thereafter, the number of records to be swapped in the first step and the percentage reduction in the number of records to be swapped at each step thereafter are all set at the discretion of the programmer.

⁴ The static model also employs a restart method. When a restart occurs the simulated annealing process begins again with a new sample of records. The restart is used so that more farm combinations can be explored. The restart method is applied if the model fails to find a satisfactory solution within the maximum number of permitted iterations.

⁵ It is important to realize that a single farm from the CV survey may appear multiple times in the simulated synthetic population of a single ED and could potentially appear in numerous EDs in the simulated population. This happens for the simple reason that we are only using the sample of 928 records to produce a simulated national farm population of 145,000. It should also be noted that the size of the matching sample can significantly influence the outcome of the synthetic data generation process. In our application, the NFS sample consists of less than 1 percent of the total farm population. The results presented in the article therefore, must be viewed as conditional upon the relatively small sample. Nevertheless, the sample is of sufficient size to demonstrate the advantages of using the combinatorial optimization methodology.

⁶ The interested reader is advised to look at Hynes and Hanley (2009) for an in-depth discussion of this contingent valuation study. To save space we only briefly discuss the results here and instead dedicate

ourselves to the comparison of the aggregation results using our combinatorial optimization methodology and more traditional aggregation approaches.

⁷ Although in our analysis we assume that there are no people with negative WTP it may be the case that some farmers would be willing to pay to have fewer or no corncrakes in the countryside. To this extent, it could be argued that zeros should be treated as left censored observations.

⁸ The Rural Environment Protection Scheme (REPS) was introduced in Ireland under EU Council Regulation 2078/92 in order to encourage farmers to carry out their activities in a more extensive and environmentally friendly manner. Approximately 43,000 farmers were actively participating in the scheme in 2006.

⁹ For the national aggregation, n , the total number of farms is equal to 145,057 (CSO, 2002)

¹⁰ An intensive farming county is defined as one where the average annual rate of organic nitrogen production per hectare is greater than 170kg while for an extensive farming county, it is less than 170kg. Intensive farming counties tend to be dominated by larger sized tillage and dairy farm enterprises while extensive farm counties would be characterised by smaller dry-stock systems.

¹¹ If we assume the simulated annealing technique create the synthetic population to closely mimic the regional characteristics of the explanatory variables in our CV function, then the aggregate WTP based on the synthetic regional data should be similar to the one we get by substituting the regional means of the explanatory variables into the national model then multiplying it by the number of farms in the region; i.e., the estimated total WTP from the synthetic data can be similarly derived by what is referred to in the literature as benefit function transfer. This however assumes that we know the means of the explanatory variables for the regions. While we can perhaps get this average regional information for size and system of farm from the census, the acquisition of the means of other important explanatory variables such as organic nitrogen production, total crops and pasture and of being a REPS farmer may be a much more difficult proposition. The advantage of the simulated annealing process is that it generates this regional information for us.

¹² See for example <http://webscripts.softpedia.com/script/Miscellaneous/General-simulated-annealing-algorithm-34475.html>.

¹³ The criteria that may affect the accuracy of aggregation include the quality of the population sample data, the methods used in modeling and interpreting the sample data, the analysts' judgments regarding research design and implementation and the closeness between the sample and the relevant population.