Cost sharing and free allocation in a nutrient trading system with heterogeneous farms: A model of the Lake Rotorua catchment

L. Timar
S. Kerr
S. Anastasiadis

GNS Science Report 2014/27
June 2015

BIBLIOGRAPHIC REFERENCE

Timar, L.; Anastasiadis, S.; Kerr, S. 2014. Cost sharing and free allocation in a nutrient trading system with heterogeneous farms: A model of the Lake Rotorua catchment, *GNS Science Report* 2014/27. 33 p.

Levente Timar, GNS Science, PO Box 30-368, Lower Hutt 5040, New Zealand/ Motu Economic and Public Policy Research, PO Box 24390, Wellington, New Zealand

Simon Anastasiadis, Stanford University, 450 Serra Mall, Stanford, CA 94305, United States

Suzi Kerr, Motu Economic and Public Policy Research, PO Box 24390, Wellington, New Zealand

ISSN 1177-2425 (Print)
ISSN 2350-3424 (Online)
ISBN 978-1-92728-51-2 (Print)
ISBN 978-0-478-19930-0 (Online)

CONTENTS

ABSTRACT ..... **III**

KEYWORDS ..... **III**

EXECUTIVE SUMMARY ..... **IV**

1.0 INTRODUCTION **1**

2.0 THE NMANAGER MODEL **2**

2.1 MODEL STRUCTURE 2
2.2 KEY ASSUMPTIONS 3

3.0 DATA PREPARATION **4**

4.0 METHODOLOGY **7**

4.1 PARAMETERISATION OF FARM PROFIT FUNCTIONS 7
4.2 IMPLEMENTING THE NUTRIENT CAP 11
4.3 THE ALLOCATION OF ALLOWANCES 12

5.0 RESULTS AND DISCUSSION **13**

5.1 AGGREGATE RESULTS 13
5.2 DAIRY SECTOR RESULTS 17
5.3 DRY STOCK SECTOR RESULTS 25

6.0 CONCLUSION **30**

7.0 ACKNOWLEDGEMENTS **33**

8.0 REFERENCES **33**

FIGURES

Figure 4.1 Illustrative quadratic profit functions. 8
Figure 4.2 The nutrient cap and its implementation. 11
Figure 5.1 The simulated market price of allowances and the perpetuity value of a permanent right to discharge (calculated at time zero). 14
Figure 5.2 Mitigation per unit area across sectors (defined by initial land use). 14
Figure 5.3 The distribution of baseline manageable nutrient loss among dairy farms. 17
Figure 5.4 The final distribution of allocation impacts among dairy farms. 18
Figure 5.5 The final (year 20) distribution of manageable nutrient loss among dairy farms. 19
Figure 5.6 The final distribution of mitigation costs among dairy farms. 20
Figure 5.7 The final distribution of total costs with no free allocation. 21
Figure 5.8 The final distribution of total costs with free allocation. 22
Figure 5.9 The relationship between baseline nutrient loss and final mitigation cost among dairy farms. 22
Figure 5.10 The relationship between baseline nutrient loss and total policy cost among dairy farms with grandparenting (GP) and sector-based averaging (SA). 23

Figure 5.11 The variation of allocation impacts with milk production. 24
Figure 5.12 The variation of allocation impacts with mean annual rainfall. 24
Figure 5.13 The distribution of baseline manageable nutrient loss among dry stock farms. 25
Figure 5.14 The final distribution of allocation impacts among dry stock farms. 26
Figure 5.15 The final distribution of manageable nutrient loss among dry stock farms. 26
Figure 5.16 The final distribution of mitigation costs among dry stock farms. 27
Figure 5.17 The final distribution of total costs with no free allocation. 28
Figure 5.18 The final distribution of total costs with free allocation. 28

TABLES

Table 3.1 The classification of Overseer blocks........ 5
Table 3.2 The sources and amount of annual nitrogen loss 6
Table 5.1 Sector-level simulation results 15

ABSTRACT

We use a nutrient trading simulation model to explore the incidence of costs across heterogeneous farm properties after mitigation actions and trade of allowances. We compare two approaches to the free allocation of nutrient discharge allowances: a grandparenting approach and a sector-based averaging approach. We parameterise the model to observations of farms in the Lake Rotorua catchment. Meeting the final nutrient target of 256 tonnes N/year requires significant land-use and land-management change throughout the catchment. Within each sector, farmers with relatively high baseline nutrient loss are better off under grandparenting, and farmers with relatively low baseline nutrient loss are better off under the sector-based averaging approach. However, final cost sharing also depends on the farmer's ability to mitigate and may not be obvious from direct consideration of the free allocation mechanism. Modelling allocation in the context of a simulated trading system allows us to consider the scale of associated fairness issues in each sector. We also discuss how the two allocation approaches satisfy potential principles for cost sharing.

JEL codes

Q52, Q53, Q58

KEYWORDS

Water quality, nutrient trading, allowance allocation, farm heterogeneity, NManager, Lake Rotorua

EXECUTIVE SUMMARY

The distribution of nitrogen discharge allowances can alter the cost incidence of water quality targets and therefore has significant equity and political implications. Modelling cannot determine the best method for allowance allocation, but it can provide critical information on some of the likely, and not necessarily immediately obvious, implications of specific allocation options.

We develop a new methodology for NManager, a simulation model of water quality policy, to compare two approaches to the free allocation of nutrient discharge allowances in a trading system for the Lake Rotorua catchment. Our main goal is to explore fairness issues in allocation, and this requires an understanding of heterogeneity at the enterprise level. The original NManager model, a representative farm model, is therefore modified to incorporate farm heterogeneity. We parameterise it to observations of farmers in the catchment and simulate a policy environment specified by the Bay of Plenty Regional Council (BoPRC).

The updated model includes heterogeneous farms, but it retains the basic structure of the original NManager. Most importantly, land management practices are modelled through quadratic profit functions; subject to these high-level profit functions, landowners are assumed to make optimal decisions; and the trading of allowances is assumed to be costless.

The virtue of this simple framework is that it yields transparency as well as some generalizable results that are robust to more complex specifications. Although NManager is not intended to be a prediction tool – for example, its cost estimates are only indicative – the results nevertheless illuminate the scale of fairness issues that must be dealt with in implementing allowance allocation in the catchment.

In estimating the direct costs of water quality policy to farmers, we include the cost of the mitigation response (measured as loss of profit), the cost of nutrient allowances purchased and the revenue from the sale of any allowances the farmer does not use. In reality, these direct costs do not reflect ultimate cost incidence. For example, farmers may be able to pass on some of the costs to consumers through higher commodity prices or to farm workers through lower wages. We are unable to consider such general equilibrium effects in NManager.

Benchmarking and other data were provided to us by the BoPRC. Implementation of farm heterogeneity in the model involves some effort to extract and prepare these data for analysis. We reclassify and manipulate the data in order to achieve consistency across our different data sets.

Farm heterogeneity in NManager is implemented through the specification of idiosyncratic profit functions. To the extent possible, we calibrate these profit functions to observed farm-level data. We also build on intuition from economic theory and invoke previous empirical research. Nevertheless, the process of parameterisation requires some potentially strong assumptions. Some of our results are sensitive to these, while others are relatively robust.

We model two approaches for the free allocation of nutrient discharge allowances: a sector-based averaging approach and a grandparenting approach. Both approaches are subject to an identical clawback over time. The clawback is set to ensure that the BoPRC's nutrient target is met in each year. Reaching the final target requires that total nutrient loss from pastoral agriculture in the groundwater catchment be reduced from the current level of 526 tonnes N/year to 256 tonnes N/year in twenty years. An intermediate target of 337 tonnes N/year is to be achieved in ten years.

The costs of meeting the nutrient target for Lake Rotorua are expected to be significant. Our simulations of nutrient loss suggest that meeting the final target requires the conversion to forestry of many dry stock farms, and the implementation of farm management techniques that could be described as similar to best practice on most dairy farms. Converting to forestry allows dry stock farmers to sell their unused allowances to dairy farmers. More heterogeneity among farmers (by increasing the variation in marginal mitigation cost around its mean) increases the potential to benefit from trade and hence it makes the ability to trade more important.

Aggregate sector-level outcomes are not affected by the manner in which free allowances are allocated to farmers, because under both approaches, each sector receives exactly the same amount of allowances. On the other hand, farmers within each sector may experience a wide range of impacts from the choice of an approach. Individual farmers from either sector may therefore express a strong preference for an allocation mechanism.[1]

It follows from the nature of the modelled allocation approaches that the distribution of relative impacts experienced from allocation is determined by the distribution of benchmarked nutrient losses around their sector-level mean. Specifically, farmers with higher-than-average nutrient loss are unambiguously better off under grandparenting, and farmers with lower-than-average nutrient loss are unambiguously better off under the sector-based averaging approach. While this observation is not in itself an argument for either allocation method, it does have implications for policy design, for it points out that the allocation mechanisms could satisfy different principles for cost sharing.

We offer three principles for cost sharing to consider, and we refer to these during our discussion of results. The 'equal sharing' principle, the 'polluter pays' principle and the 'responsibility for action' principle all represent valid (and potentially conflicting) points of view that could be taken into consideration during the design of an allocation mechanism.

The allocation of nutrient discharge allowances will always be politically contentious because the potential transfer of wealth is large. Allocation has a major impact on ultimate cost sharing – our modelling compares two allocation methods and puts these in the context of a simulated trading system to illustrate equity considerations in each sector. In addition, we highlight other related issues that may also affect the desirability of a particular allocation mechanism.

[1] More precisely, the allocation mechanism affects the wealth of owners of the land. Unless explicitly noted otherwise, we assume the farmer is also the landowner.

1.0 INTRODUCTION

In this paper, we use the NManager simulation model to compare two approaches to the free allocation of nutrient discharge allowances in a trading system with heterogeneous farmers: a grandparenting approach and a sector-based averaging approach. Both approaches are subject to an identical clawback that is specified to ensure that the nutrient target is met. We parameterise the model to observations of farms in the Lake Rotorua catchment.

The paper is structured as follows. Section 2 introduces the NManager model and discusses its key assumptions and their implications. Section 3 summarises our main data sources, and section 4 provides a detailed documentation of our methodology: it describes the approach we take in specifying heterogeneity and the manner in which we implement the nutrient cap and the allocation of allowances. Section 5 contains our simulation results as well as an in-depth discussion of these results. Finally, section 6 concludes.

2.0 THE NMANAGER MODEL

NManager is a simulation model developed by Motu and the National Institute for Water and Atmospheric Research (NIWA) as a research tool for water quality policy (Anastasiadis et al., 2011). NManager was designed to address issues around the stringency, cost effectiveness and impacts of different policies for nitrogen regulation, particularly in catchments where complex surface and groundwater hydrology slows the transportation of nitrogen from farms to the lake. The current version of NManager was calibrated to the Lake Rotorua catchment building on results from the Rotorua Taupo Nitrogen Model (ROTAN), a GIS-based catchment hydrology and water quality model.

NManager was originally set up as a representative farm model, meaning that each land use was represented by a single – representative – farm. That is, all farms of each type were assumed to be homogenous in their nutrient loss and economic outcomes. The current project calls for incorporating farm heterogeneity in NManager, which requires modifications to the model. These are described in detail in subsequent sections of the paper; here we give a brief introduction to the general structure of the model and some guidance for interpreting its simulation outcomes.

2.1 MODEL STRUCTURE

Land management practices are modelled through the use of profit functions in NManager. Landowners' profit (per hectare per year), π, is expressed as a quadratic function of nutrient discharge, x; profit is monotonic over the nutrient loss values of interest, so higher profitability is feasible through farming practices that increase nutrient loss.[2]

$$\pi = a x^2 + b x + c$$

NManager distinguishes between land initially used for dairy farming and dry stock farming, and represents the two land uses by a different parameterisation of the profit function.

Landowners may reduce their nutrient loss by changing farm management practices or converting their land to a different, less nitrogen-intensive, use. Under regulation, there may be a cost associated with discharging nutrients, and landowners will mitigate as long as the cost of mitigation (measured as loss of profit) is less than the cost of discharging nutrients. As the profit curves are concave in the amount of nutrients used, the marginal cost of mitigation increases as the amount of mitigation increases. We assume that landowners will optimise by choosing the point where the change in profit from increasing nutrient discharge by 1 kg N is equal to the cost of discharging 1 kg N. Mathematically, this point occurs where the slope of the profit function is equal to the price, P, of allowances.

$$P = 2 a x + b$$

[2] This implies that the marginal cost of mitigation is strictly positive; there is no 'free lunch': mitigation is always costly on the margin. It is a (qualitatively) accurate representation of reality if farmers are already using less nitrogen than they would use in an unconstrained environment. Farming in the Lake Rotorua catchment is subject to pre-existing water quality regulation, so positive marginal mitigation cost is expected.

This equation specifies an inverse demand function from which it is straightforward to derive landowners' demand for nutrient allowances (as a function of the price of those allowances). The supply of allowances is fixed at the quantity set by the nutrient cap. NManager takes the demand and supply for allowances and applies a numerical approximation algorithm to determine their equilibrium price in each year. A tighter nutrient cap leads to a higher allowance price, resulting in an increase in the amount of mitigation performed.

2.2 KEY ASSUMPTIONS

NManager is an optimisation model. It is built on the assumption that landowners have perfect information and perfect foresight, and that their mitigation actions are optimal (within the constraints of the model). Optimisation models have many advantages, but they tend to underestimate the costs of water quality policy because their assumptions are typically not fully met in reality.

It is important to stress that NManager is not a farm systems model. Its profit functions provide a simplistic representation of farm properties.[3] The model's main focus is not on farm-level accuracy, but rather on outcomes that arise from interactions taking place among many actors in a nutrient trading system.

The trading of nutrient discharge allowances is assumed to be costless in NManager. One practical consequence of this is that the initial allocation of allowances will not affect the economic decisions of participants: allowances will be traded to ensure mitigation takes place where it is cheapest to mitigate.[4] Allowance allocation does, however, matter for cost sharing.

On the other hand, there are reasons that the cost figures in NManager may be overestimates. They represent direct costs from a partial equilibrium view. Farmers may be able to pass some of these costs forward to consumers or backward to farm employees, so the ultimate impact on farmers is likely to be less than that suggested by the model.

There is no scope in the model for technological change. This assumption becomes potentially important over longer time horizons, and it qualifies our long-term simulation results. Improvements in technology may reduce mitigation costs over time. The long-term impacts simulated via NManager may therefore be overestimates. Similarly, other economic conditions (for example, commodity and input prices) are likely to change over time, but these changes are impossible to predict. The impact of these changes on mitigation costs is unknown.

For these reasons, NManager is not intended to be an accurate prediction tool. Its strength lies in drawing comparisons across scenarios, not in forecasting the future precisely. Consequently, although we present numerical results in section 5, they should be considered illustrative rather than predictive.

3 Although real-life decision-making processes are undeniably more complex, this simple formulation seems to describe key aspects of farmer behaviour well. It also provides simplicity and transparency in face of the paucity of evidence on actual mitigation behaviour and costs (Anastasiadis, 2011).

4 This is likely to be a reasonable approximation in the long run, but not necessarily in the short run.

3.0 DATA PREPARATION

The data necessary for this project were spread across four Geographic Information System (GIS) data sets and numerous Overseer files (Overseer is a farm management model used to estimate nutrient flows in a productive farming system and to develop on-farm nutrient budgets). In order to run the NManager model with heterogeneous farm properties, we first extract and prepare the key data from these sources. This section documents the main data preparation steps.

Two maps of land within the Rule 11 boundary were provided by the BoPRC.[5] The first map identifies land that has been benchmarked under Rule 11 as well as some of the characteristics of this land via select variables recorded in the benchmarking process (including land use, nitrogen leaching and area for each Overseer block). The second map identifies rural land that has not been benchmarked under Rule 11. In addition, we were provided with the Overseer files for nearly all of the benchmarked properties.

We supplement the two Rule 11 maps with a map drawn from ROTAN that includes all land within the Lake Rotorua groundwater catchment (National Institute of Water and Atmospheric Research, dataset, 2011) and the map of Land Use Capability (LUC) classification (Landcare Research & MAF, dataset, 2002). The latter provides an additional measure of land quality across all properties.

We assign rural properties to three groups based on these data: benchmarked properties within the Rule 11 boundary; non-benchmarked properties within the Rule 11 boundary; and non-benchmarked properties outside of the Rule 11 boundary but within the groundwater catchment boundary. The first three data rows of Table 5.1 help identify the land areas associated with each group. About 82 percent of the dairy area and 57 percent of the dry stock area within the catchment have been benchmarked. Data preparation for each of the three groups of properties is discussed in turn below.

The benchmarked properties within the Rule 11 boundary are described by both the map of benchmarked properties and the Overseer files. For these properties, we have sub-property information at the Overseer block level. This data includes the type, size, nitrogen loss and geophysical characteristics – such as mean slope, rainfall and altitude – of each block. Using the block type variable, we divide each property into endogenous and exogenous blocks according to the classification in Table 3.1. Blocks that are used for pastoral, arable or horticultural uses are categorised as endogenous. We expect that these are the blocks a farmer would actively manage for nitrogen loss. Given this classification, we determine for each property its total area and nitrogen loss using both endogenous and exogenous blocks, and we determine its average slope, rainfall, altitude and LUC class using endogenous blocks only.

---

5 Rule 11 is a series of rules in the Bay of Plenty Regional Water and Land Plan aimed at addressing water quality issues resulting from land use activities. It applies to catchments of lakes with degraded water quality.

Table 3.1 The classification of Overseer blocks

Classification	Overseer block type
Endogenous	Crop, Cut and Carry, Fodder, Fruit Crop, Pastoral
Exogenous	House, Non-productive, Riparian, Trees, Uncategorised

The block size reported for benchmarked farms within the Rule 11 area does not always correspond to the geometric area given by the map. The inconsistency likely arises because data for some properties that are crossed by the Rule 11 boundary include land on both sides of the boundary (Alastair MacCormick, personal communication). We use the reported area when calculating per-hectare average values from the Overseer data, and we use the geometric areas when incorporating the data into the NManager model.

In deriving farm-level measures of nitrogen loss and land area by adding up Overseer blocks, we lose some information as to the underlying land uses that make up a farm property. From the perspective of modelling mitigation, this does not necessarily matter because land-use is not an explicit choice variable in NManager.

We require a measure of production for each farm in order to estimate its profitability. We extract milksolid production for dairy farms and average stocking rate for dry stock farms from the Overseer benchmarking files. Farms that report milk production are classified as dairy farms for our analysis. Farms that report the number of stock units grazed, without any milk production, are classified as dry stock farms. Farms that report no measure of production are classified as exogenous and are excluded from our NManager analysis (though we still account for their pastoral nitrogen loss). Because we are unable to control for multiple land uses within a farm, we likely underestimate the productivity of mixed farms (for example, of farms that both milk a dairy herd and raise stock for slaughter and of farms that have a significant cropping focus). As we note later, mixed farm systems pose a potential issue for allocating allowances as well. However, very few properties in our data set could truly be called mixed-use farms – the vast majority of properties are predominantly dairy or predominantly dry stock.

We lack Overseer files for a small number of benchmarked properties. Where the Overseer files are missing, we determine whether a farm is dairy, dry stock or exogenous based on its constituent block types (which we are able to identify from the geographic data). Those properties identified as dairy or dry stock farms are assigned the average production per hectare for their respective type.

For lack of relevant farm-specific data, the non-benchmarked properties within the Rule 11 boundary cannot be modelled explicitly. We use the ROTAN land-use map to determine how much of the non-benchmarked land is used for dairy and dry stock farming. We create two representative farms to capture these properties: a representative dairy farm and a representative dry stock farm. These have the characteristics of the average farm within the corresponding sector.

Likewise, the properties outside the Rule 11 boundary but within the groundwater boundary cannot be modelled explicitly. Following the same approach as above, we create two representative farms for this area (one for dairy and one for dry stock). We keep these farms separate from the previous representative farms because their contribution to lake loads (and hence the amount of allowances they need to surrender) is different. Results from ROTAN suggest that 53 percent of nitrogen enters the lake through groundwater. Farms outside the Rule 11 boundary contribute to lake loads through groundwater only, so we assume that 53 percent of the nitrogen from these farms reaches the lake.

The ROTAN model suggests that current pastoral nutrient loss in the groundwater catchment is 526 tonnes N annually. Some of this is considered unmanageable. In this project, there are three sources of unmanageable nitrogen exports: the loss from exogenous blocks on dairy and dry stock farms within the Rule 11 boundary; the loss from pastoral blocks that report no productive use; and a loss of 4 kg nitrogen per hectare per year on all land (this amount is unmanageable because it persists even if pastoral land is converted to forestry). Table 3.2 identifies the amount of nutrient loss associated with each of these sources. The remaining 422 tonnes N enter the NManager model (this includes the approximate amount of nutrient loss from non-benchmarked properties).

Table 3.2 The sources and amount of annual nitrogen loss

Source	kg N
Baseline 4 kg N for all endogenous land	85,226
Exogenous nitrogen from dairy and dry stock farms (for example, from buildings and non-productive land)	16,656
Pastoral Overseer blocks on exogenous properties	2,088
Endogenous nitrogen included in NManager	422,030
Total pastoral nitrogen	526,000

We calibrate the amount of nutrient leaching per hectare for the representative farms we model to ensure that the nitrogen loss included in NManager is consistent with the nitrogen loss suggested by ROTAN. The two representative dairy farms are assigned a leaching of 45.74 kg nitrogen per hectare per year (this matches the average nutrient loss of the benchmarked farms we classify as dairy within the Rule 11 boundary).[6] The two representative dry stock farms are assigned a leaching of 14.93 kg nitrogen per hectare per year (this value is required to ensure that total nitrogen loss aggregates to 526 tonnes N in the groundwater catchment; it is lower by about 3 kg/ha than the average for benchmarked dry stock farms). We also account for the fact that only about half of the nutrients from properties outside the Rule 11 boundary contributes to lake loads. By constructing the representative farms in this manner, we are able to account for all pastoral lake loads in the catchment and also for the mitigation that might be performed by these non-benchmarked farms.

[6] This figure is slightly lower than the 49 kg N/ha sector average value supplied to us by the BoPRC. The discrepancy has to do with our classification of farms into the two sectors; it only has a small quantitative impact on our results.

4.0 METHODOLOGY

In this project, we consider an export trading regime under which landowners must hold a quantity of nutrient discharge allowances that is sufficient to cover nitrogen losses from their property each year. We devote this section principally to the specification of farm heterogeneity through the use of profit functions. We also describe how we implement the nutrient caps necessary to achieve the BoPRC's water quality targets over time.

4.1 PARAMETERISATION OF FARM PROFIT FUNCTIONS

The original NManager model is based on a single representative farm for each land-use. This project requires that we incorporate farm heterogeneity into NManager. Given the structure of the model, heterogeneity operates through the specification of idiosyncratic profit functions that define the relationship between profit and nitrogen leaching on each farm. These relationships are unobserved in the real world: they are based not only on profitability given current farming practices, but also on profitability under a host of alternative farm management options and land-use change. Because we do not observe outcomes under alternative management options, we are required to make some relatively strong assumptions to specify the profit functions. To the extent possible, our parameterisation draws on previous research in the Lake Rotorua catchment (Smeaton et al., 2011; Perrin Ag Consultants Ltd & AgResearch, 2012). Regardless, it is important to emphasise at the outset that accurately describing the profit function of any given farm is impossible: farms are complex systems managed by idiosyncratic owners who respond to technology, input costs, commodity prices, weather and a range of other factors. The best we can hope for is to encapsulate relationships that are reasonable on average. Because of the many unobserved and unpredictable factors that affect farm outcomes in real life, the farms in NManager should be viewed not as the specific properties they are modelled on, but as entities representing properties with similar observed characteristics.

We maintain the quadratic profit functions that are already coded into NManager. Quadratic functions are computationally easy to solve and have some desirable behavioural implications that are consistent with economic theory. For example, with a quadratic profit function, the marginal cost of mitigation increases with the amount of mitigation performed. Quadratic functions are fully defined by three parameters. Our approach is to specify nutrient losses and the associated level of profits for two points on the curve: in the baseline – benchmarked – state, and under maximum mitigation (i.e., full conversion to forestry). Fixing the third parameter determines the curvature of the quadratic function.

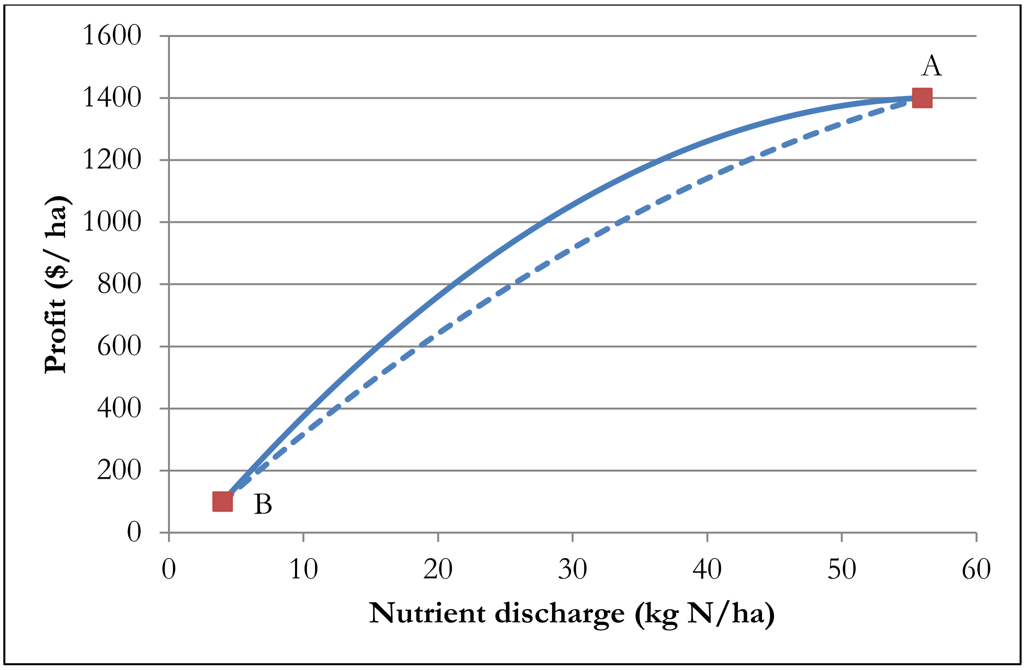

Figure 4.1 Illustrative quadratic profit functions.

Figure 4.1 illustrates two hypothetical profit functions that have been drawn for identical baseline (point A) and fully mitigated (point B) outcomes.[7] The quadratic functional assumption implies that the locus of points given by the combination of potential mitigation options and the profit levels associated with these options can be connected by a smooth curve. There is a smooth transition between on-farm mitigation and land-use change, for example. This may not be as restrictive an assumption as it seems at first because it is possible for a farmer to apply certain mitigation practices on a portion of the farm's land area.

The first derivative (or slope) of the profit function corresponds to the marginal cost of mitigation: the amount of profit foregone by reducing nutrient loss incrementally. The three parameters of the profit function thus also determine the intercept and slope of the marginal mitigation cost curve. (Intuitively, the intercept and slope represent the lowest nutrient price needed to induce any mitigation and the rate of change in mitigation with changes in the nutrient price, respectively.) In the example above, starting at point A, the farm characterised by the solid profit curve will carry out more mitigation initially (that is, at low levels of stringency) because mitigation is cheaper on this farm (the slope of its profit function is lower in the baseline state).

We parameterise the profit function of each farm based on observed and projected farm characteristics. Nutrient discharge and profit (per hectare) in forestry use (the fully mitigated state) are assumed to be identical everywhere. This assumption fixes point B at the same location for each farm. We assign the amount of nutrient leaching specified by the ROTAN model to this point: 4 kg N/ha. The profitability of forestry was specified by the BoPRC to be $150/ha.

[7] This example is purely illustrative.

The next step in the parameterisation process is to identify the baseline outcome; this corresponds to point A in Figure 4.1. Farm-specific information on nutrient loss is available through the BoPRC's benchmarking study, and information on production is available from the Overseer files, but we lack farm-specific information on profits.

For both dairy and dry stock farms, we establish an initial relationship between farm production and profit according to unit record data behind the MAF monitor farm reports (Ministry of Agriculture and Forestry, dataset, 2010). For dairy farms, out of the variables we have for the benchmarked farms, nutrient efficiency (output produced per unit of nutrient use) appears to be the best single predictor of profit per hectare. For dry stock farms, the stocking rate appears to be the best single predictor of profit per hectare. We use regressions of profit per hectare on these predictors to estimate the initial level of profit for each farm, subject to the constraint that profit is at least $160 per hectare. This constraint is necessary for the internal consistency of the model.[8] Following this estimation, we scale the fitted profit values for each farm type to align our average results with those of the Farmer Solution Project (Perrin Ag Consultants Ltd & AgResearch, 2012). The second and final data columns of Table 5.1 show that baseline profit is $1,804 per hectare for the average dairy farm and $343 per hectare for the average dry stock farm.

For both dairy and dry stock farms, our approach potentially underestimates the variation that exists in the profits of benchmarked farms because it is based on fitted values from an ordinary least squares regression (performed on a different sample of farms). For example, higher nutrient efficiency for a dairy farm will unambiguously lead to a higher estimated profit. Similarly, a higher stocking rate will unambiguously be associated with higher profit on a dry stock farm.[9]

The remaining step is to specify the curvature of the profit function. There is very little data we can draw on here because this step essentially requires determining, at the farm level, how profits change with additional mitigation. As explained below, we expect farmer skill to affect the curvature of the profit function, so we attempt to develop a proxy for farmer skill from observed farm-level data.

At an intuitive level, we might expect more skilled farmers to have a lower initial marginal cost of mitigation. Some potential reasons for this may be that

- skilled farmers could be expected to be more responsive to price signals
- skilled farmers may be better at mitigating and may have more capacity to adapt
- less-skilled farmers may be more resistant to changing farming practices and hence they may mitigate slowly initially.

These arguments suggest that higher farmer skill may tend to increase the curvature of the profit function. This has plausible behavioural implications. For example, it implies that skilled farmers respond to the policy first by adopting mitigation at lower nutrient prices. It also implies that at moderate nutrient prices, it is less skilled farmers who will tend to convert to

[8] Recall that profit in forestry use is assumed to equal $150/ha – a dry stock profit below this value would imply negative mitigation costs.

[9] This is not expected to affect our results significantly. The relationships we rely on to estimate profits are moderately strong, and there is large variation in the explanatory variables (nutrient efficiency for dairy and the stocking rate for dry stock) among our sample of benchmarked farms. However, it is important to keep in mind that we are not able to accurately predict profitability for any given farm – we regard farms in our sample in a statistical sense.

forestry: skilled farmers are able to operate profitably under tighter constraints, and they may also have more capital invested in the farm which would make them less likely to change land use.

Farmer skill is, of course, unobserved, so it cannot be used directly in parameterising profit functions. We follow the methodology developed in Anastasiadis and Kerr (2012) to develop a proxy for the unobserved variable. They consider a model where the nutrient efficiency of a farm is a function of farmer skill and a mix of other factors – some of these, such as land and climate conditions, cannot be managed by the farmer. In order to consider heterogeneity in nutrient efficiency due to farm management practices, Anastasiadis and Kerr determine, using a regression framework, how much of the difference in nutrient efficiency between farms is due to differences in inputs that cannot be managed by the farmer. Any residual variation between farms is then attributed to factors that can be managed by farmers and therefore are affected by changes in farm management practices. This residual variation is considered to capture the effect of farmer skill.

We are unable to build on the regression results of Anastasiadis and Kerr directly because we do not have data on comparable variables. We therefore perform an analogous regression of nutrient efficiency on the observed geophysical attributes of each farm in our dataset: the average slope, rainfall, altitude and LUC class.[10] The fitted values from this regression represent the degree of efficiency that can be explained by exogenous variation in these geophysical factors. The residual captures the effect of unobserved factors – we loosely interpret this as a measure of farmer skill.[11] The effectiveness of the approach depends on how well the residual identifies the nutrient efficiency that is due to factors that can be managed by farmers. Due to the small amount of variability within our sample, the regressions in our application have weak explanatory power. This means that the residual, our measure of farmer skill, causes most of the variation in nutrient efficiency.

Given proxies for farmer skill, we determine the curvature of farms' profit functions as follows. We use empirical results from Smeaton et al. (2011) to determine plausible bounds on the curvature of the profit functions. These results suggest that the curvature of dairy farm profit functions may vary between 20 and 100 percent of the maximum possible curvature, while the curvature for dry stock farms may vary between 50 and 100 percent of the maximum possible curvature (the maximum possible curvature can be determined mathematically given that the cost of mitigation is assumed to always be greater than zero). Given these bounds on curvature, we distribute farmers across these ranges based on our measure of farmer skill.

Anastasiadis and Kerr (2012) review the existing literature on mitigation and farm profits and provide some justification for our approach. While multiple prior studies have considered mitigation costs, we cannot combine their results to specify our profit functions as different authors have made different assumptions, and in general, insufficient information is provided to standardise the results.

---

[10] For dairy farms, our measure of nutrient efficiency is milksolid production per unit of nutrient loss; for dry stock farms, we use stock units per unit of nutrient loss.

[11] If skilled farmers have better land (in terms of the geophysical attributes we observe), our specification suffers from omitted variable bias. If this were the case, the estimated effect of observed land quality would be biased upward. As described in the next paragraph, our parameterization relies on farmer skill in an ordinal sense, so the potential bias is not of great concern.

In summary, two pieces of data define the profit function of a dairy farm: nutrient loss and nutrient efficiency (the latter affecting both estimated baseline profit and curvature). The profit functions imply that those with higher nutrient loss in the baseline are able to mitigate more cheaply. The effect of nutrient efficiency is ambiguous: on the one hand, higher efficiency implies greater farmer skill (higher profit function curvature) leading to cheaper initial mitigation; on the other hand, higher efficiency also implies higher baseline profit leading to more expensive initial mitigation through a generally steeper profit function.

For dry stock farms, three pieces of information are used to define the profit function: nutrient loss, the stocking rate (to determine baseline profit) and nutrient efficiency (for curvature). All else equal, those with higher nutrient loss, a lower stocking rate and higher nutrient efficiency can carry out mitigation at a lower cost initially.

4.2 IMPLEMENTING THE NUTRIENT CAP

Total nutrient loss from pastoral agriculture within the groundwater catchment has been estimated at 526 tonnes N/year by the ROTAN model (Rutherford et al., 2011). The Bay of Plenty Regional Council aims to reduce this to 256 tonnes N/year; the target is to be met fully in 20 years. It is desired that seventy percent of the reduction in nutrient loss takes place in the first 10 years. Accordingly, we set the cap at 337 tonnes N/year in the tenth simulation year. For simplicity, we fit a piecewise linear function to these three data points. That is, in each time period between years 1 and 10, the cap is reduced by 18.9 tonnes N, and in each time period between years 11 and 20, the cap is reduced by 8.1 tonnes N. Changes in the cap over time are illustrated by the series labelled ‘catchment cap’ in Figure 4.2.

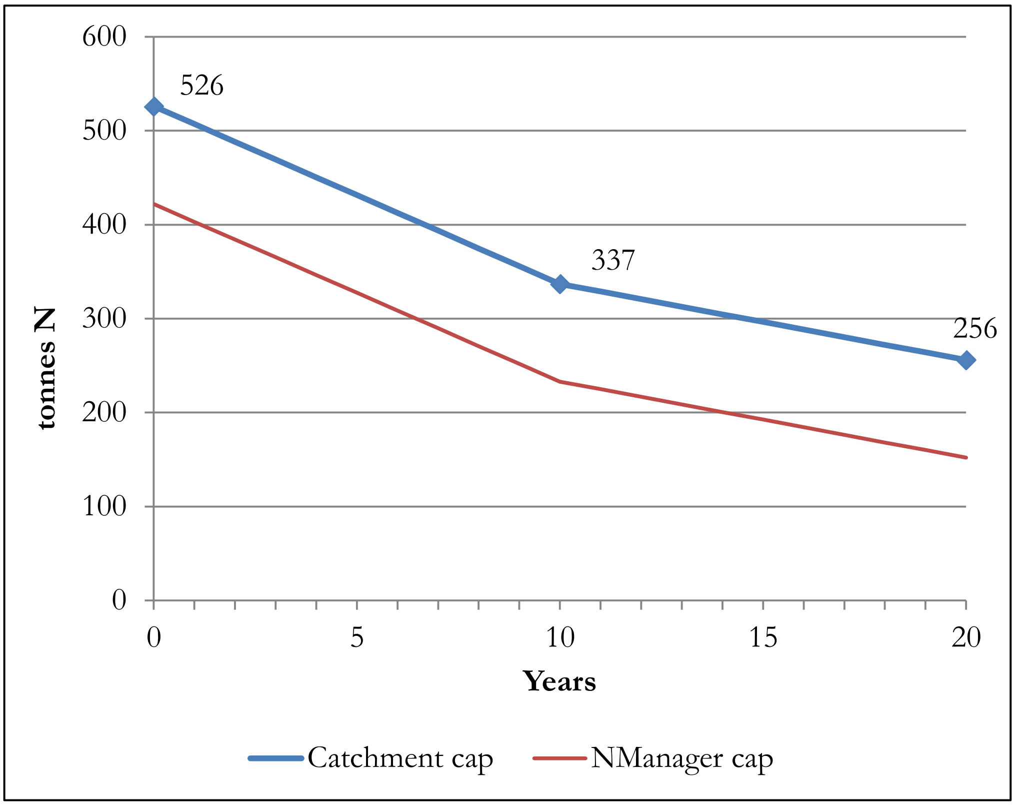

Figure 4.2 The nutrient cap and its implementation.

As noted previously, some of the nutrient loss associated with land use is unmanageable. For example, even full conversion to forestry does not completely eliminate nutrient leaching. Following the ROTAN model, we assume that unmanageable nutrient loss, i.e., leaching under forestry use, is uniformly 4 kg N/ha (Rutherford et al., 2009). Further, a small proportion of nutrient loss is classified as unmanageable because it is associated with non-productive land or exogenous properties. In all, we consider 103.97 tonnes of N unmanageable (impossible to mitigate) in the study area. From a modelling perspective, only the manageable portion of nutrient exports enters NManager, so the cap implemented in the model is lower than the Regional Council's cap by the amount of total unmanageable N in each year. This is made explicit in Table 3.2, and it is also shown in Figure 4.2. Reaching the target of 256 t N requires a reduction of around 51 percent in total nutrient loss, but this corresponds to a much higher, 64 percent, reduction in manageable nutrient loss.

4.3 THE ALLOCATION OF ALLOWANCES

We model two approaches for the free allocation of nutrient discharge allowances in the catchment.

1. A sector-based averaging approach, under which each hectare of land receives an allocation based on the mean amount of baseline nutrient loss within the sector it is currently in. Under this approach, the allocation to each hectare of dairy is the same, and the allocation to each hectare of dry stock farmland is the same.
2. A grandparenting approach, under which each parcel receives an allocation of free allowances in proportion to the amount of baseline nutrient loss specific to the parcel, as established through the benchmarking work of the BoPRC.[12]

The nutrient cap determines the supply of allowances. Under both approaches, the allocation is therefore subject to a clawback that ensures the nutrient cap is met exactly in each year. Reductions in the NManager cap are implemented by a proportional reduction in the amount of allowances allocated to each farmer. To be precise, the reduction in allowances is based on the farmer's benchmarked manageable (not total) nutrient loss. This ensures that the free allocation is always sufficient to cover unmanageable losses. For instance, consider a one-hectare property with a total baseline nutrient discharge of 6 kg N/ha. A cap requiring a 50 percent reduction in manageable nutrient loss relative to the baseline would result in a free allocation of 5 units to the owner of this property: 4 units for unmanageable and 1 unit for manageable nutrient loss (the latter representing a 50 percent reduction relative to the benchmarked amount). The NManager cap shrinks by 3.5 to 7.5 percent each year. This decrease also corresponds to the reduction in free allocation for manageable nutrient loss relative to the previous year's amount.

Farmers with a nutrient loss below their sector's average receive more allowances under the sector-based averaging approach, and vice versa. This result trivially follows from the definition of the allocation methods, and it fundamentally determines the distribution of impacts experienced by farmers from the choice of an allocation approach. As long as allowance trading is costless, this distribution does not depend on the manner in which we model mitigation behaviour (so it is robust to the parameterisation of profit functions).

[12] Both of these approaches involve the free allocation of allowances based on historical nutrient loss and can therefore be considered different forms of grandparenting.

5.0 RESULTS AND DISCUSSION

In this section, we present the results of the NManager simulation. Although we discuss most simulation outcomes quantitatively, it is worth reiterating that NManager is not primarily a prediction model – the strength of the framework is in the formal exploration of qualitative relationships. The discussion in this section is therefore subject to the qualifications laid out previously in sections 2.2 and 4.1. We discuss aggregate (catchment-level and sector-level) outcomes before proceeding to consider the impact on heterogeneous farmers within each land-use sector in turn.

The incidence of costs can be modified through the free allocation of nutrient discharge allowances. To motivate the interpretation of results, we consider three potential, possibly conflicting, cost-sharing (or equity) principles and refer to them from time to time during our discussion. People often feel that a more equal distribution of costs is more equitable. Equal sharing could, however, be defined in many different ways: for example, on the basis of costs per person, per farm or per hectare. In this application we consider that the 'equal sharing' principle is better satisfied when per-hectare costs are distributed more equally within a group of farmers; the 'polluter pays' principle is satisfied when those with higher nutrient loss bear a larger burden; and the 'responsibility for action' principle is satisfied when, on the one hand, those who choose to have high nutrient loss (for example, through the adoption of intensive production technologies) bear more costs than those who have high nutrient loss due to factors outside their control, and on the other hand, those who have low nutrient loss because they undertook costly mitigation in the past are rewarded relative to those who did not incur such costs. In addition to these, there exist other valid principles for cost sharing that may need to be weighed and balanced during the design of an allocation mechanism (Kerr & Lock, 2009; Kerr et al., 2012).

5.1 AGGREGATE RESULTS

Figure 5.1 reproduces the simulated price path of nutrient discharge allowances.[13] As the supply of allowances decreases with the tightening cap, their market price rises steadily. Beyond the twentieth simulation year, we do not project further changes in the allowance price because the cap remains constant at 256 tonnes N. The figure also identifies the perpetuity value of a permanent allowance: this is the constant stream of income one would need to pay to acquire a right to discharge 1 kg nitrogen in each future year forever. The perpetuity is calculated using a 5 percent discount rate and considering the simulated price path of allowances with an infinite time horizon. The present value of the stream of income represented by the perpetuity – the amount a permanent right could be sold for today – is $415.[14]

[13] Throughout this paper, we define an allowance as a right to discharge 1 kg of nitrogen in a specific year. The price of an allowance in a given year can be thought of as the price for which such a right could be rented out during that year.

[14] For context, we note that this is higher than the standard price established by the Lake Taupo Protection Trust for each kg N permanently reduced in the Lake Taupo catchment. After discounting, the corresponding upfront price in the Lake Taupo catchment is estimated to be approximately $300 (Duhon et al., forthcoming). There is, of course, no reason for nutrient prices to be similar across catchments.

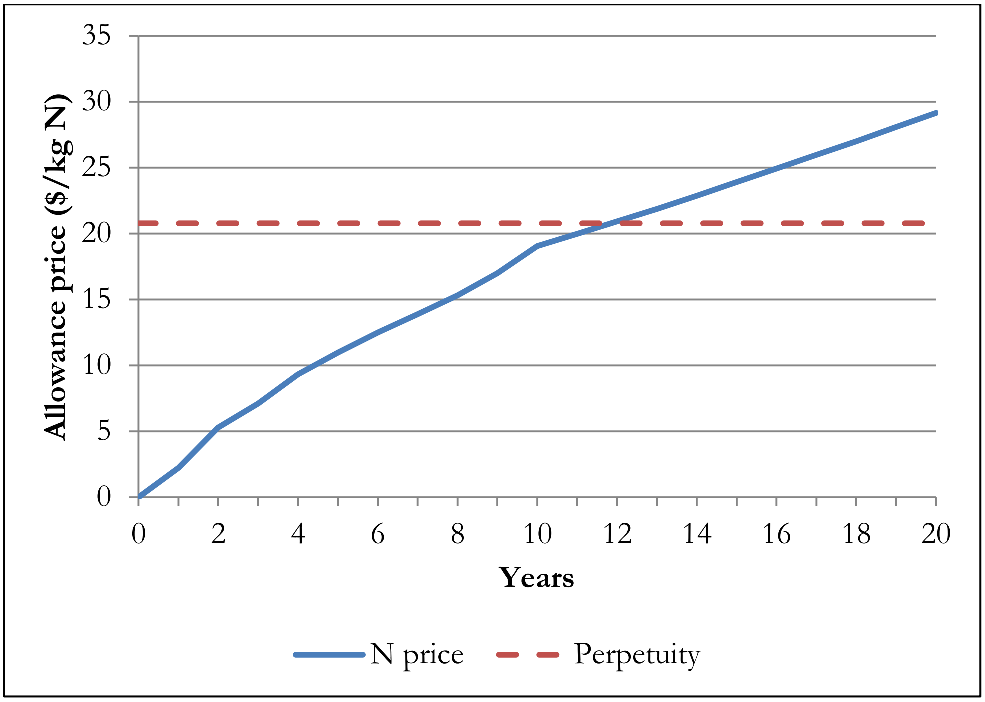

Figure 5.1 The simulated market price of allowances and the perpetuity value of a permanent right to discharge (calculated at time zero).

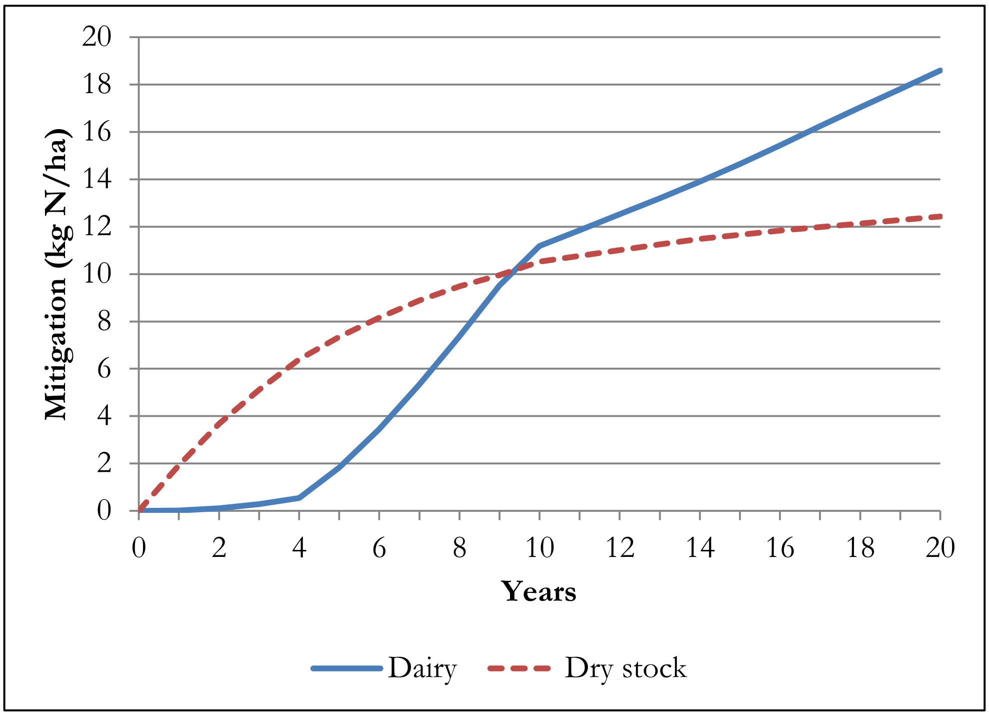

Figure 5.2 Mitigation per unit area across sectors (defined by initial land use).

At low nutrient prices (below about $10 per kg N), nearly all of the mitigation is carried out on dry stock farms; this is shown in Figure 5.2.[15] By the tenth simulation year, the dry stock sector is projected to mitigate nearly three quarters of its initial manageable nutrient loss (or 58 percent of its total nutrient loss); about a third of the sector's farms are projected to fully convert to forestry use by this time. This implies that further mitigation options in the sector are limited and meeting the nutrient cap is feasible only through additional mitigation by dairy farmers. Consequently, as the cost of allowances rises, mitigation performed in the dairy sector increases sharply.

Table 5.1 displays simulation results in more detail for particular years at the sector level (where sectors are defined by their initial land use). For context, the first several rows of the table also include simple descriptive statistics for our study area. The columns labelled 'total' show various outcomes (nutrient loss, mitigation, costs) at the scale of the groundwater catchment. The columns labelled 'per ha' show average (per hectare) outcomes for benchmarked properties only.[16] All entries for baseline nutrient loss show manageable nutrients only – total nutrient loss can be found by adding 4 kg N/ha.

Table 5.1 Sector-level simulation results

Outcome/Sector	**Dairy**		**Dry stock**	
	Total[1]	**per ha[2]**	**Total[1]**	**per ha[2]**
Total area (ha)	5,349		17,044	
Rule 11 area (ha)	4,927		15,154	
Benchmarked area (ha)	4,374		9,781	
Baseline nutrient loss (kg N/year)	215,019	41.75	207,011	14.04
Estimated baseline profit ($)	9,647,194	1,803.71	5,848,648	343.15
Year 10				
Mitigation (kg N)	51,345	11.18	137,655	10.52
Mitigation cost ($)	716,583	154.76	1,227,466	82.87
Allowance cost ($)	855,943	143.08	-855,938	-80.64
Total cost ($)	1,572,525	297.84	371,528	2.24
Value of free allocation ($)	2,260,895	438.97	2,176,692	147.65
Year 20				
Mitigation (kg N)	88,550	18.59	181,450	12.43
Mitigation cost ($)	1,615,599	333.49	2,269,185	127.46
Allowance cost ($)	1,428,951	236.81	-1,428,946	-100.51
Total cost ($)	3,044,550	570.30	840,238	26.94
Value of free allocation ($)	2,258,312	438.47	2,174,205	147.48

[1] Total figures are for all farms within the groundwater catchment.

[2] Per hectare averages pertain to benchmarked farms only.

[15] This potentially contradicts the findings of the Farmer Solutions Project (Perrin Ag Consultants Ltd & AgResearch, 2012) which suggest that the initial marginal cost of mitigation is low on dairy farms. Our profit functions imply that the marginal cost of mitigation is lower on dry stock farms, so it is these farms that mitigate first. This result is to a large extent driven by the large difference in the baseline profitability of dairy and dry stock properties.

[16] In terms of nutrient leaching per hectare, our representative dry stock farms are slightly different to the average benchmarked dry stock farm. The reason for this is explained at the end of section 3.

To any farmer, the total cost of the policy can be broken down into two components: the cost of the mitigation undertaken (measured as loss of profit) and the net cost of any allowances purchased. The latter component can be negative (representing a benefit) if the farmer is able to sell any unused allowances. Accordingly, total cost in Table 5.1 is calculated as the sum of the mitigation cost and allowance cost rows.

The aggregate amount of mitigation by both sectors is sufficient to meet the intermediate target of 337 tonnes N/year in year 10 and the final target of 256 tonnes N/year in year 20. Most of the reduction in nutrient exports – approximately 73 percent in year 10 and around 67 percent in year 20 – is provided by dry stock farmers. The sector faces higher total mitigation costs than the dairy sector. However, by performing this mitigation, dry stock farmers are able to sell valuable allowances to dairy farmers. The revenue from the sale of allowances largely offsets their cost of mitigation.

Dairy farmers also mitigate to reduce their nutrient losses, but the marginal cost of mitigation in the sector is higher than elsewhere: dairy farmers are better off by purchasing allowances from outside the sector (i.e., by paying dry stock farmers to perform mitigation for them). Because of this demand for additional allowances, the total cost to the sector is higher (both per hectare and overall) than the cost of mitigation carried out directly. We estimate that, on average, the policy's proportional impact on profits is higher in the dairy sector. (It is, however, much lower than it would be without trading.)

For reference, Table 5.1 also includes the value of free allocation received by each sector. If the allowances were auctioned off to the highest bidder, total costs would increase by this amount. By coincidence, the value of free allocation is similar in simulation years 10 and 20. Although the number of allowances received is lower in year 20, the market price of each allowance is higher – this increase in price almost exactly offsets the reduction in the amount of allocation.

To assess the effect farm heterogeneity has on outcomes at the sector level, we ran an auxiliary simulation using the original, representative farm, version of NManager (with appropriately calibrated farms). Mitigation costs are higher in both sectors with representative farms, on aggregate by about 30 percent in the tenth simulation year and by about 20 percent in the final (twentieth) simulation year. Qualitatively, this is a general result. Heterogeneity, by definition, provides additional opportunities for cheap (below average-cost) mitigation.

The results in Table 5.1 are not specific to either allocation scenario: none of the catchment- and sector-level outcomes are affected by the manner in which free allowances are allocated. There are three reasons for this. First, the nutrient cap is identical across the scenarios, which implies that the same amount of mitigation needs to take place. Second, with free trading, mitigation decisions are not affected by the initial allocation of allowances; mitigation will be carried out where it is cheapest to mitigate. Third, under both allocation approaches we consider, each sector receives exactly the same amount of allowances (and the price of these allowances is the same). For these reasons, the approaches differ in terms of their impacts within sectors, but not in terms of their impacts across sectors.

5.2 DAIRY SECTOR RESULTS

In this subsection, we take a closer look at simulation results specific to the dairy sector, and also provide a discussion of the impacts of free allocation.

The figure below reproduces the distribution of nutrient loss per annum per hectare on farms classified as dairy. It is directly based on benchmarking results, but it includes the manageable portion of nutrient leaching only.[17] A dashed vertical line shows the mean nutrient loss within the sector. We start out by presenting this distribution because, as we demonstrated in section 4.3, it is fundamentally related to the distribution of impacts experienced from the choice of one allocation scenario over the other. The figure shows that there is large variation in baseline nutrient losses within the sector, and it is this variation that determines the relative impact of the two allocation options we consider.

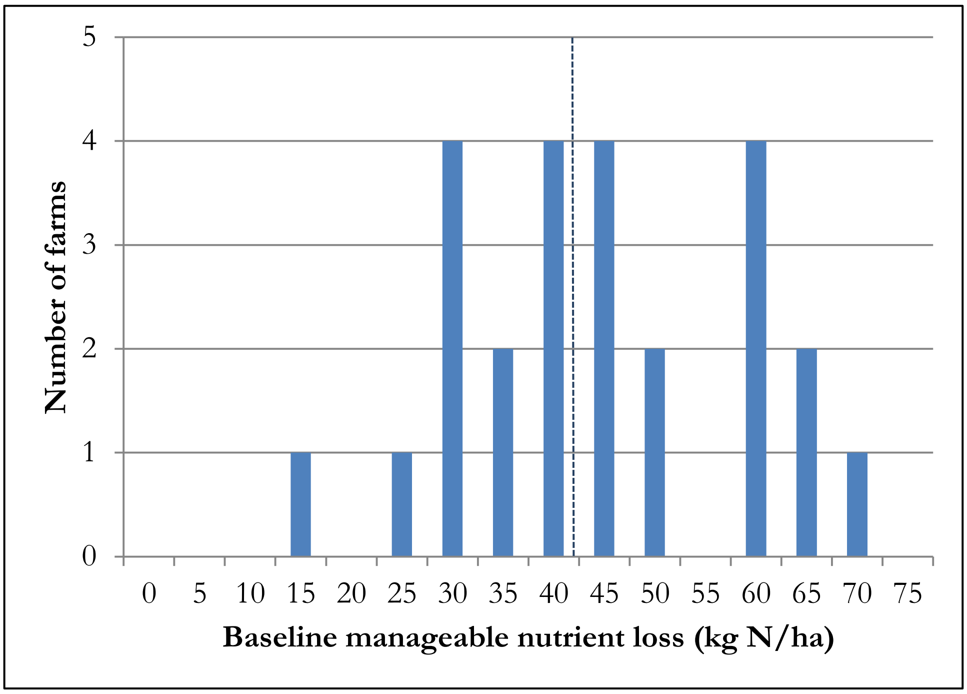

Figure 5.3 The distribution of baseline manageable nutrient loss among dairy farms.

The strength of a farmer's preference for an allocation method is directly related to the amount by which the farmer's baseline nutrient loss differs from the sector-level average. Dairy farmers with a nutrient loss below the sector's average (around 42 kg N/ha for the manageable portion) receive more allowances under the sector-based averaging approach. Conversely, farmers with a nutrient loss above the sector's average receive more allowances under the grandparenting approach. Those with low baseline leaching therefore unequivocally prefer a sector-based allocation approach to grandparenting, and vice versa.[18] Those with an average level of nutrient loss are indifferent between the allocation methods. Consequently, despite the fact that our profit functions are highly parameterised, the results that consider cost differences between the two allocation approaches are robust.

[17] All figures that follow in this section have been adjusted in the same manner and include manageable nutrient losses only.

[18] The low outlier in the histogram of Figure 5.3 is a mixed dairy and dry stock farm. For lack of data, we are unable to model sub-farm blocks separately, and our method of classification (which is directly based on the Overseer classification) assigns this farm to the dairy sector. Mixed farms are potentially problematic for at least two reasons. While we are able to account for all sources of nutrient loss on these farms, the profit

The wide distribution of benchmarked nutrient losses in Figure 5.3 suggests that the method of allocating nutrient discharge allowances matters to most dairy farmers. The histogram in Figure 5.4 illustrates the range of potential impacts experienced from the choice of an allocation approach in absolute (monetary) terms.[19] The horizontal axis indicates the difference in simulated total costs, per hectare of land, under sector-based averaging versus grandparenting – we call this the impact of allocation. (For comparison, recall that $570 per hectare is the mean total policy cost in the sector.) Farmers with a positive cost difference are better off under grandparenting, and vice versa. The horizontal axis is therefore related to the relative benefit a dairy farmer experiences from a grandparenting allocation of nutrient discharge allowances compared to the alternative approach. The absolute size of the impact depends on the value of free allocation (which, in turn, is determined by the amount of free allocation and the market price of allowances).

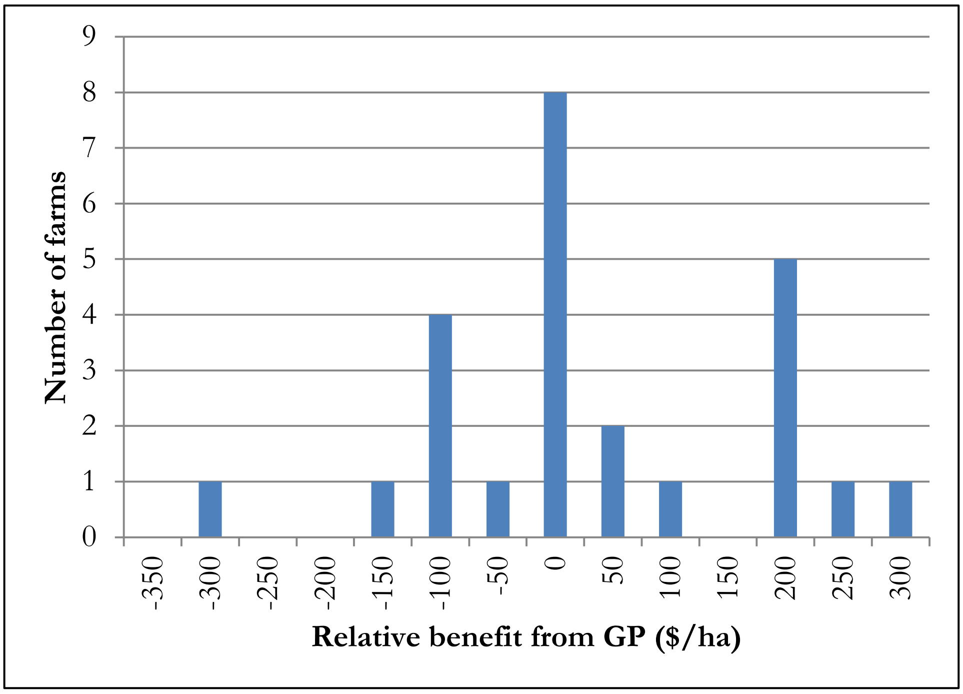

Figure 5.4 The final distribution of allocation impacts among dairy farms.

If allocation is tied to the land (and there are no market imperfections such that initial allocation would affect the economic decisions of participants), then we expect its value to be capitalised into farm land values (Allan & Kerr, forthcoming). The allocation impact is therefore entirely an impact on owners of land (who may be different from the person farming the land). Whether a windfall gain or a loss of equity, this impact would not be passed on to others through general equilibrium mechanisms.

---

functions of these farms are constructed on the basis of their main land use only. Further, mixed farms, by virtue of being assigned to a single sector, receive the sector-based allocation characteristic of their main land use for the whole farm area. This is a problem because it inflates the difference between the two allocation approaches. If the definition of sectors is similarly imprecise in an actual trading system, similar problems will arise. Therefore, when mixed farms are common, it may be appropriate to consider sub-farm level approaches to allocate allowances.

19 Unless explicitly noted otherwise, all simulation results presented in this section refer to the final, twentieth, simulation year.

While the absolute impact depends on the nutrient price, it needs to be stressed that the relative impact (the shape of the distribution) is entirely determined by the distribution of benchmarked individual nutrient losses around the sector's mean. This is true as long as mitigation costs to each farmer are exactly equal across the scenarios.[20] The manner in which we incorporate farm heterogeneity into the model makes no difference to the distribution of impacts: regardless of the form heterogeneity takes, the distribution owes its shape solely to the variation in baseline nutrient losses around the sector-level mean.[21]

For context, we also review the distribution of final (mitigated) nutrient losses, and of mitigation costs and total costs within the dairy sector.

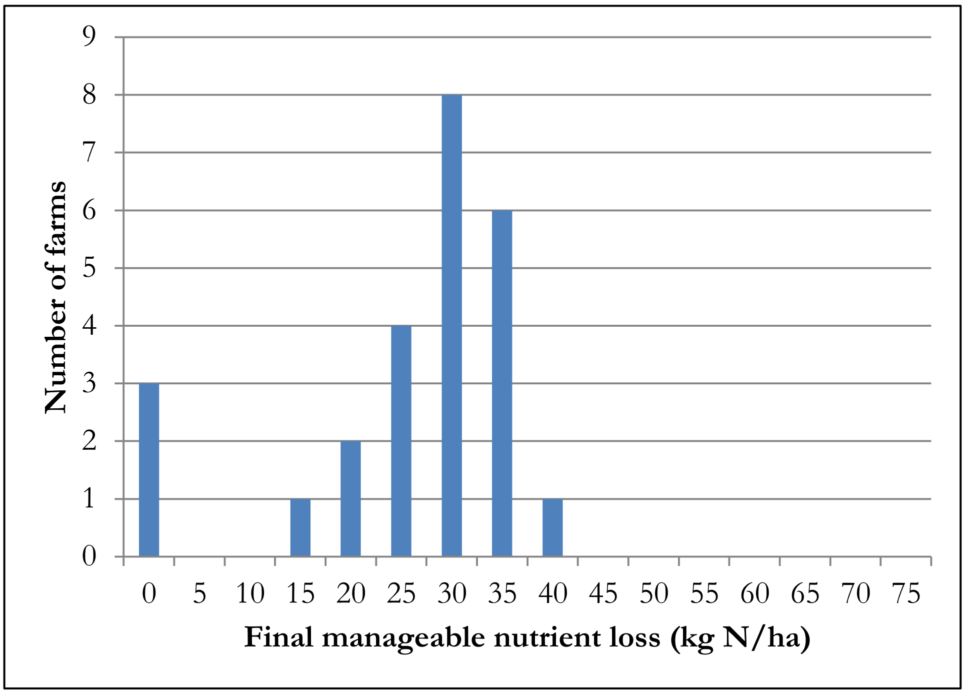

Figure 5.5 The final (year 20) distribution of manageable nutrient loss among dairy farms.

Figure 5.5 displays the histogram of dairy sector nutrient exports in the final simulation year (and for ease of comparison it uses the same horizontal scale as the baseline histogram in Figure 5.3). The final distribution has a much narrower range, suggesting that those who lose more nutrients initially tend to mitigate more. The amount of mitigation performed (per unit land area) therefore varies widely across farms. In our simulations, only a handful of properties convert fully to forestry (indicated by the bar at zero manageable nutrient loss). The amount of nutrient loss on most other farms indicates that they likely remain dairy farms (albeit highly mitigated dairy farms). The cost of achieving this mitigated state varies across farm properties as shown by Figure 5.6.

[20] The initial allocation may affect the mitigation response if trading allowances is not costless. It is also possible that the sale of allowances provides significant capital income to some farmers, allowing them to overcome capital constraints that affect their behaviour without free allocation.

[21] That is, the specification of profit functions does not affect the relative impact a farmer experiences from the choice of an allocation method, so the results in Figure 5.4 are robust to most simplifying assumptions we have made in modelling.

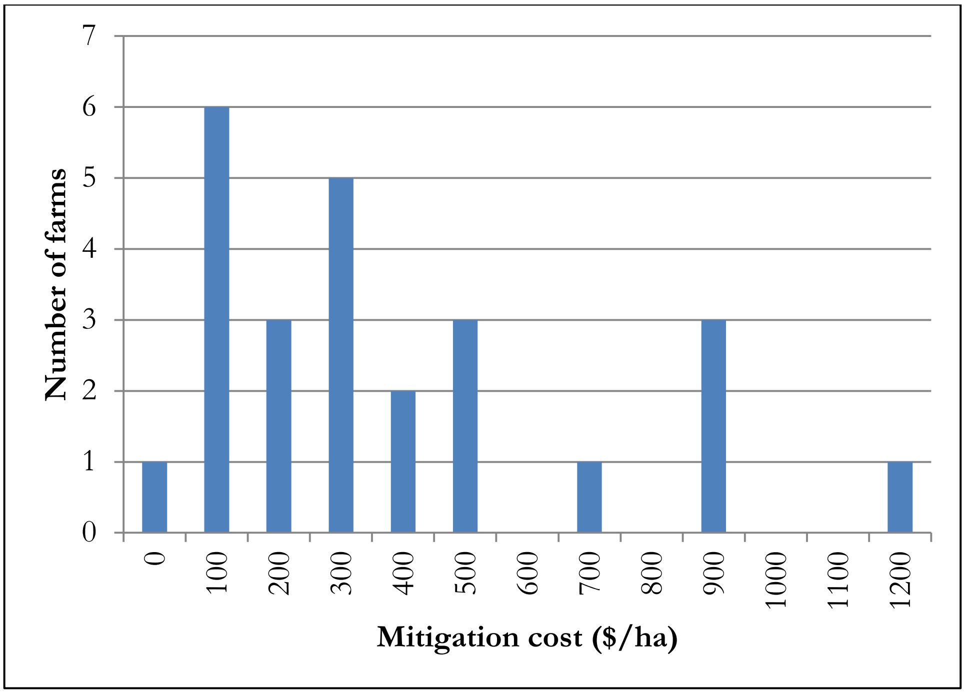

Figure 5.6 The final distribution of mitigation costs among dairy farms.

Simulated mitigation costs are sensitive to the methodology we employ in constructing the profit functions. Due to the sometimes strong assumptions we are required to make in the process, the precise value of mitigation cost to certain farmers is probably of less interest than the general shape of the distribution. In our simulations, some dairy farmers face a very high initial marginal cost of mitigation – at the simulated market price of nutrient allowances, these farmers perform very little mitigation directly (and hence they have low total mitigation costs). Farms that mitigate little tend to be associated with low levels of observed nutrient exports, so it is not implausible to think that these farms already follow best practice and hence their marginal mitigation cost are high (to be more precise, nutrient loss is only one side of the coin: marginal mitigation cost also depends on the nutrient efficiency of the farm). On the other hand, some farmers with high nutrient loss perform a lot of mitigation and face high total mitigation costs per hectare. While the specific form of the quadratic profit functions employed in NManager may contribute to some of these findings, the simulated cost distribution underscores that the sector is highly heterogeneous and that a diverse range of mitigation actions can be expected to be undertaken by dairy farmers.

Next, we examine the distribution of total costs. These differ from mitigation costs by the net cost of allowances. To illustrate the impact of free allocation we first consider a hypothetical scenario with no free allocation (Figure 5.7). Here, allowances can be thought of as being auctioned off to the highest bidder.

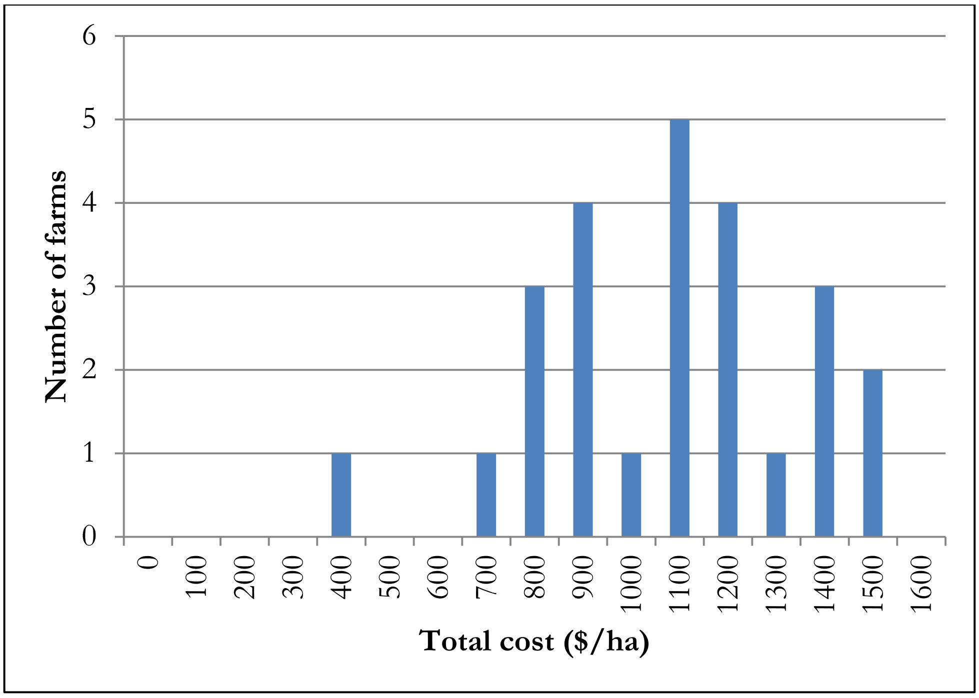

Figure 5.7 The final distribution of total costs with no free allocation.

Free allocation reduces total costs to dairy farmers significantly. This is not surprising, and it is easily verified by comparing the impacts in Figure 5.7 and Figure 5.8.

The method of free allocation has a systematic effect on the distribution of total costs in our results. Figure 5.8 suggests that the grandparenting approach leads to more equal cost sharing within the dairy sector (i.e., it is associated with a narrower range of impacts). That is, grandparenting better satisfies the ‘equal sharing’ principle.[22] Intuitively, the reason for this is that more allowances are allocated under grandparenting to those who tend to mitigate more – farmers who bear higher total mitigation costs therefore get a higher benefit from free allocation. From this perspective, the grandparenting approach can be thought of as providing compensation to those who undertake more mitigation. On the other hand, if the reason these famers mitigate more is that they had performed less mitigation in the past (or that they followed more intensive farming practices), then the more equal distribution of impacts under grandparenting may not be politically desirable. A key issue in designing an allocation mechanism is therefore to determine the source of differences in benchmarked nutrient discharge rates.

[22] However, equal cost sharing is not necessarily equitable cost sharing.

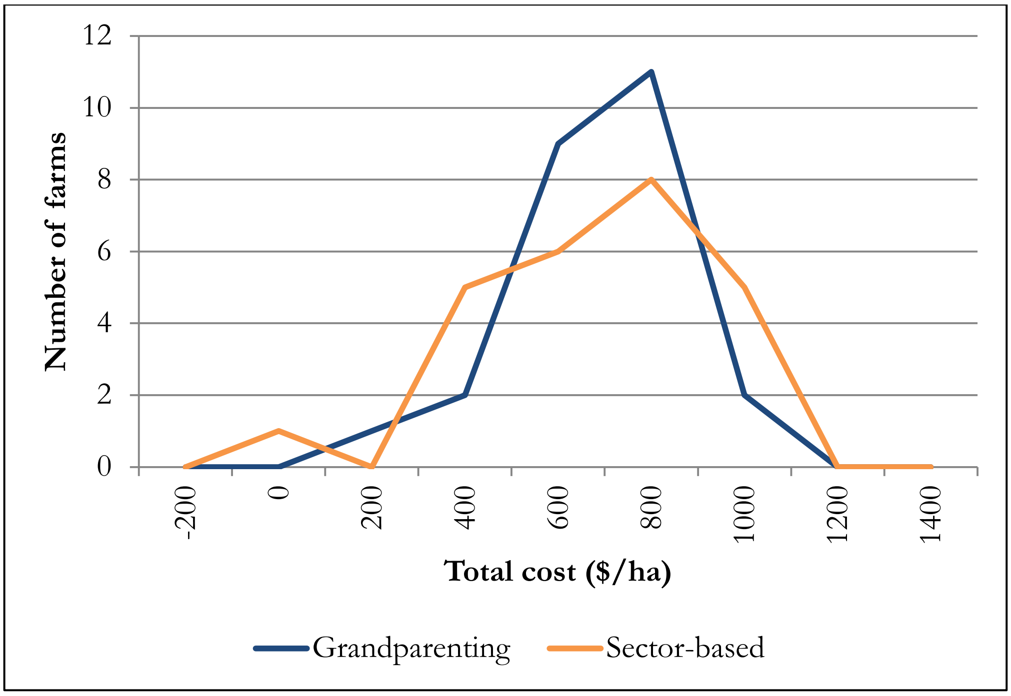

Figure 5.8 The final distribution of total costs with free allocation.

The finding of more equal cost sharing under grandparenting is sensitive to parameterization and is not necessarily generalisable. It only holds if those with higher initial nutrient loss bear higher total mitigation costs. This is clearly the case for dairy farms in our simulations of the final target, as shown by Figure 5.9.

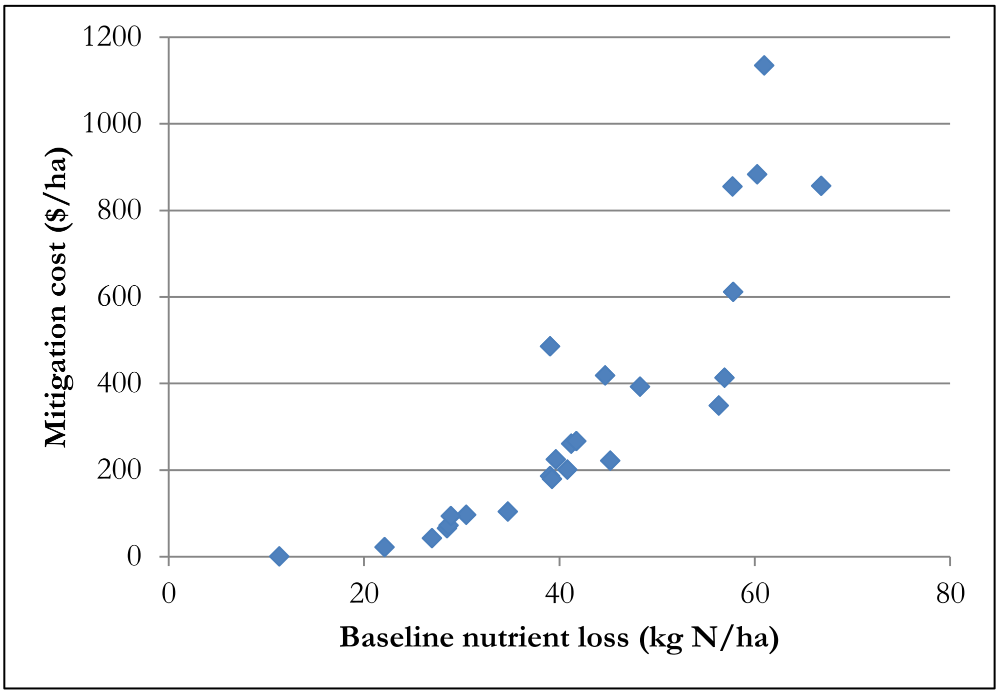

Figure 5.9 The relationship between baseline nutrient loss and final mitigation cost among dairy farms.

In addition to mitigation cost, the total cost of the policy also includes the net cost of allowances to farmers – unlike mitigation cost, this component varies with allocation in our modelling. It can be seen from Figure 5.10 that those who pollute more initially face a higher burden under sector-based averaging than under grandparenting, and vice versa.[23] That is, farmers with high rates of historical nutrient loss are held responsible to a larger extent under a sector-based allocation approach. Therefore, if 'polluter pays' is defined on the basis of initial nutrient loss, sector-based averaging could be considered more strictly consistent with the principle.[24]

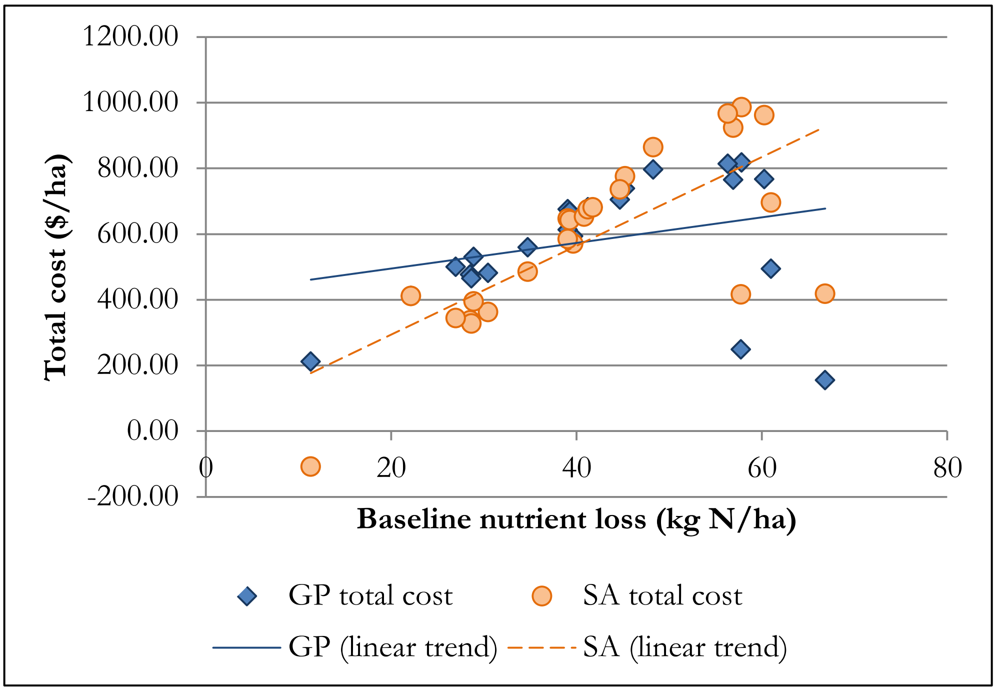

Figure 5.10 The relationship between baseline nutrient loss and total policy cost among dairy farms with grandparenting (GP) and sector-based averaging (SA).

We have explained that at the aggregate level, there is no difference to the sector between the two allocation scenarios. The wide distribution of impacts in Figure 5.4 indicates that the manner in which allowances are allocated among dairy farmers may have a great effect on cost sharing within the sector. In particular, on a per-hectare basis, the impact of allocation is significant in relation to baseline profits and total policy costs for some farmers. For example, the simulated impact is in some cases as large as one fifth or even one quarter of baseline profits. These farmers can be expected to express strong preferences for a particular allocation mechanism.

Are the winners and losers of grandparenting (relative to sector averaging) different in any systematic way? We know that their observed baseline nutrient losses are different, but what factors might explain these differences? As the final step of our investigation of the dairy sector, we explore how the simulated impact of allocation varies with selected observed farm characteristics. While the questions we pose above should be of interest to policy makers, we do not claim to provide the answers. The main purpose of our admittedly simplistic descriptive analysis is to illustrate the 'responsibility for action' principle. Figure 5.11 and Figure 5.12 plot the allocation impact against milk production and mean annual rainfall, respectively.

[23] Total cost is relatively low under both approaches for three of the highest-leaching farms. As suggested by figure 5.5, these farms convert to forestry use in our simulations: they face high mitigation costs, but are then able to sell their allowances reducing their overall cost burden.

[24] There is a much weaker relationship between the total cost burden and final nutrient loss.

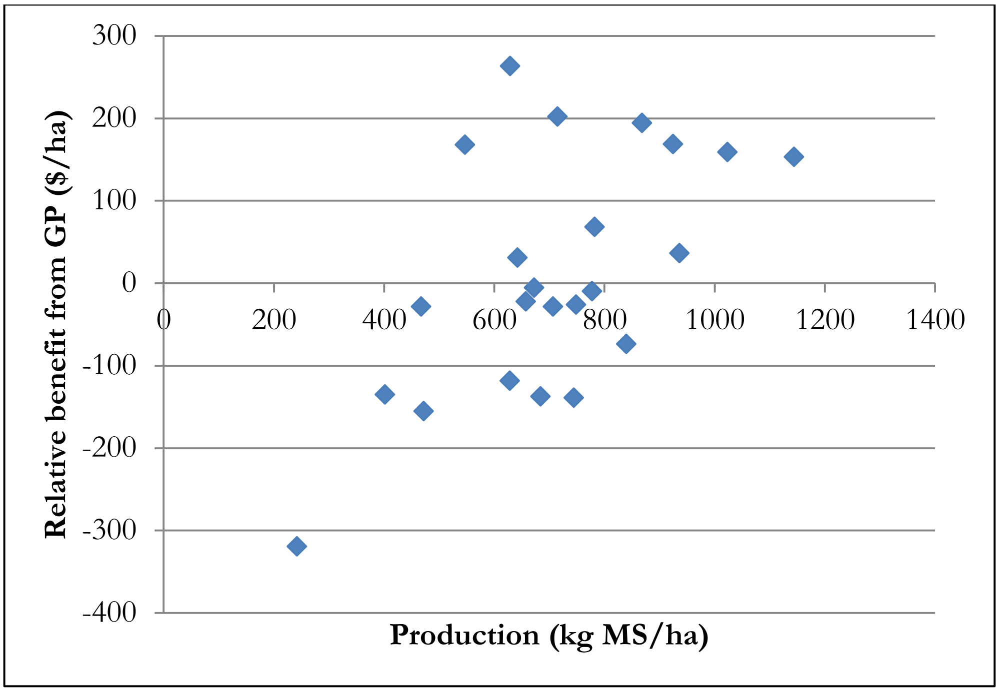

Figure 5.11 The variation of allocation impacts with milk production.

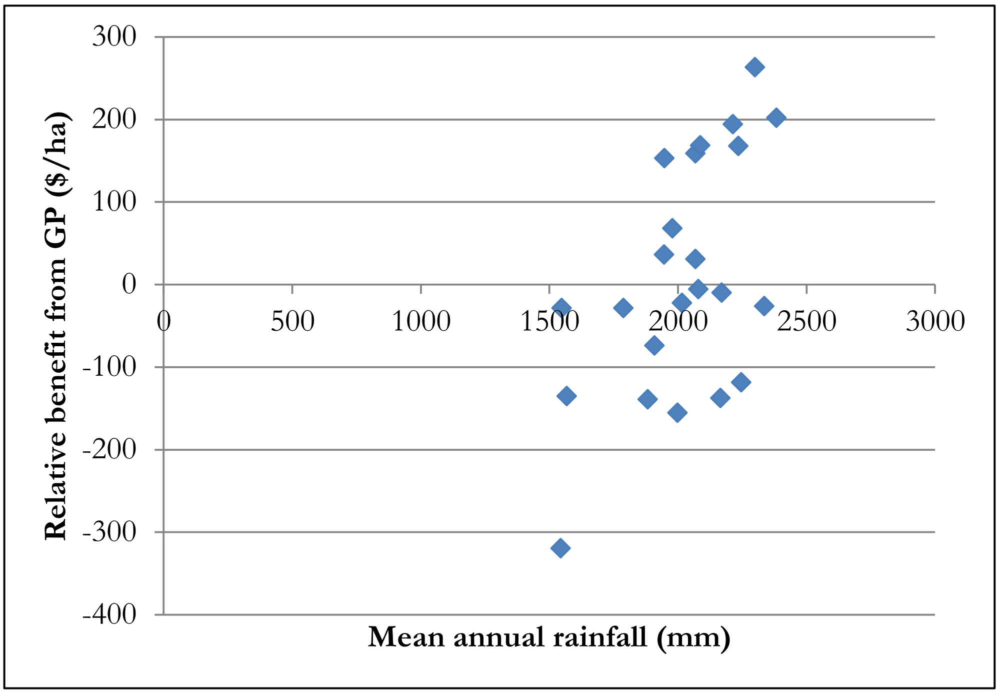

Figure 5.12 The variation of allocation impacts with mean annual rainfall.

Although there is a lot of variation in these figures, it appears that there may be a positive relationship between milksolid production and the dairy farmer's preference for grandparenting. Likewise, there may be a positive relationship between rainfall and the farmer's preference for grandparenting.[25] These relationships would suggest that farmers who produce more milk as well as farmers whose property receives more rain tend to have higher-than-average levels of nutrient loss. They would therefore be better off under grandparenting.

25 In both cases, the slope of a regression line is significantly different from zero at the 1% level.

The significance of these observations lies in the fact that while farmers can manage the volume of production on their farm, they have no control over the amount of rainfall their property receives. This invokes the 'responsibility for action' principle for cost sharing. A potential argument for the grandparenting approach is that it does not disadvantage farmers who have high rates of baseline nutrient leaching due to factors outside their control – for example, the geophysical and climatic environment of their property. On the other hand, a potential argument for the sector-based averaging approach is that it rewards farmers who have low rates of baseline nutrient leaching because of choices they had made prior to benchmarking – for example, forgoing land-use intensification and hence profits from additional milk production.

A key consideration in designing an allocation mechanism is therefore the source of differences in benchmarked nutrient discharges. If a significant proportion of these differences can be attributed to geophysical factors exogenous to farm management, then it may be argued that grandparenting some portion of allowances is justified to ease the burden on farmers who own land that is more prone to losing nutrients.

5.3 Dry Stock Sector Results

This subsection considers simulation results and the impact of allocation for dry stock farms in the study area. It closely parallels our previous analysis of the dairy sector. Accordingly, unless the conclusion is different than before, we provide less discussion here.

Figure 5.13 shows the distribution of initial nutrient loss per hectare on farms we classify as dry stock. The initial distribution is of interest because it determines the gains from having allowances allocated via one method over the other. In dry stock farming too, farmers whose nutrient loss is below the sector's average (denoted by a dashed vertical line in Figure 5.13) would benefit from the sector-based allocation approach, and vice versa. This is illustrated in Figure 5.14.

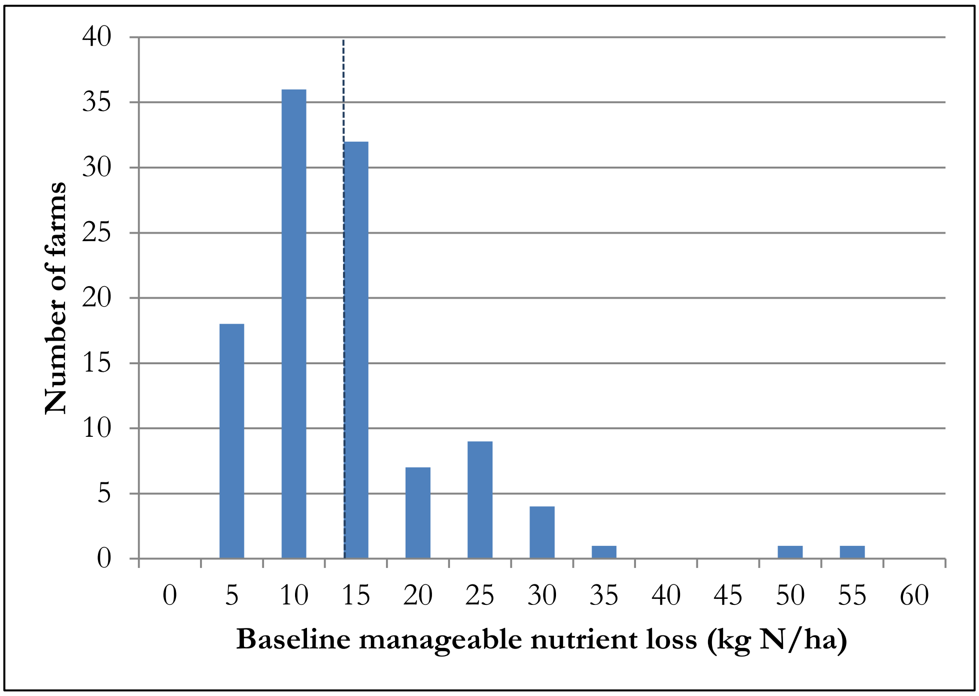

Figure 5.13 The distribution of baseline manageable nutrient loss among dry stock farms.

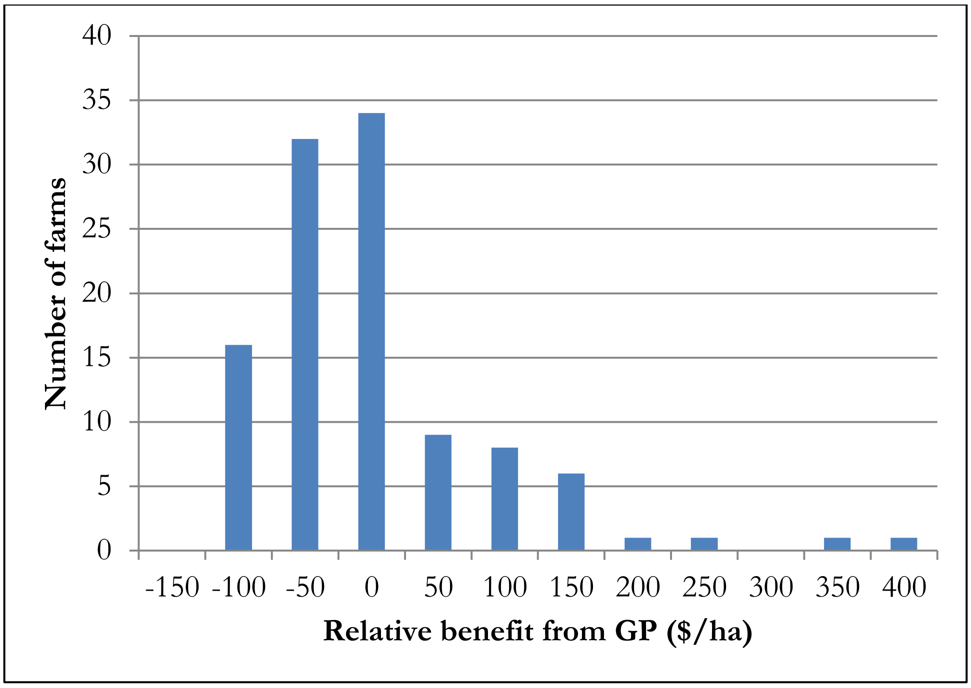

Figure 5.14 The final distribution of allocation impacts among dry stock farms.

The method of allocation is important to many dry stock farmers, as shown by the distribution of allocation impacts in Figure 5.14. (For reference, note from Table 5.1 that the mean mitigation cost is $127 per hectare and the mean policy cost is $27 per hectare within the sector.)

We project that dry stock farms mitigate most of their manageable nutrients. Many of them fully convert to forestry by the final simulation year, as indicated by the bar at zero (manageable) nutrient loss in Figure 5.15.

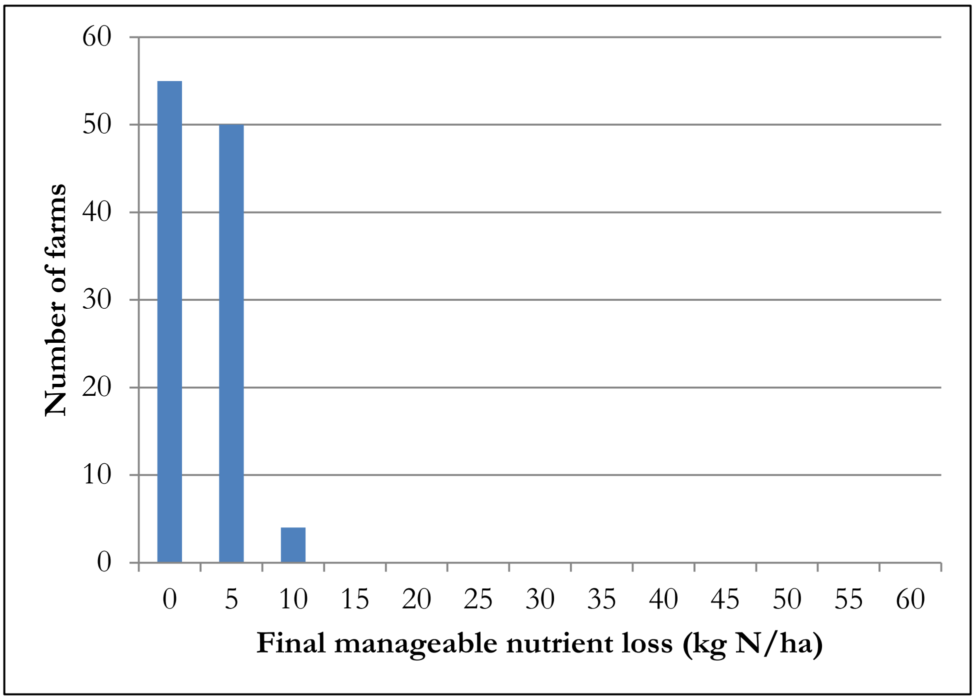

Figure 5.15 The final distribution of manageable nutrient loss among dry stock farms.

While simulated mitigation behaviour obviously depends on the manner in which we parameterise profit functions, the particular assumptions we make actually matter little for the above finding. To verify this, consider the co-variation of profit and nutrient loss across the sectors. On a per hectare basis, the mean profit of dairy farms is more than five times higher, while their mean nutrient loss is only about three times larger. This necessarily implies that mitigation, including land-use based mitigation, will be relatively cheap on dry stock farms. In particular, forestry is a viable alternative to some dry stock farmers: profits in forestry use (under full mitigation) may not be much lower than initially. This contributes to the large mitigation response of the sector.

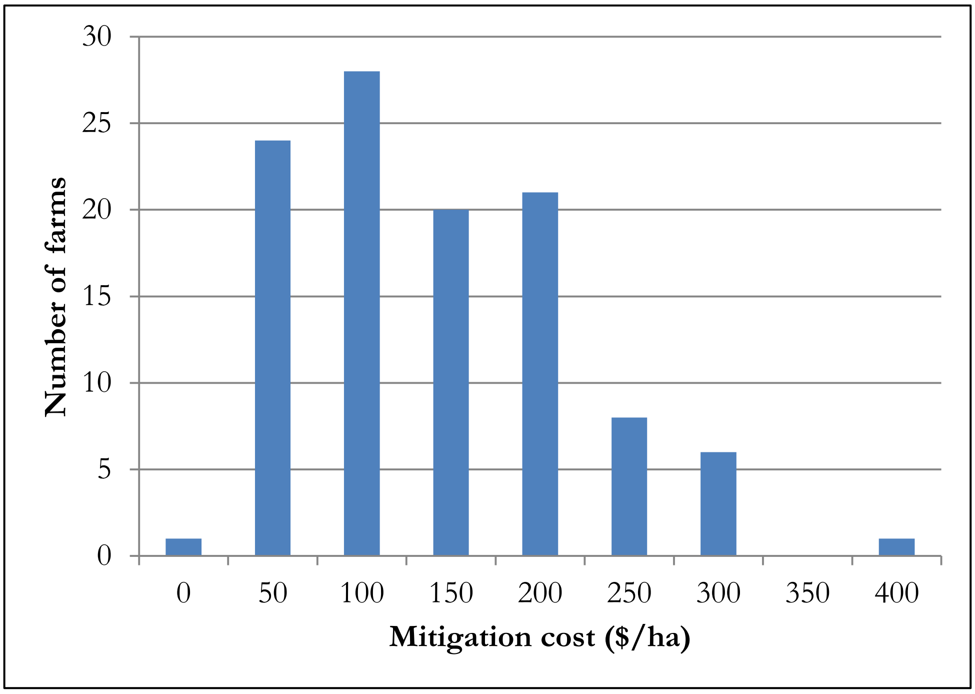

Figure 5.16 The final distribution of mitigation costs among dry stock farms.

The distribution of (total) mitigation costs is shown in Figure 5.16. Compared to dairy farming, mitigation costs are much lower in the dry stock sector – they rarely reach $300–400 per hectare, whereas in dairying they may exceed $700–1200 per hectare.

Next, we also present the distribution of total costs (where total cost is defined as the cost of mitigation plus the net cost of allowances) – first in the absence of free allocation (Figure 5.17), and then under the two free allocation approaches (Figure 5.18).

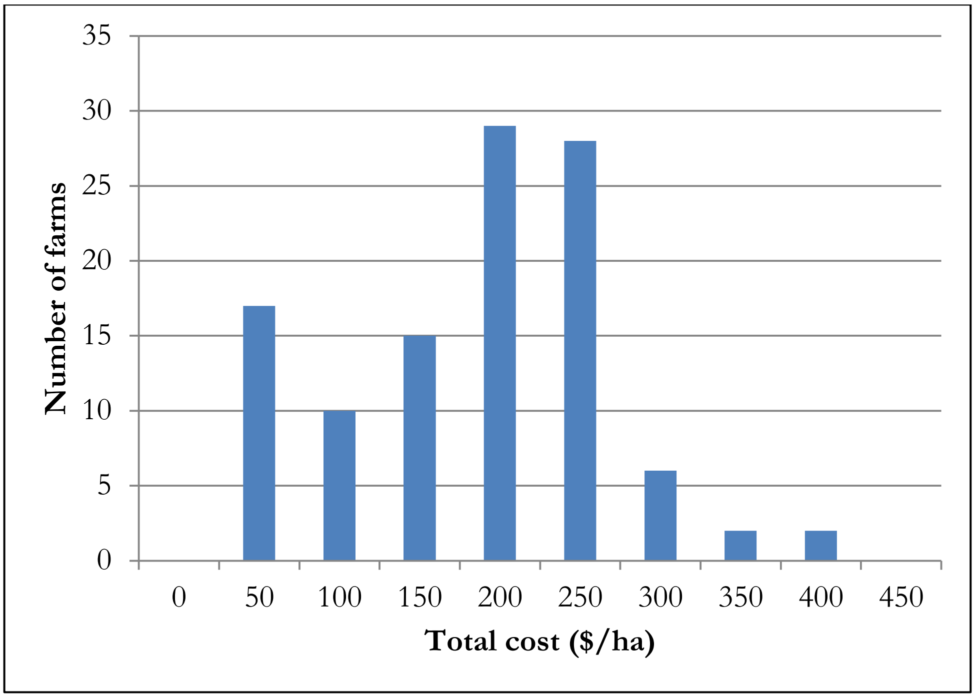

Figure 5.17 The final distribution of total costs with no free allocation.

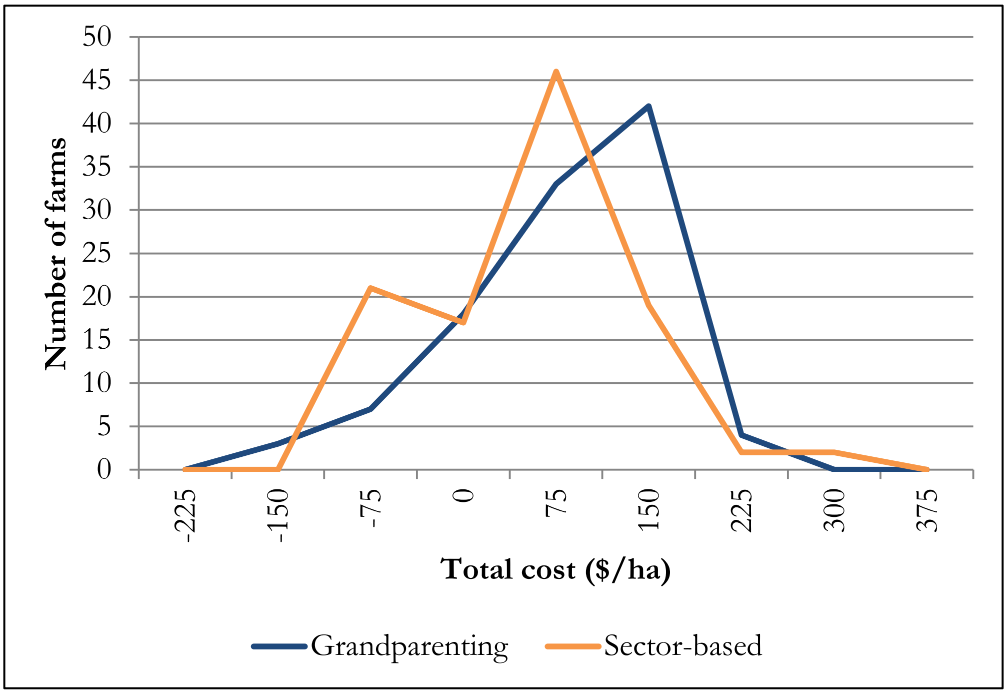

Figure 5.18 The final distribution of total costs with free allocation.

As in dairy farming, the method of free allocation affects the simulated distribution of total costs. However, it is not clear that the grandparenting approach leads to more equal cost sharing here: it does not obviously reduce the range of impacts experienced by farmers. The implication is that dry stock farmers who have high initial nutrient loss do not necessarily have high mitigation costs. This is possible if higher leaching is not necessarily associated with higher profit within the sector – a plausible proposition given the large heterogeneity with respect to geophysical land attributes and farming methods among dry stock farms.

Our remarks on general principles for cost sharing in the dairy sector apply here as well. In general, however, it might be more difficult to identify winners and losers based on observed farm characteristics in the dry stock sector. This is because of the large amount of heterogeneity that exists within the sector.

6.0 CONCLUSION

We use the NManager simulation model to compare two approaches, grandparenting and sector-based averaging, for the free allocation of allowances in a nutrient trading system with heterogeneous dairy and dry stock farmers. Where possible, we parameterise the model to data on farms in the Lake Rotorua catchment.

The distribution of relative allocation impacts experienced by farmers (or landowners) depends fundamentally on the distribution of benchmarked nutrient losses around the sector-level mean. This follows from the mechanism of free allocation under the alternative approaches. Therefore, results for relative outcomes under the two approaches are robust to the assumptions we make in parameterising profit functions. The nature of the simulation model suggests that most other results should be viewed as illustrative rather than predictive.

The cost of meeting the BoPRC's nutrient target for the lake is high. Our simulations suggest that meeting the final target requires the conversion to forestry of many dry stock farms, and the implementation of farm management techniques that could be labeled 'best practice' on most dairy farms. Although farmers might be able to pass on some of the costs associated with these mitigation responses to customers or workers, their ability to do so would likely be limited because other catchments do not have nutrient regulations.

Under both free allocation approaches, the sale of allowances provides net revenue to dry stock farms – in effect, through the sale of allowances, dry stock farmers are paid to perform mitigation for the dairy sector. More heterogeneity among farmers increases the potential to benefit from trade, and hence it could reduce the overall cost of compliance.

Free allocation significantly reduces the costs to each sector. But the method of allocation (grandparenting or sector-based averaging) has no bearing on the costs experienced at the sector level. Therefore, neither sector (as a whole) will argue for one approach over the other. On the other hand, landowners in both sectors will experience a wide range of impacts from the choice of an approach, and are therefore expected to express strong individual preferences for a particular allocation mechanism.

Farmers with higher-than-average nutrient loss are unambiguously better off under grandparenting, and farmers with lower-than-average nutrient loss are unambiguously better off under the sector-based averaging approach. The financial impact depends on the market value of allowances. Numerical results indicate that for many farmers, the difference between the allocation approaches can reach hundreds of dollars per hectare – this is substantial in relation to estimated profits.

For mixed-use farms, our results are sensitive to the manner in which we define sectors. In particular, mixed farms are assigned to a particular sector, and this assignment determines their sector-based allocation. Unless the allocation mechanism is refined to account for multiple land uses on a farm, this issue may pose difficulties not only in modelling but in an actual policy environment as well. Therefore, when mixed farms are common, sector-based allocation approaches may be impractical. Grandparenting is not affected by this issue.

To motivate the interpretation of results, we discuss some of them through the lens of alternative cost sharing principles: the ‘equal sharing’ principle, the ‘polluter pays’ principle and the ‘responsibility for action’ principle. All three represent valid perspectives that could be taken into consideration during the design of an allocation mechanism. In addition, other principles may further qualify the arguments for either allocation mechanism.

We find that the grandparenting approach leads to more equal final cost sharing within the dairy sector. This result is potentially sensitive to parameterization, and it is not necessarily true in all circumstances: for example, the relationship appears to be less strong in simulations of the dry stock sector. In fact, the set of farms that bear the highest mitigation costs may change with the nutrient price. It is possible, for instance, that as the nutrient price rises, mitigation shifts from on-farm management toward land-use change, and that the farms that engage in these responses are different. This suggests that the relationship between cost sharing and allocation could be ambiguous even within a given sector.

Farmers with high rates of historical nutrient loss face a higher burden under a sector-based averaging allocation. Therefore, when the principle is defined on the basis of baseline nutrient loss, this approach is more strictly consistent with ‘polluter pays’.

Because the variation in benchmarked nutrient loss ultimately determines the variation in relative allocation impacts, the source of differences in nutrient discharge rates within each sector should be of interest to policy makers. A potential argument for grandparenting allowances is that this approach does not disadvantage farmers who have high rates of baseline nutrient loss due to factors outside their control – for example, due to geophysical and climatic conditions on their farm. Conversely, a potential argument for sector-based allocation is that it provides a higher reward to farmers who have low nutrient loss because they have already undertaken some mitigation. These arguments illustrate the ‘responsibility for action’ principle.

Although we do not model one explicitly, it should be noted that mixed approaches to allocation are feasible. For example, allowances could be allocated by grandparenting only up to the amount of the sector’s average nutrient loss; remaining allowances could then be auctioned off. It is also possible to gradually shift from grandparenting to a sector-based allocation (or to another different approach) over time. If a significant proportion of nutrient leaching can be attributed to factors outside farmers’ control, then permanently grandparenting some portion of allowances may be desirable. The approach taken could also be different across the two sectors: in our application, sector level-outcomes do not depend on the method of free allocation, so it is possible to devise sector-specific allocation methods in order to meet desired cost sharing principles. More generally, the same principles and issues could also be applied to allocation between sectors.

Several other important considerations could affect the design of an allocation mechanism. Throughout our discussion, we assume that mitigation responses are not affected by the amount of free allocation received. This may not be true if, for example, the trade of allowances is costly, farmers face capital constraints, or there are uncertainties associated with regulation (Fowlie & Perloff, 2013). Trading may not be costless, especially in a relatively small catchment where the supply and demand for allowances is limited, and especially initially, while participants learn to use the system. Under such circumstances, allocating allowances to those who use them reduces the costs associated with trading (and increases the efficiency of the system if these costs would have prevented some participants from trading). It would, for example, be possible to argue in this case that increasing the allocation of dairy farmers may improve short-run cost-effectiveness.

It is also possible that the sale of allowances provides significant capital income to some farmers, allowing them to overcome capital constraints that otherwise affect their behaviour. In this case, too, the allocation could potentially affect mitigation behaviour. This would support a possible equity argument of giving allowances to those who will optimally incur high mitigation costs. Likewise, existing irreversible investments (sunk costs) that affect nutrient loss may affect the attractiveness of one approach over the other.

The allocation of nutrient discharge allowances will always be politically contentious because the potential transfer of wealth is large. Allocation has a major impact on ultimate cost sharing, and in some cases it is easy to identify winners and losers from the choice of an allocation method. However, final cost sharing may not be immediately obvious from direct consideration of the free allocation mechanism because another key determinant is the farmer's ability to mitigate. Our modelling illustrates differences between the two allocation methods and illuminates the scale of fairness issues that will be dealt with in gaining acceptance in each sector. It also identifies some practical issues that may need to be addressed in working toward a final allocation scheme with stakeholders. This is important to ensure that stakeholders feel key principles have been addressed and unintended outcomes are avoided.

7.0 ACKNOWLEDGEMENTS

Funding for this project was provided by the Bay of Plenty Regional Council.

8.0 REFERENCES

Allan, Corey and Suzi Kerr. (forthcoming). "Show me the profits: Examining the relationship between agricultural profitability and rural land values in New Zealand over time," Motu Working Paper (draft), Motu Economic and Public Policy Research, Wellington.

Anastasiadis, Simon and Suzi Kerr. 2012. "Mitigation and Heterogeneity in Management Practices on New Zealand Dairy Farms", draft paper prepared for the NZARES conference 2012, Motu Economic and Public Policy Research, Wellington.

Anastasiadis, Simon; Marie-Laure Nauleau; Suzi Kerr; Tim Cox and Kit Rutherford. 2011. "Does Complex Hydrology Require Complex Water Policy? NManager Simulations for Lake Rotorua," *Motu Working Paper* 11-14, Motu Economic and Public Policy Research, Wellington.

Duhon, Madeline; Hugh McDonald and Suzi Kerr (forthcoming). "Nitrogen Trading in Lake Taupo: An evaluation of innovative water quality policy", Motu Working Paper (draft), Motu Economic and Public Policy Research, Wellington.

Fowlie, Meredith and Jeffrey M. Perloff. 2013. "Distributing Pollution Rights in Cap-and-Trade Programs: Are Outcomes Independent of Allocation?" *Review of Economics and Statistics*, vol. 95, no. 5, pp. 1640-1652.

Kerr, Suzi; Catherine Leining and Justine Sefton. 2012. "Roadmap for Implementing a Greenhouse Gas Emissions Trading System in Chile: Core Design Options and Policy Decision-Making Considerations," *Motu Working Paper* 12-14, Motu Economic and Public Policy Research, Wellington.

Kerr, Suzi and Kelly Lock. 2009. "Nutrient Trading in Lake Rotorua: Cost Sharing and Allowance Allocation," *Motu Working Paper* 09-09, Motu Economic and Public Policy Research, Wellington.

Landcare Research and MAF. "Land Use Capability (LUC) Map 2002," obtained by Motu Research 2002. Restricted dataset. Details online at http://www.motu.org.nz/building-capacity/datasets.

Ministry of Agriculture and Forestry. "Monitor Farms Underlying Unit Record Data 2007/2008 - 2009/2010," obtained by Motu Research 2010. Restricted dataset. Details online at http://www.motu.org.nz/building-capacity/datasets.

National Institute of Water and Atmospheric Research. "ROTAN Land Use Data," obtained by Motu Research 1-1-2011. Restricted dataset. Details online at http://www.motu.org.nz/building-capacity/datasets.

Perrin Ag Consultants Ltd and AgResearch. 2012. "Farmer Solutions Project," Report prepared for the Bay of Plenty Regional Council.

Rutherford, Kit; Chris Palliser and Sanjay Wadhwa. 2009. "Nitrogen Exports from the Lake Rotorua Catchment – Calibration of the ROTAN Model," National Institute for Water and Atmospheric Research, Hamilton.

Rutherford, Kit; Chris Palliser and Sanjay Wadhwa. 2011. "Prediction of Nitrogen Loads to Lake Rotorua Using the ROTAN Model," *NIWA Client Report HAM2010-134*, National Institute for Water and Atmospheric Research, Hamilton.

Smeaton, Duncan; Tim Cox; Suzi Kerr and Robyn Dynes. 2011. "Relationships Between Farm Productivity, Profitability, N Leaching and GHG Emissions: a Modelling Approach," 73, *New Zealand Grasslands Association*, pp. 57-62.

www.gns.cri.nz

Principal Location

1 Fairway Drive
Avalon
PO Box 30368
Lower Hutt
New Zealand
T +64-4-570 1444
F +64-4-570 4600

Other Locations

Dunedin Research Centre
764 Cumberland Street
Private Bag 1930
Dunedin
New Zealand
T +64-3-477 4050
F +64-3-477 5232

Wairakei Research Centre
114 Karetoto Road
Wairakei
Private Bag 2000, Taupo
New Zealand
T +64-7-374 8211
F +64-7-374 8199

National Isotope Centre
30 Gracefield Road
PO Box 31312
Lower Hutt
New Zealand
T +64-4-570 1444
F +64-4-570 4657