Multifunctional Analysis of Spatially Targeted Environmental Policy

Leppert, Daniel Jakob

Multifunctional Analysis of Spatially Targeted Environmental Policy

Daniel Leppert

of

University College

A thesis submitted in fulfilment of the requirements for the degree of Doctor of Philosophy

Department of Economics

Durham University

United Kingdom

November 11, 2025

Abstract

Growing tensions between economic priorities and protection of nature highlight the importance of cost-effective environmental policies. Amidst mounting climate impacts and higher inflation, policymakers around the world are working to meet environmental objectives while limiting the burden on taxpayers. There are important spatial dimensions to many critical environmental problems, including air pollution, flooding, and pollinator declines. This thesis demonstrates that adverse incentives may jeopardise the effectiveness of environmental policy when geographic conditions allow firms to export pollutants beyond the regulator's jurisdiction. Using a custom air pollution dispersion model, this work calculates the interstate SO_2 pollution from coal-fired power plants across the United States between 1997 and 2020. It exploits a natural experiment to show that firms exporting pollutants beyond the regulator's jurisdiction respond less to a tightening of emission caps. The following research explores so-called spatially targeted policies that seek to account for heterogeneous policy impacts in different geographies. The focus of this thesis is environmental land management (ELM) schemes that compensate farms to retire cultivated land. It advances a novel multifunctional cost-effectiveness analysis of hypothetical schemes by combining cost estimates via discrete choice experiments (DCEs) with benefit estimates from hydrological and ecological models. This thesis demonstrates that tradable and spatially targeted ELM contracts are likely to deliver measurable improvements in both natural flood management and pollinator services. In addition, simulating multiple spatial configurations of ELM features illustrates how small, evenly distributed natural features may cost-effectively circumvent coordination costs among farms. This thesis demonstrates the value in integrating hypothetical DCEs with spatial simulation models.

To all animals — large and small

Contents

List of Tables

List of Figures

Table 1: *List of abbreviations*

ARP	Acid Rain Program
ATT	Average treatment effect on the treated
BIC	Bayes information criterion
CAA	Clean Air Act (of the United States)
CAIR	Clean Air Interstate Rule
CSAPR	Cross-State Air Pollution Rule
CS	Countryside Stewardship scheme
DCE	Discrete choice experiment
DD	Difference-in-differences
DDD	Triple differences
Defra	Department for Environment and Rural Affairs (of the United Kingdom)
ELM	Environmental Land Management
EPA	Environmental Protection Agency (of the United States)
GIS	Geographic information system
KKT	Karush-Kuhn-Tucker conditions
LC	Latent class
MMNL	Mixed multinomial logit
MNL	Multinomial logit
NAAQS	National ambient air quality standards
NFM	Natural flood management
PSM	Propensity Score Matching
SFI	Sustainable Farming Initiative
SIP	State Implementation Plan
UK	United Kingdom
US	United States
WTA	Willingness-to-accept
WTP	Willingness-to-pay

Table 2: *List of Greek letters*

α	Farm output elasticity of non-land inputs
β	Farm output elasticity of land inputs
β_k	Coefficient for model parameter k
$\hat{\beta}$	Priors for taste parameters
γ	Pollinator dependence of crops
δ_r	Share of pollutants to region r
δ_s	Utility offset for latent class s
ϵ	Error term for linear model
λ	Endowment elasticity of demand for land
μ	Lagrangian multiplier (shadow cost)
π	Payment
τ	Transaction cost
ρ	Pollinator survival rate
σ	Standard deviation
ϕ	Connectivity insensitivity ratio ($V'(\ell)/V'(n)$)
A	Statistical significance cutoff
B	Statistical power cutoff
Δ	Excess emissions over emission cap
Ω	Variance-covariance matrix

Declaration and Copyright

I certify that this thesis is my own work. Ashar Aftab assisted in the survey design by advising on questionnaire items, DCE attribute selection and best practices. Dr Aftab also contributed to the development of the research agenda and provided feedback on the manuscripts. Riccardo Scarpa contributed econometric expertise, including advise on model selection and hypothesis testing, and feedback on manuscripts. Professor Scarpa also provided support in recruiting farmers for the DCEs, facilitated access to electoral records used to collect farm addresses, and managed relations with respondents and the farming community. Professor Scarpa also facilitated a number of in-person interviews. Sim Reaney consulted on the application and validity of `SCIMAP-Flood` for this work. All sources and materials used in the preparation of this thesis have been properly acknowledged and cited. No material contained in the thesis has previously been submitted for a degree at Durham University or any other institution.

Research for this thesis was partially funded by the Natural Environment Research Council (NERC) and the Economic and Social Research Council (ESRC) via grant [NE/W007495/1] Synthesizing evidence in the economics of farm environmental biodiversity.

Acknowledgements

Completing a doctoral thesis, while at times a solitary endeavour, is never the product of a single person. I would not have reached this stage without the people who supported me in various ways along the journey.

First, I wish to thank my supervisors, Ashar Aftab and Riccardo Scarpa. Your involvement and enthusiasm in this project have been invaluable. You have both had my back throughout. Thank you for promoting my better ideas, and for having patience with the ideas that did not go so well.

I am also grateful to my earlier mentors. In particular, Carl-Johan Lagerkvist at the Swedish University of Agricultural Sciences. Thank you for seeing potential in my work, for hiring me as a research assistant, and facilitating those important first opportunities to prove myself.

I am grateful to Anna and Anders, my parents, for giving me the freedom to explore, discover and grow at my own pace. Thank you for raising me in a home full of books, Lego, and love.

Thank you Zoe, my partner, for all your love and support through these years. A romantic relationship is no less complicated than the economic relationships covered herein, and you teach me something new every day.

Chapter 1

Introduction

1.1 Motivation

Europe and North America are experiencing growing political pressure to reduce the economic burden of environmental regulation. During periods of high inflation, the price s of energy and food products sit at the heart of businesses' and voters' concerns about their economic well-being. Price inflation is perceived as an unambiguously negative phenomenon, often attributed to government policies that directly or indirectly concern the environment. Examples include emission taxes (Ewald et al., 2022), bans on fracking for natural gas (Joskow, 2013), and mandates on environmental land management for food producers (Holstead et al., 2017). Inflation has also been found to deepen political polarisation, with conservatives being more likely to emphasise the role of government regulation as the cause (Binetti et al., 2024).

Against this backdrop of public unease about higher costs for essential energy and food, countries across the world are facing growing exposure to environmental risks. A business-as-usual scenario for Europe projects flood damages of \$30-\$60 billion annually by the year 2100. For 2020, the Association of British Insurers reports £817 million in flood-related losses for the UK alone (Bates et al., 2023). Additionally, absolute costs from flooding are increasing as a result of agricultural land use and economic development on flood-prone land (Dottori et al., 2023). Crops valued at \$195-\$387 billion globally (Porto et al., 2020) are increasingly at risk due to declining populations of wild pollinating insects (Powney et al., 2021). Humans have benefited from the fragile symbiotic relationship between pollinators and flowering fruits, nuts and berries since our hunter-gatherer ancestors. Today, the use of animal pollinated biofuel crops is growing, with the cultivation area of oilseed rape, sunflowers and soybeans increasing by 32% across Europe between 2005 and 2010 (Breeze et al., 2015). Loss of natural habitats, resulting from intensification and expansion of agriculture, have been contributing to these declines (Xiao et al., 2016).

Damage to the environment also has direct consequences for human health. Chay and Greenstone (2003b) attribute the marginal milligram of particulates per m^3 of air to 4-8 infant deaths per 100,000 live births. Since the Chay and Greenstone (2003b) study was conducted, environmental regulation, such as pollution permits, has contributed to significant improvement in air quality across the U.S., and lives saved as a result.

I am emphasising these environmental risks not to delegitimise concerns about regulations and their potentially inflationary effects. For example, the European tradable permit scheme (ETS), set up to regulate carbon emissions, has been found to fuel inflation in the EU (Känzig, 2023), to the detriment of households and firms. Harding et al. (2021) found that conservation zones restricting where agricultural firms were allowed to clear old growth forest did not only raise the price of agricultural outputs. They also observed secondary effects, where unprotected forests suffered more intense deforestation. I want to illustrate how there are real trade-offs between different groups whose voices all compete for the ears of policymakers.

In light of these trade-offs, I argue that the economics discipline, which has a long tradition of emphasising careful analysis of marginal costs and benefits, can contribute to a better understanding of the most cost-effective ways to achieve environmental goals. Failure to design policies that are efficient, targeted, and transparent will result in misallocation of public funds, and in further alienation of key stakeholders from the challenge of environmental protection.

This thesis is motivated by two observations. First, trade-offs between the costs of regulation and the environmental services at risk from economic activity demand exploration of more cost-effective policy designs. Second, ever-greater availability of high-resolution geography- and land use data enable spatially explicit estimation of how and where environmental benefits occur. Together, these observations invite research into *spatially targeted* environmental policy. The aim of such policy is to incentivise abatement of environmental damage in those places where the environmental costs are greatest.

1.2 Research statement

This thesis consists of three separate contributions to the research literature, each studying one of three different environmental problems, and each with an important spatial dimension. Chapter 2 studies sulphur dioxide emitted from coal-fired power plants and how air pollution can escape borders and hence regulation. Chapter 3 is a bridging chapter which sets up the data and empirical methodology that is shared between chapters 4 and 5. These chapters evaluate the multifunctional benefits of environmental land management (ELM) schemes. ELM schemes refer to policies obligating farmers to create natural features (e.g. planted trees, hedgerows, grass strips, or retirement of grassland from grazing) in exchange for a government payment. Chapter 4 explores how spatially targeted ELM contracts can mitigate flooding in downstream towns and villages, specifically by creating natural flood management features such as planted trees and regeneration. Finally, chapter 5 studies the cost-effectiveness of ELM schemes in terms of promoting insect pollination and conservation of habitats.

Each chapter analyses the environmental problem in focus through the same theory lens. Firms seek to minimise the cost of production while accommodating the demands of the market. The production process results in some environmental damage which is not (fully) suffered by the firm, but impacts the wider society. The *Polluter Pays Principle* is a key idea in environmental regulation, arguing that the party responsible for pollution should bear the cost of its environmental damage. On the issue of deciding the appropriate cost, Coase (1960) pioneered a long-standing and influential economics literature which showed how the government can achieve the optimal outcome for the whole society when property rights are defined, transaction costs are zero, and the marginal cost of abatement equals the marginal social cost from economic activity. That is, when the damage to the wider society, e.g. in the form of diminished air quality, associated with a single unit, produced via a polluting process, equals the market price of that unit.

However, the marginal cost is not always trivially estimated, and further complexities arise when the marginal cost curve is not the same across producers of environmental externalities. An important cause of differences in social cost is geography, which has been shown theoretically as early as Montgomery (1972). While many different settings have been explored theoretically, including air pollution with differentiated costs (Fowlie & Muller, 2019) and agricultural pollution in diverse catchments (Kampas et al., 2013), the empirical literature mostly focuses on low-resolution spatial differences in socioeconomic and demographic variables (Fowlie et al., 2012; Holland & Yates, 2015). Interdisciplinary research integrating environmental and geophysical modelling in economic cost-benefit analyses can inform design of spatially targeted environmental policy. This thesis addresses the following research questions:

Question I: A long-standing body of work has hypothesised that tradable pollution permits with a cap on overall emissions is an efficient policy mechanism to internalise the social cost of environmental damage from economic activity (Fowlie & Muller, 2019; Montgomery, 1972; Xepapadeas et al., 1997). Empirical studies have demonstrated effectiveness in terms of overall emissions of nitrogen oxide (Fowlie et al., 2012), sulphur dioxide (Schmalensee & Stavins, 2013) and carbon dioxide (Känzig, 2023). However, theoretical research has observed that cap-and-trade programs may result in emissions in excess of the cap if compliance is not enforced by the regulator (Stranlund & Chavez, 2000). This may be the case if emissions from a polluting firm, due to its location, are received outside the jurisdiction of the regulator, such as state- or local government. How do firms respond to spatially differentiated compliance costs and what is the resulting environmental impact?

Question II: Spatially targeted cap-and-trade programs have been proposed to address heterogeneous damages. Such programs introduce trading ratios, akin to exchange rates, that reflect the relative marginal damages between two firms that may trade in permits. In theory, this encourages greater abatement among high marginal damage sources, as these firms receive more money for each permit that they sell. Such a scheme is only optimal when the regulator has full information about the spatial distribution of marginal damages, so that trading ratios are assigned correctly (Holland & Yates, 2015). Agriculture has been identified as a sector where the spatial targeting of current policies is insufficient. While externalities such as pollutant runoff and habitat fragmentation frequently occur at the landscape scale, regulation via so-called environmental land management (ELM) typically only target the farm (Nguyen et al., 2022). What are the efficiency gains from spatially targeted permit trading over a non-targeted regime?

Question III: A cap-and-trade program, targeted or otherwise, involves the trading of permits among firms and therefore the possibility of transaction costs (Stavins, 1995, 2003). In the case of air pollutants from the energy sector, transaction costs have been empirically insignificant (Schmalensee & Stavins, 2013, 2017). However, evidence from the energy sector may not be directly applicable to trade in ELM contracts among farmers. Farms are frequently resource-constrained, and transaction costs have been identified as a barrier even in bilateral agreements between a farmer and the government agency (Peterson et al., 2015). Transaction costs have also been found to inhibit voluntary coordination between farmers (Banerjee et al., 2017). How do transaction costs impact the feasibility of a hypothetical market in ELM obligations?

Question IV: Spatially targeted ELM in agricultural regions can also be achieved via voluntary coordination, where farmers are incentivised with a bonus payment to coordinate land use change where it is most impactful (Kuhfuss et al., 2016; Parkhurst & Shogren, 2007). Such an agglomeration bonus stands in relation to the transaction cost involved for farmers, which inhibits collaboration (Banerjee et al., 2017; Nguyen et al., 2025). So far, few studies involving active farmers outside of the lab have focused on the determinants of transaction costs and how they can be reduced (Nguyen et al., 2022). In particular, what role does social- or professional networks play in farmers' perceived barriers to coordination?

Finally, the fifth research question is more practical and relates to the latter observation that inspired this thesis. Availability of high-resolution spatial data, increased computing power, and function packages made specifically for geospatial analysis in programming languages such as `R` and `Python`, open new research avenues. In particular, simulation of results from the aforementioned environmental policies:

Question V: Environmental objectives of ELM schemes such as runoff reduction or pollinator conservation are difficult to quantify (Bartkowski et al., 2021). Digital technologies, including GIS and remote sensing, are becoming part of the regulator's toolbox to aid compliance monitoring, data exchange and analysis (Ehlers et al., 2021). How can spatially explicit simulation models contribute to cost-benefit analysis of spatially targeted schemes? In particular, what can such models tell us about the trade-offs between variations in the coverage, type and spatial configuration of natural features?

1.3 Thesis outline

An outline for the remainder of this thesis is shown in figure 1.1. Its academic contributions are presented in three self-contained but connected chapters, in addition to chapter 3 which ties together the common methodological elements of the latter two. Each chapter begins with introduction and background sections, covering the relevant literature, policy environment, and research gaps. They are followed by explanations of the theoretical model, predictions, and hypothesis tests. I go on to present the results and discuss limitations, contributions, and policy recommendations.

Starting from the top of figure 1.1, chapter 2 is a quasi-experimental study into how polluting firms respond to a non-targeted cap-and-trade program. The policy in focus is the Clean Air Interstate Rule (CAIR), a market in permits for SO_2 emissions from coal-fired power plants. Announced in 2005, CAIR would cover a region of 26 eastern US states. The rule was later vacated after a court found that the non-targeted design of the program did not comply with the Clean Air Act provision to regulate interstate air pollution.

Using a model of non-targeted cap-and-trade with cost-minimising firms, I hypothesise that, in the absence of a credible mechanism to punish cross-border pollution, upwind sources respond less to reductions in the emission cap. I develop a custom air pollution dispersion model, GAUSSMOD, which allows me to attribute changes in SO_2 concentrations to individual power plants. I calculate the interstate SO2 pollution from 493 coal-fired power plants across the United States between 1997 and 2020.

In a difference-in-differences setup with plants not covered by CAIR in the control group, I estimate the treatment effect of the program on overall- and cross-border SO_2 emissions and find a 30% reduction in overall emissions but none in cross-border pollution. Instead, geographic factors rather than emission rates were the primary driver of interstate pollution. I report heterogeneous treatment effects where the reduction in overall emissions attributed to CAIR is lower among plants that transport emissions outside their state.

Chapters 4 and 5 both depart from the conclusions of chapter 2, which emphasise the risks of non-targeted cap-and-trade programs. Each chapter simulates hypothetical policies that address negative externalities arising from agricultural land use. These externalities, water runoff and habitat fragmentation, each depend significantly on local conditions and the spatial configuration of ELM features. Each chapter contributes a cost-benefit analysis of the respective ELM scheme in focus. In each case, the cost-side of the analysis is done using discrete choice experiments and the benefits are estimated with simulation models.

I conduct a discrete choice experiment (DCE) with a sample of farmers from the north of England. Respondents in the DCEs are asked to consider two different types of ELM schemes. Each hypothetical scheme involves the creation of natural features on the farm in exchange for a payment. The features include planted trees and regenerated vegetation arranged in different spatial configurations, ranging from in-field, disconnected patches to field-edge corridors. The first scheme introduces a market in ELM contracts, allowing farmers to trade obligations with trading ratios that reflect the flood mitigation potential of their land. The second scheme introduces a bonus payment for voluntary collaboration between farms. Neighbouring farmers may coordinate the placement of the ELM features. By observing respondents' choices from among a set of pre-defined schemes, I am able to elicit preferences for individual attributes. The DCE allows me to estimate how sensitive farmers are to variations in e.g. natural feature types, placement, transaction costs, and coordination demands.

Common across chapters 4 and 5 is the sample of surveyed farmers as well as the set of simulated ELM scenarios. In both chapters, results from the DCEs are used to estimate the required cost associated with ELM scenario. The amount of compensation demanded by farmers indicates the necessary cost to the government. This is the cost side of the cost-benefit analysis.

Also common across chapters 4 and 5 is the set of hypothetical ELM projects. The features of these projects are constructed using permutations of the DCE options. I develop an algorithm to simulate the land use change resulting from farmers' enrolment in each hypothetical project. The counterfactual land cover maps in each scenario are used as inputs to quantify environmental benefits in the next stage. As shown in the flowchart, the benefit estimations are unique contributions in each chapter.

In an effort to avoid repetition, these common elements are treated in chapter 3 which bridges the transition from chapter 2 to chapters 4 and 5. This part of the thesis, which is shown in the mid section of the flowchart in figure 1.1, does not contain any scientific results. Instead, it introduces the reader to the relevant policy background for the following two chapters and discusses the survey and sampling methodology, as well as the experimental design.

Chapter 4 introduces an economic model of farm behaviour when the regulator sets a catchment-wide cap on water runoff generated by agricultural land use. Farmers can trade ELM obligations with trading ratios that reflect the relative runoff generation potential at each farm. I use the hydrological model `SCIMAP-Flood` to simulate the trading ratios and the counterfactual flood risk reduction resulting from each ELM project with and without trading.

Chapter 5 explores the impact of spatially coordinating ELM features on crop pollination. I use a model to simulate bee foraging and population dynamics for each ELM scenario. Biological modelling allows me to compare changes in pollination services resulting from each spatial configuration of ELM features. I evaluate the benefit of coordination between farms, which facilitates improved habitat connectivity across the landscape. Combined with costs derived from the DCEs, simulated benefits make up the *multifunctional* cost-effectiveness analysis which is a novel contribution from this work. I discuss its relevance for how we think about cost-effectiveness of spatial targeting and the implications for environmental policy.

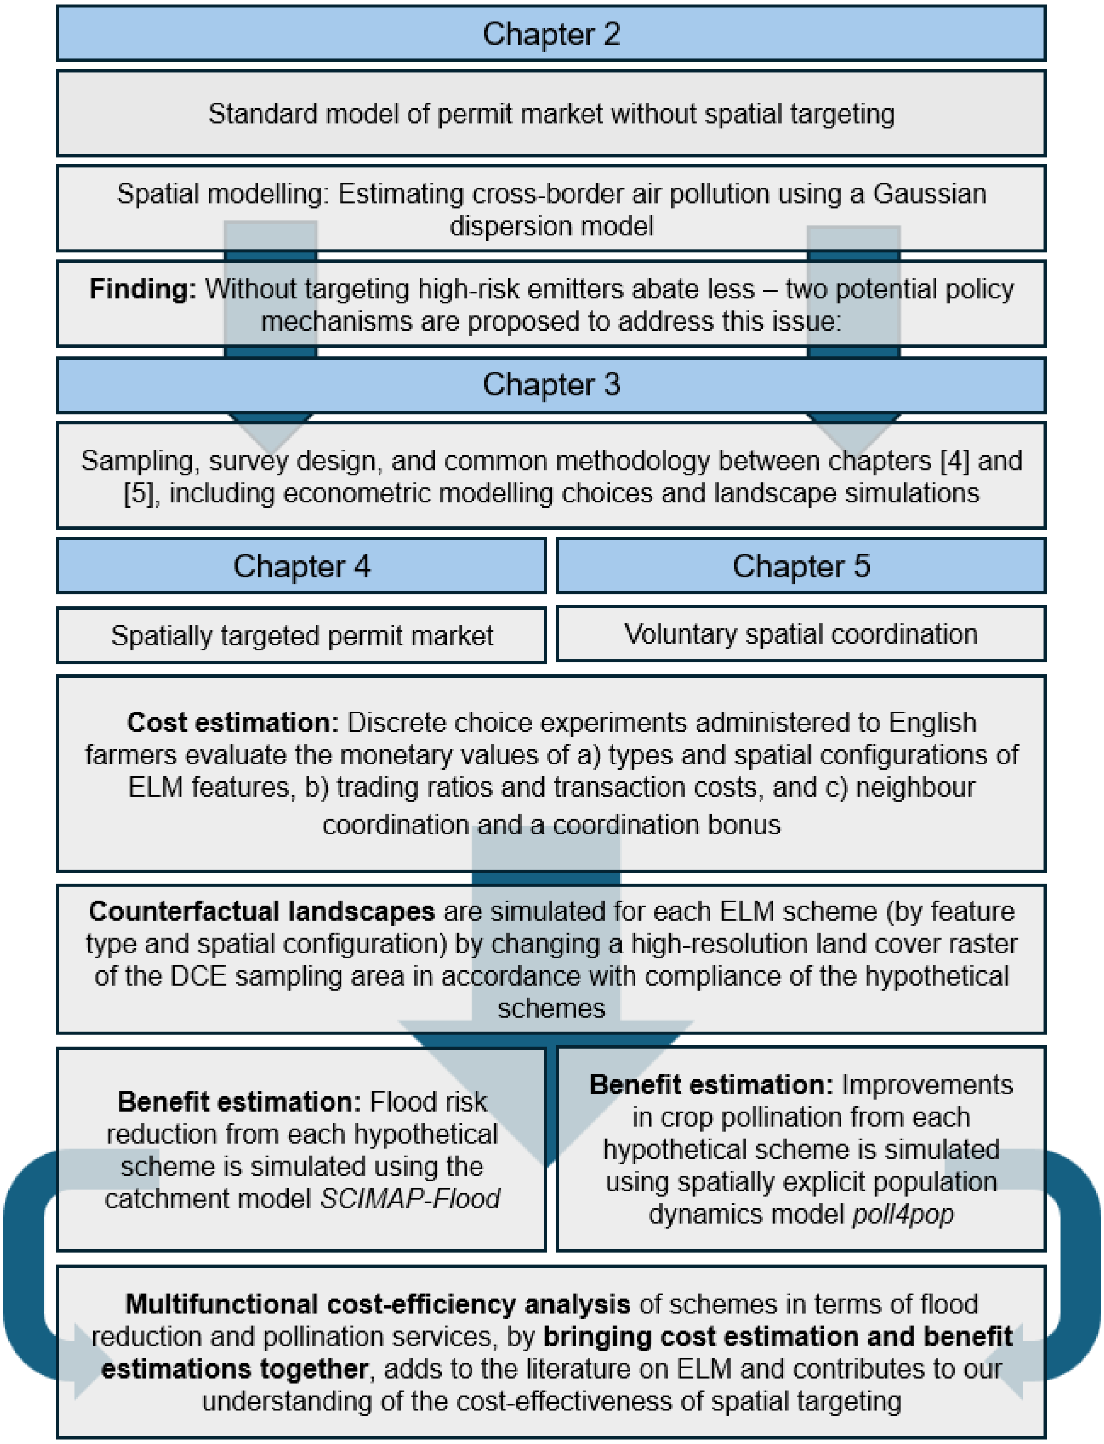

Figure 1.1: *Flowchart of thesis structure and contributions*

1.4 Scope

This research focuses on environmental policy in a European and North American context. This scope is motivated in part by the particular tensions between economic concerns and environmental protection that feature prominently in European and US discourse. My work also relies heavily on modelling using high-resolution environmental data which is largely limited to high-income countries (Arguez et al., 2012; Tanguy et al., 2021). The scarcity of high-quality data and mismatch between local institutional knowledge to research funding (Wintrup, 2022) results in a bias of research efforts towards western economies. A diverse body of work studies similar issues in other regions, including in East Asia (Cai et al., 2016; Heo et al., 2023; Liu et al., 2024), Africa (Benjamin & Sauer, 2018) and Latin America (Harding et al., 2021). It is important that policy recommendations are tailored to local economic and environmental conditions. That is not to say that the lessons from my research are entirely inapplicable to other contexts. For example, I find that coal-fired power plants in the US respond less to tightening of emission caps when they export a significant share of pollutants outside the state where they are regulated (Leppert, 2023). In China, Cai et al. (2016) similarly find evidence that upstream factories close to local authority borders contribute more to river pollution.

Chapters 4 and 5 are concerned with ELM schemes featuring so-called action-based payments. Such payments are conditional on farmers taking particular actions, such as planting trees or retiring land from intense grazing. An alternative, *result-based payments*, has attracted more interest in recent years. When the UK Government in November 2020 published *The Path to Sustainable Farming*, setting out the post-Brexit agenda for agriculture, it was unequivocally stated that a guiding principle would be a ”focus on achieving [environmental] outcomes” (Cardwell, 2023). Under results-based ELM schemes, payments are conditional on achieving a particular environmental result. As Bartkowski et al. (2021) emphasise, results-based ELM schemes are not common in practice. The authors attribute this to the sophisticated monitoring and measurement of outcomes that is required. This thesis aims to explore if spatial targeting can improve the cost-effectiveness of presently dominant schemes. For the purposes of cost-effectiveness analysis supported by choice experiments, it is also advisable to base the ELM options on schemes that farmers are familiar with (Johnston et al., 2017). Results in chapter 3 show that deviating significantly from current schemes increases the risk of inconsistent preferences. For these reasons, results-based payments are outside the scope of this thesis. However, I suggest that the simulation methods used here can also be applied in the measurement of outcomes in results-based ELM schemes.

Finally, although policy evaluation lies at the heart of this work, I do not attempt to completely monetise the benefits attributed to the hypothetical schemes in chapters 3 and 4. Even relatively direct benefits, such as the expected annual reduction in damage to life and property attributed to natural flood management, are highly complex to calculate. Indeed, the hydrological connectivity model I use to estimate runoff potential and determine trading ratios is insufficient for this task (Reaney, 2022). Monetisation of pollination services attributed to ELM schemes depends heavily on current pollinator abundance (Kleijn et al., 2006, 2015). Reduction of habitat fragmentation in agricultural landscapes may also contribute to non-market benefits (Correa Ayram et al., 2016).

Attempting to value natural goods that are not regularly traded in the economy presents new challenges (Hoyos, 2010). When using hypothetical stated preference surveys to estimate farmers' required compensation to enrol land in ELM schemes, I am asking them to value costs that they experience regularly and are core to their business: The loss of a defined area of productive land, the fencing of field edges, engagement with Defra, extension advisers and other stakeholders, etc. By contrast, when asked to put a value on e.g. restoration of woodland, a hiker may think of natural beauty and bird song, while an ecologist may value the provision of habitats for some obscure endangered species.

As economists ventured into stated preference studies for non-market valuation, Vatn and Bromley (1994) framed this inability to align respondents' perception of choice attributes as an insurmountable issue. Nevertheless, the non-market valuation literature persists and is growing more sophisticated, using lab- and revealed preference methods (Hanley & Perrings, 2019).

This thesis sets all of these issues to one side, and focuses on comparing scenarios in terms of required costs and expected environmental benefits. Policymakers and regulators can then judge its outcomes in terms of their goals and priorities, along with those of their constituents.

Chapter 2

Heterogeneous externalities in a pollution permit market without spatial targeting

2.1 Introduction

A key consideration in any attempt at regulating air pollution is its ability to effortlessly cross administrative and legal boundaries. A comprehensive theory of cross-border externalities was proposed as early as Montgomery (1972), who showed that the abatement effort mandated by the regulator ought to be higher for upwind sources that contribute to ambient pollution in downwind receptor regions. Indeed, in maintaining air quality or other environmental standards across regions with cross-border pollution, the optimal regional tax rate is a function of the downwind externality (Xepapadeas, 1992a). The logical question that follows is how does the regulator identify pollution sources that contribute to degrading the environment also in other regions? The United States' Environmental Protection Agency, the EPA, maintains close monitoring of ambient air quality using a network of monitors, as do environmental agencies in many industrialised countries. However, even when the government has broad authority to monitor emissions and gather accurate information, it is not always trivial to determine how much pollution from which source ends up where (Wei et al., 2018). It is rarely simply a matter of distance between source and receptor point, but as shown by e.g. Zheng et al. (2014), geography and meteorology also play important parts. While a large literature has studied the externalities firms impose on society, such as public health (Chay & Greenstone, 2003b; Fowlie et al., 2012; Schlenker & Walker, 2016) and urban amenity values (Zheng et al., 2014), comparatively less attention has been devoted to how firms that face different geographic conditions respond to regulation (Kampas et al., 2013) and to geography as an influencer of efficient policy.

Despite an established theoretical literature (Fowlie & Muller, 2019; Montgomery, 1972; Xepapadeas, 1992b) raising the issue, cross-border pollution remains salient in practice. For example, recent work by Heo et al. (2023) emphasise the problem, reporting that cross-border air pollution from China significantly increases mortality and morbidity in South Korea. Between 2016 and 2018, the US states of Connecticut, Delaware, Maryland, and New York each petitioned EPA to regulate pollution sources in upwind states that allegedly interfered with the petitioners' air quality standards (Gerrish, 2020). The efficient policy response depends on the primary driver of cross-border pollution.

While spatially non-targeted instruments can be effective in cases where cross-border pollution depends primarily on the emission rate, spatially targeted policies are preferable when geography is a significant driver (Holland & Yates, 2015; Xepapadeas, 1992b). With incomplete information about the leading cause of cross-border pollution, the most effective policy response is uncertain. This chapter clarifies this uncertainty in the context of US state-level standards for ambient air pollution, determined by the EPA and regulated under the federal Clean Air Act. EPA established the Acid Rain Program (ARP) under Title IV of the 1990 CAA amendments to reduce power sector emissions that cause acid rain (Stavins, 2003). Specifically, the ARP targets SO_2 emissions through cap-and-trade. The cap-and-trade system, currently covering over 2,000 electricity generating units across the United States, is widely regarded a success story in US environmental regulation, having contributed an estimated 10.8 million tonne reduction in SO_2 between 1990 and 2010, or 67% (Schmalensee & Stavins, 2013).

In 2005, the Clean Air Interstate Rule (CAIR) was promulgated under the federal law to limit the interstate transport of SO_2, an air pollutant contributing to acid rain primarily from burning fossil fuels, across 27 eastern states. However, CAIR was short-lived. In 2014 it was vacated following a 2008 ruling by the D.C. Court of Appeals in favour of North Carolina, which argued that the cap-and-trade system made downwind states powerless to combat emissions from upwind sources outside their jurisdiction (Kruse, 2009). Because upwind plants were able to purchase permits to cover their emissions, they could keep contributing to ambient pollution in a downwind state.

The 2011 Cross-State Air Pollution Rule (CSAPR) which succeeded CAIR following the legal challenges and remains in effect, attempts to target sources in upwind states by restricting the market for permits to within-state trading (Shouse, 2018). Recognizing that cross-border pollution can produce spillover harms (Heo et al., 2023), it is motivated to examine CAIR's impact in this respect. Using an atmospheric air pollution dispersion model suggested by Mendelsohn (1980), but rarely used to evaluate the need for spatially targeted regulation of air pollution (Jaramillo & Muller, 2016), I identify individual electric utilities that contribute to ambient SO_2 in downwind states.

Using a canonical difference-in-differences experimental design with utilities to be covered by CAIR SO_2 caps in the treatment group and remaining ARP regulated utilities as controls, I estimate the effect on interstate SO_2 pollution from tightening of emission caps. Because geography does not change over time, treatment timing captures the effect of emission reductions on downwind pollution.

The rest of the article is structured as follows: I first provide the policy background for CAIR and CSAPR, as well as the legal arguments that led the D.C. Court of Appeals in 2008 to rule that CAIR was ineffective at protecting downwind states. Secondly, I present the economics behind environmental externalities and the difference between spatially targeted and non-targeted permit allocation. I go on to describe the theory behind the Gaussian air pollution dispersion model (Zannetti, 2013) I develop and apply for the first time in combination with a natural experiment. A difference-in-differences design is appropriate for this problem because its separation of observations into two groups and two time periods can simultaneously handle two sources of bias. First, the post-treatment period dummy in the DD term addresses the *selection bias.* This is the bias arising from the fact that CAIR was not a random collection of states. The policy sought to target a region where SO_2 pollution was a particular problem. Second, the treatment group dummy in the DD term deals with *omitted variable bias.* This bias arises from common national trends in incentives unrelated to CAIR, such as GDP or the cost of abatement technologies. The model, referenced below as GAUSSMOD, is developed and optimized for replicability, and presented for an interdisciplinary and policy-oriented audience.

These sections arrive at the conclusion that the cross-border externality is a function of three variables: The rate of emissions at the source, typical weather conditions, and the geographic conditions. My identifying assumption is that while source emission rates change over time, geography does not (Fowlie et al., 2012). This allows me to more convincingly isolate any treatment effect caused by a reduction in emissions as a result of CAIR. To rule out unobserved abatement heterogeneity between groups I also estimate CO_2 emissions, which are not differentially regulated. Then, I present the data and dispersion model output. Finally, I present and discuss the results.

2.2 Background

The Clean Air Act is the United States' primary federal law to reduce nationwide air pollution. Initially enacted in 1963 the law, henceforth CAA, has been praised as a success of early U.S. environmental policy, for example in terms of health outcomes (Chay & Greenstone, 2003b). A collection of major amendments to the law came into force in 1990 (Waxman, 1991), and included tradeable permits in nitrogen oxides (NO_x) and sulphur dioxide (SO_2). A cap-and-trade system under Title IV of the CAA, also known as the 'Acid Rain Program' regulates acidifying pollutants, mainly from coal-burning power plants, by allocating permits to emitters and allowing reallocation via auction to improve economic efficiency. (McCubbin, 2009) Allowances under Title IV are regulated by the Environmental Protection Agency (EPA) under §7408(a) of the Clean Air Act. The Acid Rain program has involved two phases, beginning in 1990 and 2000 respectively. Title IV also requires sources to install a continuous emission monitoring system (CEMS) and annually report emissions to the EPA and state regulators (Ellerman et al., 2000).

In Phase I, individual emissions limits were assigned to the 263 most SO_2 intensive generating units at 110 plants operated by 61 electric utilities, and located largely at coal-fired power plants east of the Mississippi River. After January 1, 1995, these utilities could emit sulphur dioxide only if they had adequate allowances to cover their emissions. During Phase I, the EPA allocated each affected unit, on an annual basis, a specified number of allowances. The initial allowances were not auctioned but grandfathered based on sources' share of heat input during the baseline period 1985-1987. By Phase II, almost all coal-fired power plants were covered by the system. If trading permits represents a carrot in the system, the stick is a penalty of $2,000 per ton of emissions that exceed any year's allowances and a requirement that such excesses be offset the following year (Stavins, 2003).

Largely considered successful, it is estimated that between 1990 and 2008, the majority of reductions in U.S. air pollution was due to changes in environmental regulation (Shapiro & Walker, 2018). The federal CAA regulates individual states' emissions via the National Ambient Air Quality Standards (NAAQS) where they are responsible for maintaining caps on ambient concentrations of air pollutants. The NAAQS for SO_2 is 75 ppb, measured as the 99th percentile of 1-hour daily maximum concentration, averaged over three years. The EPA requires that individual states submit so-called State Implementation Plans (SIPs) detailing how they will comply with the national standards for each pollutant set under §7408 (Potoski, 2001).

Building on the success of the acid rain program, the EPA in 2005 introduced the Clean Air Interstate Rule (CAIR), which mandated that states and the federal government work together to address regional pollution. Constructed upon the previous pollution credit programs in the ARP, CAIR created a regional trading program to reduce interstate pollution (Pleune, 2006). The EPA determined which states would participate in the regional program based on whether they made a "significant contribution" to non-attainment of NAAQS for downwind states (Glasgow & Zhao, 2017). However, there was not a designation of individual plants as high or low risk of significant contributions, and one does not yet exist.

The 1990 amendments to the CAA also added provisions specifically to combat externalities due to spatial diffusion of air pollutants. This "Good Neighbour" provision states that an upwind state may be ruled in violation of Title IV if pollutants from point sources move to downwind states in such quantities that they impede the ability of the downwind state to meet its allowances under §7408 and its implementation plans (Gerrish, 2020; McCubbin, 2009). Although EPA found that out-of-state sources would cause non-attainment in 2010 (the States' deadline under the CAA for reaching attainment), EPA determined that it would not be feasible to reduce the out-of-state emissions by that time. Instead, CAIR required the reduction to be implemented in two phases. States would implement the first phase of reductions by 2009 for NOx and by 2010 for SO_2. A second set of reductions would bring the level of out-of-state contributions to air quality non-attainment to an acceptable level by 2015. After a downwind state has filed a complaint of a Good Neighbour violation under section 126, EPA has 60 days to respond.

If EPA determines action is necessary, the upwind state must address the emissions in their SIP, effectively reducing the permits its emitters are allowed to use. Failure to do so could, if the Good Neighbour provision is enforced, make the violating firm liable to pay the $2,000 per excess tonne SO_2. Since there is no borrowing of permits from future allocation to plants allowed under Title IV (Schennach, 2000), plants in the upwind state must either invest in abatement or buy permits at auction.

2.2.1 The collapse of CAIR: *North Carolina v. EPA*

An additional event on the timeline of interstate SO_2 regulation is of particular note. In the 2008 case *North Carolina v EPA*, the D.C. court of appeals ruled in favour of the state and a number of electric utilities, arguing that CAIR had several flaws, and because the EPA had adopted it as one, integral action, the rule in its entirety must be vacated and remanded to the EPA. The court's opinion was that CAIR could not properly respect the 'good neighbour provision' requiring sources to take responsibility for their contribution to non-attainment of NAAQS in the downwind state. One flaw found by the court was in CAIR's trading programs for SO_2, which it said essentially amounted to a "regionwide approach" which failed to prohibit sources "within the State from contribut[ing] significantly to non-attainment in any other State..." (Kruse, 2009) because sources could purchase enough SO_2 allowances to cover current emissions, resulting in no change (Tait, 2009). The result of the cap-and-trade system, North Carolina and a number of downwind power companies argued, is that downwind states and firms can do very little in terms of policy to address non-attainment of NAAQS, if significant contributions to ambient air pollution come from out-of-state sources that can buy permits to make up the difference. As summarised in Kruse (2009), the D.C. Circuit decided that the CAIR trading program went beyond the mandate of the CAA because the regional program did not address sources from one specific state contributing to non-attainment in another specific state.

In 2011, the Obama administration announced the Cross-State Air Pollution Rule (CSAPR) which replaced CAIR in 2015 and involves the same eastern states. CSAPR attempted to address the legal issues in CAIR by allowing only *within-state* trade in permits (Chan et al., 2012). As of 2021, there have been a number of section 126 petitions: Between 2016 and 2018, Connecticut, Delaware, Maryland, and New York each petitioned the EPA to regulate pollution from an upwind state. The EPA denied all four petitions.

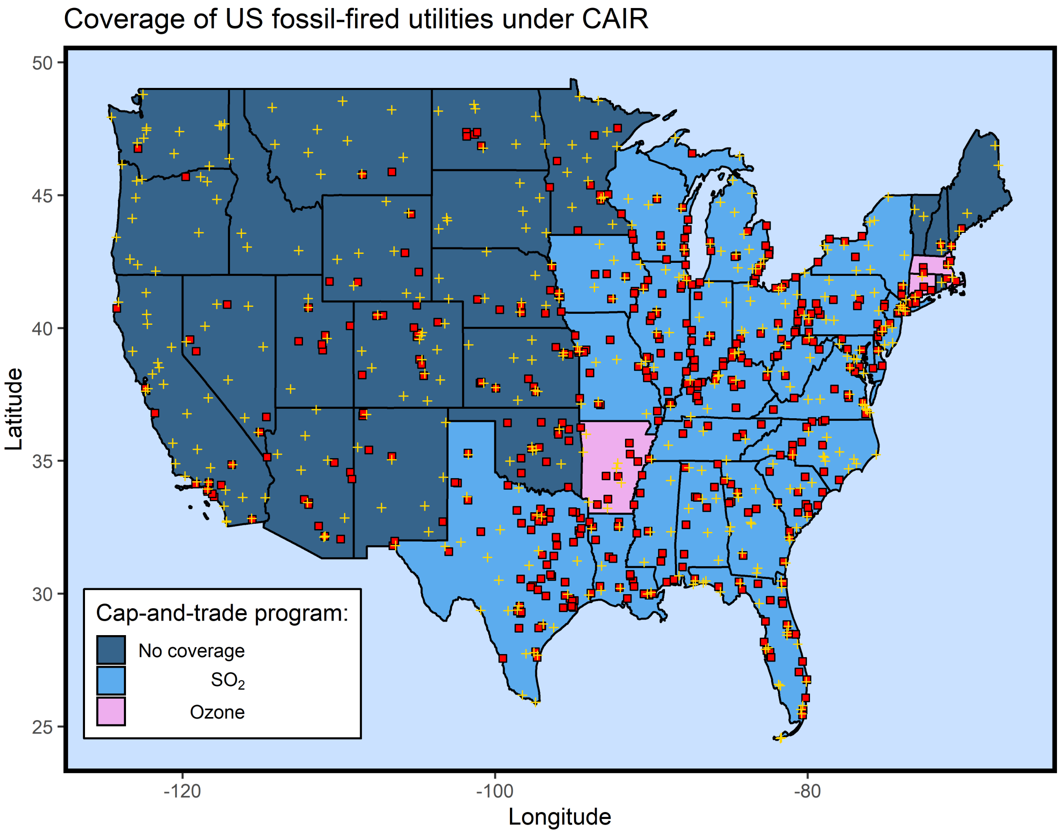

Figure 2.1: *CAIR coverage for 493 fossil-powered electric utilities. Neighboring weather stations (+) provide hourly weather inputs for the dispersion model* GAUSSMOD.

Delaware, Maryland, and New York challenged those denials in court. In 2020, the D.C. Circuit denied Delaware's petition, granted Maryland's petition in part, and vacated EPA's denial of New York's petition (returning the petition to EPA for reconsideration) (Gerrish, 2020). The unwillingness of the federal regulator to grant section 126 petitions may be interpreted by emitters as a signal that violations are unlikely to be investigated and punished (Harstad & Eskeland, 2010). If legal action does not come from the federal level, state regulators have no incentive to pursue cross-border emissions transported from sources in their own state.

2.3 Theoretical Framework

This article contributes to an ongoing empirical literature on the effectiveness of cap-and-trade programs (Barreca et al., 2021; Chan & Morrow, 2019; Glasgow & Zhao, 2017) by focusing on the less studied aspect of cross-border pollution (Chen et al., 2022) and by combining causal inference with geophysical modelling. In formulating an initial hypothesis, and throughout the remainder of this article, I make a number of assumptions about the way firms respond to changes in the expected cost of polluting the air. The natural experiment takes place in an economy with one environmental regulator and many polluting power plants. Plants are distributed across several regions, each with administrative borders and responsibility for maintaining limits on pollution set by the regulator.

In the standard cap-and-trade model, and in the absence of interstate pollution rules, the regulator determines ambient air quality standards according to its own evaluations of the social damage function, then introduces an emissions cap to achieve the ambient standards. Once the EPA allocated emission permits to coal-fired power plants and allowed trading in permits between plants it effectively introduced a market price for SO_2 emissions (Montgomery, 1972; Xepapadeas, 1992a). Questions about how to refine market-based policy designs to account for pollution transport and associated spatial variation in marginal damages have been the subject of contentious debate over the last decade. It was not until 2014 when a US court ruled that regulations limiting harmful emissions should proceed, the uncertainty around estimating damages from nonuniformly mixed pollutants notwithstanding (Fowlie & Muller, 2019). As such, CAIR was designed around sources trading permits at a uniform price.

When the representative firm is a price-taker on the permit market, it chooses its abatement level such that its marginal abatement cost equals the market price for permits P^T. The regulator does not know the firm's abatement cost function, and so initial allowances $\tilde{e}_i$ are not allocated based on the firm's marginal abatement cost but, in the case of the Acid Rain Program, on its share of heat input (Stavins, 2003). To enforce compliance, the Clean Air Act allows the EPA to impose a fine of $f = \$2,000$ per tonne in excess of the cap. Meanwhile, firm i chooses its abatement efforts and the amount of permits q_i to buy in order to minimise their individual total cost $c_i(e_i)$ which is a function of its emissions e_i:

$$\min_{e_i \geq 0} c_i(e_i) + P^T \times q_i + f(e_i - q_i) \ \text{ s.t. } e_i \geq \tilde{e}_i + q_i > 0 \tag{2.1}$$

where $c'(e) < 0$ which means that reducing emissions increases the cost. In other words, the marginal abatement cost is positive. Following Stranlund and Chavez (2000), I impose the restriction that all firms hold permits and that the number of permits held by the firm do not exceed its emissions. The Lagrangian:

$$\mathcal{L} = c_i(e_i) + P^T \times q_i + f(e_i - \tilde{e}_i - q_i) - \mu(e_i + \tilde{e}_i - q_i) \tag{2.2}$$

yields the Karush-Kuhn-Tucker (KKT) conditions:

$$\partial\mathcal{L}/\partial e = c'(e_i) + f - \mu = 0 \tag{2.3}$$

$$\partial\mathcal{L}/\partial q = P^T - f + \mu = 0 \tag{2.4}$$

$$\partial\mathcal{L}/\partial\mu = \mu \geq 0; \, \mu \times (q_i + \tilde{e}_i - e_i) = 0 \tag{2.5}$$

The complementary slackness condition (2.5) reveals that $e = \tilde{e} + q$ has a positive shadow cost μ. Substituting μ in equation (2.4) with $(f - P^T)(q_i + \tilde{e}_i - e_i)$ shows that full compliance $q_i + \tilde{e}_i = e_i$ only occurs when the fine exceeds the market price for permits, P^T. It also illustrates that a higher initial allowance $\tilde{e}_i$

results in a lower demand for tradable permits at all levels of P^T. Lower demand for permits across the market results in a lower equilibrium price and abatement. When the emissions cap is reduced as anticipated by CAIR states, following its announcement in 2005, average abatement costs rise and with them the permit price. Irrespective of its compliance status, the firm will stop investing in abatement once the marginal abatement cost equals the market price of permits. This is because the marginal abatement cost rises with the abatement effort, while the price for permits does not depend on the individual firm's choices (Stranlund & Chavez, 2000). On the issue of market power in the permit market, Hintermann (2017) shows that price manipulation by dominant firms primarily results in pass-through of abatement costs onto consumers and taxpayers. Overall, a reduction in the emissions cap is still expected to increase the price for permits. Accordingly, granted only the assumption that the threat of penalties for non-compliance with the CAIR caps is credible, I make the following proposition:

Proposition I: An increase (decrease) in the market price of permits results in a decrease (increase) in average emissions across power plants.

2.3.1 The firm's response to cross-border pollution

Now suppose that only some share $\delta \in [0, 1]$ of the firm's pollution stays within the region (such as a state) where it is regulated. This can result from proximity to a state border and prevailing winds in that direction. Supported by historical accounts of the CAIR period (Glasgow & Zhao, 2017; Schmalensee & Stavins, 2013), I assume that states did not reliably fine excess emissions from sources outside their borders. In this setting, the objective function of a firm located in region r becomes:

$$\min_{e_i \geq 0} c_i(e_i) + P^T \times q_i + f\left(\delta_{ir} e_i - \tilde{e}_i - q_i\right) \text{ s.t. } \delta_{ir} e_i \geq q_i > 0 \tag{2.6}$$

Repeating the minimisation procedure from equations (2.3) trough (2.5) we find that emissions are set such that:

$$-\frac{\partial c(e_i)}{\partial e_i} = \delta_{ir} P^T \tag{2.7}$$

Theory predicts that as a larger share of pollution is transported out of the state in which the polluter is located (δ tends towards 0) a higher permit price is required for the upwind firm to switch from permits to abatement. This is because in the event of non-compliance of an amount Δ tonnes above its allocated emission cap, the firm only expects to be penalised for a fraction of total excess emissions $\delta \times \Delta$.

Although NAAQS are determined at the federal level, states have autonomy regarding the implementation and, crucially, enforcement of the SIP (Stavins, 2003). The dynamics of interstate pollution control within the federal US is therefore comparable to the international case. Maler (1989) applies game theory to the European acid rain problem and does not find cooperative equilibria without international transfers. These results suggest that upwind states would regulate only the amount of pollution which remains within their borders. Granted the assumption that externalities affecting downwind states do not affect the enforcement of State Implementation Plans, I state the second proposition:

Proposition II: The firm does not expect to be fined for emissions exceeding its allowances if the excessive pollution is transported out of the state in which it operates.

2.3.2 Gaussian dispersion modelling

To quantify downwind SO_2 dispersion from each coal-fired power plant, I develop GAUSSMOD, a three-dimensional Gaussian dispersion model, in `Python 3.6`. The Gaussian model is one of the simplest dispersion models for point-source air pollutants. The plume dispersion equations featuring Gaussian distributed dispersion were first derived in Sutton (1947) and have become increasingly popular. In the advent of stringent environmental control regulations, there was an immense growth in the use of air pollutant plume dispersion calculations between the late 1960s and today (Zannetti, 2013). Gaussian models are popular because they are mathematically tractable, easy to implement, and rely on widely available data. They offer advantages over simple trajectories used in e.g. Heo et al. (2023) because they allow for estimation of cross-border concentrations.

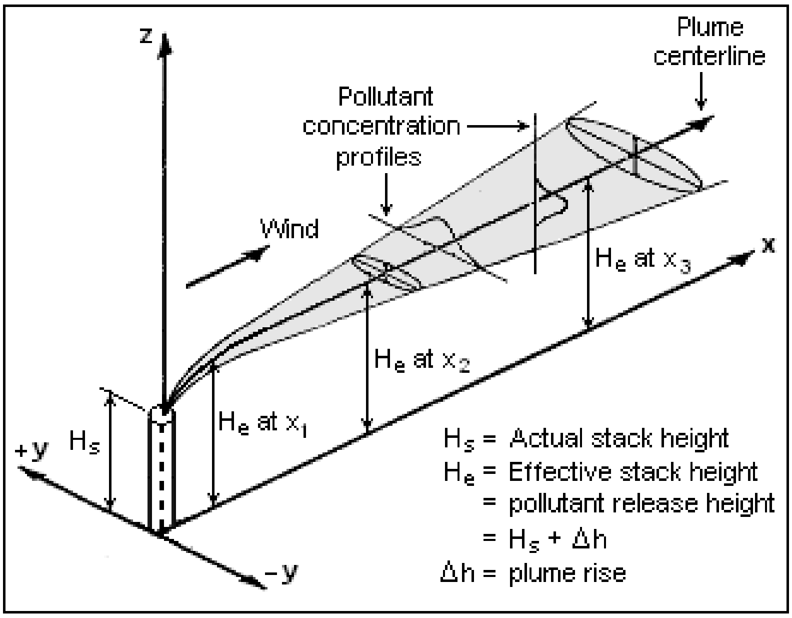

Figure 2.2: *The plume centreline vector* **x** *runs in the wind direction angled v degrees. The pollutant concentration follows a Gaussian distribution along the dispersion vector* **y** *extending perpendicular from the plume centreline and the vertical height vector* **z**. *Image from Leelőssy et al. (2014)*

In this paper, I implement the Gaussian model from Abdel-Rahman (2008) and U.S. EPA (1989) and apply it to SO_2 emissions. The plume dispersion equations are as follows:

$$C(x,y,z) = \frac{Q}{u} \cdot \frac{f}{\sigma_y\sqrt{2\pi}} \cdot \frac{g}{\sigma_z\sqrt{2\pi}} \tag{2.8}$$

where $f = \exp[-y^2/(2\sigma_y^2)]$ is the crosswind dispersion parameter and $g = \exp[-(z-H)^2/(2\sigma_z^2)]$ is the vertical dispersion. Q is the emissions rate expressed in grams per second. C is the concentration of emissions, in g/m^3, at any receptor located x meters downwind from the emission source, y meters crosswind from the emission plume centreline, and z meters above ground level. σ_y is the horizontal standard deviation of emissions dispersion, while σ_z is the standard deviation in the vertical. σ_y and σ_z are functions of the atmospheric stability class (i.e. a measure of the turbulence in the ambient atmosphere) and of the downwind distance to the receptor.

The two most important variables affecting the degree of pollutant emission dispersion obtained are the height of the emission source point and the degree of atmospheric turbulence. The more turbulence, the greater the degree of dispersion. For a description of the six stability classes A-F used in this model that depend on wind speed and cloud cover, see Pasquill (1961). The equations for σ_y and σ_z are:

$$\begin{aligned} \sigma_y(x) &= \exp(I_y + J_y ln(x) + K_y[\ln(x)]^2) \\ \sigma_z(x) &= \exp(I_z + J_z ln(x) + K_z[\ln(x)]^2) \end{aligned} \tag{2.9}$$

where I, J, and K are coefficients that depend on the stability class at the stack location (Seinfeld & Pandis, 2016), Ch. 18. Equation (2.9) shows that both crosswind dispersion and vertical dispersion are functions of distance downwind from the pollution source, with lower concentration in both dimensions further from the smoke stack.

Equation (2.8) also shows that the concentration at ground level can be reduced by increasing the height of the smoke stack H. The effective height H_e of the smoke centreline is the sum of the stack height and the plume rise at a given distance x from the smoke stack. The plume rise is determined by the downwind horizontal distance from the stack and the buoyancy factor, which describes the upward force exerted by the gas on the air above. (Beychok, 2005) The buoyancy factor F is calculated using the following equation:

$$F = g \times v_e \times R^2 \times \frac{T_g - T_a}{T_a} \tag{2.10}$$

where $T_g - T_a$ gives the temperature difference between the exit gas and the surrounding air.

Table 2.1: *Variables and Physical Constants*

g	Gravity of Earth	9.8 m/s^2
v_e	Gas Exit Velocity	m/s
T_a	Temperature of Air	°K
T_g	Temperature of Exit Gas	°K
R	Radius of Flue Stack	m

Because hot gases rise faster, a large temperature gradient between the sulphur dioxide and ambient air will allow the pollutant to rise higher before the temperatures equalise and wind speed and direction dominate as drivers of plume trajectories. Similarly, a high gas exit velocity will have the same effect (Beychok, 2005). The model uses the plume rise equation from Briggs (1982) where the plume rise $\Delta h = 1.6F^{1/3}x^{2/3}h^{-1}$ and thus the effective stack height $H_e = H_s + \Delta h$. [2]

[2]Estimating empirical plume rise equations has proved challenging. Carson and Moses (1969)

2.4 Data

The raw data used in this article is exclusively from publicly available sources. Replication code and documentation, including the source code for the dispersion model GAUSSMOD, are published online. [3] Hourly data on wind speed, wind direction, ambient temperature and cloud cover were obtained from the Global Historical Climatology network (Menne et al., 2012). The hourly 30-year normals dataset includes 1991-2020 averages for every hour, totalling 8,760 hours. After incomplete time series had been removed, complete records remained for 423 weather stations across the continental United States. The normals are constructed from hourly observations, and quality assurance checks are routinely applied to the full dataset, although Menne et al. (2012) acknowledge that the data are not homogenized to account for artefacts associated with the various eras in reporting practice at any particular station (i.e., for changes in systematic bias). Hourly data were aggregated into 12-hour daytime (07.00 - 18.59) and night-time (19.00 - 06.59) averages. Normals in wind direction, speed and cloud cover over a 30-year period were used because they are the most indicative of hourly variation in these variables across any given year (Arguez et al., 2012). To account for climate trends, observed air temperature daily time series were used instead of normals following Leppert et al. (2021). Daytime temperature was calculated as a weighted average of maximum and minimum temperatures (0.75*TMAX + 0.25*TMIN) and night-time as 0.25*TMAX + 0.75*TMIN.

---

compare 15 formulas using stack and atmospheric data and find large variation in average plume rise, from 35.2 to 151.9 meters. Briggs (1965) suggests that "...the rise of most hot plumes is caused almost entirely by buoyancy due to heat; the most important stack parameter for such plumes is the buoyancy flux F, proportional to the heat flux." Briggs later showed in Briggs (1982) that in usual atmospheric conditions, the plume rise peaks some distance x_f downwind from the stack beyond which $\Delta h = 1.6F^{1/3}x_f^{2/3}u^{-1}$. The so-called Briggs plume rise equations remain popular in Gaussian disperson models (Beychok, 2005) and are used also here.

[3] https://github.com/DanielLeppert/EEPS_cross-border_SO2

Data on plant characteristics were obtained from the U.S. Energy Information Administration which publishes data collected from all coal-fired power utilities in annual EIA-767 and EIA-923 surveys. The surveys include data on net generation, heat input, stack height, stack radius, mean exit gas velocity, and mean exit gas temperature. The environmental compliance form also provide self-reported plant-level spending on flue gas desulphurisation (FGD).

While self-reports come with the usual caveats, EIA form data have been used in previous research on coal-fired utilities' emissions accounting (Quick, 2014) and remain the most comprehensive publicly available reports. Data on annual SO_2 emissions and permit holdings for coal-fired power plants across the CAIR/CSAPR region were collected from the Air Markets Program data supplied by the U.S. EPA. Plant-level emissions data are available from the conception of the Acid Rain Program in 1995 through to today, and include values from firms' own reports as well as EPA monitoring.

Utility codes that uniquely identify each plant are consistent across EIA and EPA datasets and allow me to track individual utilities through changes in the surveys over the years. Emissions, net generation, operational flue gas desulphurisation spending (filters, scrubs, sorbent and labour) have missing entries as completed surveys were not received by the EIA for every utility in every year. There is a small discontinuity in 2007 when the EIA-923 form superseded the EIA-906, EIA-920, FERC 423 and EIA-423. This change improved coverage. Schedule 2 of the EIA-923 collects the plant level fuel receipts and cost data previously collected on the FERC and EIA Forms 423. Several approaches exist to deal with missing data. The researcher might collect more data themselves, drop observations containing missing data in at least one variable from the sample, or use one among a number of imputation methods (Little & Rubin, 2019). As the first option is not feasible and the second presents an avoidable loss of power, I compare the summary statistics from the imputed data with the complete analysis data, where entries containing missing data are removed.

Missing values were imputed based on the remaining plant characteristics while accounting for plant- and yearly fixed effects using multivariate imputation with the R MICE package. Multivariate imputation is commonly used in survey data and can provide smaller variance than alternative methods with small sample sizes ($< 10,000$) (Yadav & Roychoudhury, 2018). The MICE (Multiple Imputation by Chained Equations) algorithm is implemented in four steps (Van Buuren & Groothuis-Oudshoorn, 2011):

1. Missing values are imputed with a simple method such as imputing the mean
2. The imputed means are returned to missing, for only one variable Y at a time
3. The non-missing observations of the current Y are regressed on the other variables as predictors
4. Regression coefficients for each predictor are used to impute missing values in Y, which is then itself used as a predictor in case of further variables containing missing data

Table 2.2 shows summary statistics from the imputed dataset next to the complete data. Comparing means and medians shows that distributions for several variables are skewed towards zero. Following suggestions in Little and Rubin (2019), I therefore use predictive mean matching in step 1) which is implicit and does not require specifying the distribution of the target variable. A Jarque-Bera test rejects a normal distribution for all variables in both samples (p-values < 0.01). Deviation from normality does not in itself invalidate regression analysis, but may be exaggerated by outliers in the sample and should be handled with care in model specification.

Table 2.2: *Summary statistics*

Variable	**Min**	**Mean**	**Median**	**Max**
		Dropped NAs		
SO_2 emissions (kt / year)	0.00	16.5	7.25	284
CO_2 emissions (kt / year)	0.25	4,609	3,153	27,231
Generation (GWh / year)	0.00	4,292	2,820	25,054
Heat input (BBtu / year)	0.02	45,876	32,477	265,410
Operating time (hours / day)	0.00	45.3	41.2	231.6
Distance to state border (kilometers)	0.002	50.4	35.4	268
		Imputed NAs		
SO_2 emissions (kt / year)	0.00	16.4	7.18	285
CO_2 emissions (kt / year)	0.25	4,377	2,843	27,231
Generation (GWh / year)	0.00	4,596	3,141	25,054
Heat input (BBtu / year)	0.00	45,563	32,200	265,410
Operating time (hours / day)	0.00	45.3	41.6	231.6
Distance to state border (kilometers)	0.002	50.2	34.5	268

Note: Total operating time across generators of a plant may exceed 24 hours.

Figure 2.3 shows scatter plots of four covariates against SO_2 emissions. I plot a log-log specification which best fits the linear model given the distributions of covariates. Figure 2.3 shows that the imputed sample (red) contains more outlier observations. Specifically, they arise from imputed zeros in unobserved emissions data. Weighing the risk of overstating standard errors using the imputed sample against the modest loss of power (8,452 versus 8,557 observations) I proceed with the smaller sample without imputation.

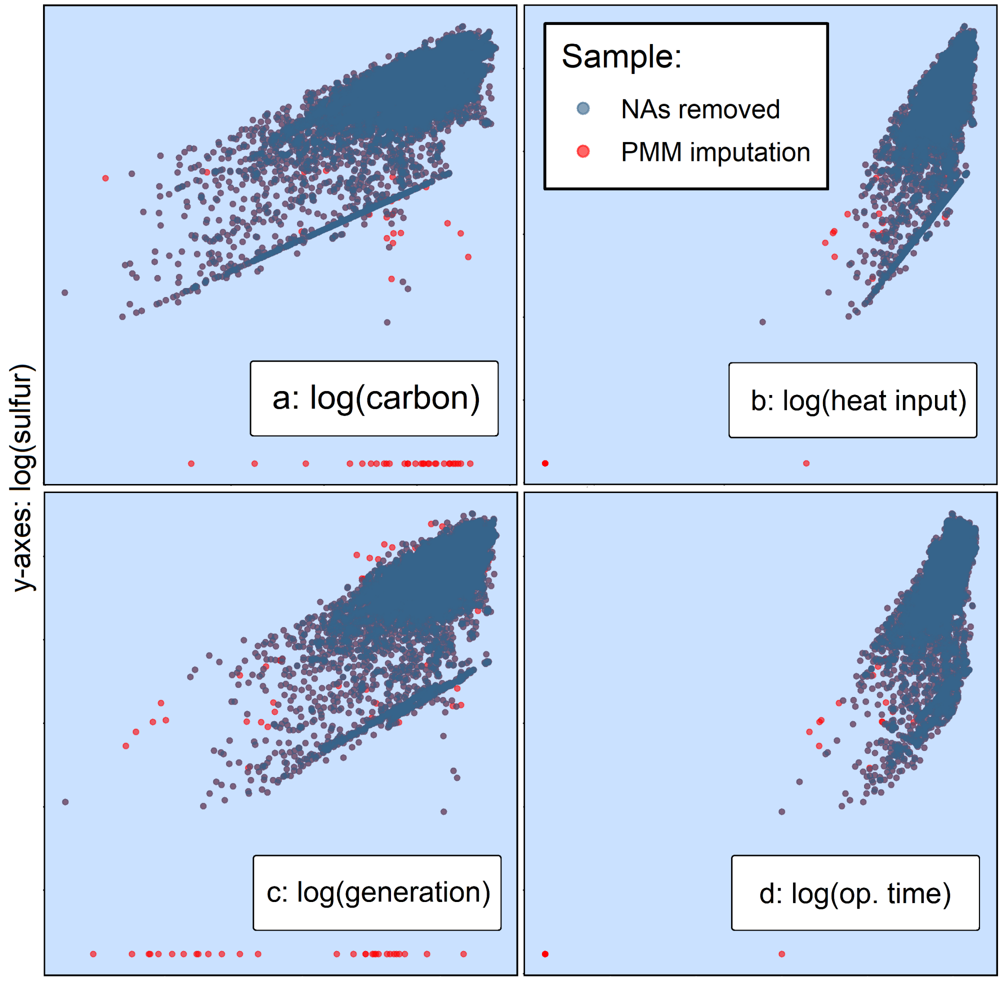

Figure 2.3: *Scatter plot by sample of sulphur against a) carbon, b) heat input, c) generation and d) operating time. The bimodality in sulphur arises from lower emissions by plants mixing coal-fired generators with oil-fired combustion.*

2.5 Method

This paper aims to estimate the effect of tightening the cap on SO_2 emissions on overall- and cross-border pollution. According to Proposition I, the announcement of CAIR should increase the market price for permits as affected firms scramble to comply with the lowered emission cap. Compliance was incentivized via a $2,000 fine per excess tonne of SO_2 and the enforcement mechanism involved mandatory installation of CEMS and emission reporting (Ellerman et al., 2000). Higher permit prices relative to marginal abatement costs (the cost of flue gas desulphurisation, such as limestone wet scrubbers, has declined throughout the study period, for both treatment and control groups (Chestnut & Mills, 2005)) are expected to increase abatement in CAIR states compared with unaffected emitters. Equation (2.8) states that SO_2 dispersion correlates positively with emission rates. I can therefore state in conjunction with Proposition I the first null hypothesis:

Hypothesis I: The announcement of CAIR caused no change in average cross-border SO_2 emissions from the power sector.

Rejecting hypothesis I would confirm that emission rates are important drivers of cross-border pollution, possibly alongside time-invariant factors like the locations of point-sources. The 2008 *North Carolina v. EPA* ruling established that interstate trade in permits between sources invalidates protection against cross-border pollution. A separate enforcement mechanism exists via the Good Neighbour provision wherein downwind states can petition the EPA to penalise cross-border sources. However, as emphasized in Harstad and Eskeland (2010), the reluctance of the EPA to grant Section 126 petitions call into question the likelihood of penalties. Based on Proposition II that firms do not expect to be fined for excess emissions that are transported out of their home state, I formulate the following null hypothesis:

HYPOTHESIS II: Plants contributing cross-border transport of SO_2 emissions did not respond differently to the CAIR announcement.

Rejecting hypothesis II would provide evidence that interstate polluters make less effort to comply with emission caps. In a natural experiment with electric utilities covered by CAIR in the treatment group and remaining ARP utilities as controls, inference relies first on identifying upwind power plants and estimating their cross-border emissions. I do this by feeding hourly data on SO_2 emission rates and local weather conditions for coal-fired power plants in 27 eastern states into a custom Gaussian air dispersion model GAUSSMOD.

Defining cross-border pollution

The cross-border SO_2 is defined as the average SO_2 concentration (μg/m^3) dispersed from a given plant outside of the state in which it is located. Based on heat input, stack flue characteristics and local weather conditions, GAUSSMOD calculates the concentration measured at ground level (1.5 meters) where health impacts are typically measured (World Health Organization, 2006). Dispersion is calculated across a 50,000 m^2 area around the plant, with a resolution of 1,000 m^2 following De Kluizenaar et al. (2001). Figure 2.4 displays the average daily SO_2 dispersion for two large coal-fired power plants, Barry Electric Generating Plant in Alabama and George Neal South Power Plant in Iowa. Over an average day, pollution from George Neal is transported across the Iowa-Nebraska border. Figure 2.4 illustrates how location and weather trends affect the problem of cross-border pollution. Quality control of GAUSSMOD is reported in appendix 2.7

2.5.1 Causal Identification and Estimation

Difference-in-differences (DD) is a method designed to estimate the causal impact of a policy on some outcome, such as cross-border pollution. It is known as a quasi-experimental method, because it attempts to approximate randomised controlled experiments, arguably the gold standard of empirical science, using observational data outside of a controlled lab setting. It requires observations from before and after some policy intervention, from the treatment group and unaffected controls.

CAIR raised the price of permits for SO_2 emissions by reducing the supply relative the nationwide Acid Rain Program via a new regional cap-and-trade program (Shouse, 2018). An increase in the permit price is expected to cause an increase in abatement, because power companies are willing to accept a higher abatement cost. The increase in the permit price following the announcement of CAIR in 2005 appears clearly in figure 2.5. I define years prior to 2005 as a pre-treatment period, while years following CAIR introduced in 2005 are in the post-treatment period. The two periods produce the first difference in the DD setup. Crucially, CAIR was a regional program covering power plants in 27 states. Plants covered by the rule are labelled as treated, while remaining plants serve as a control group.

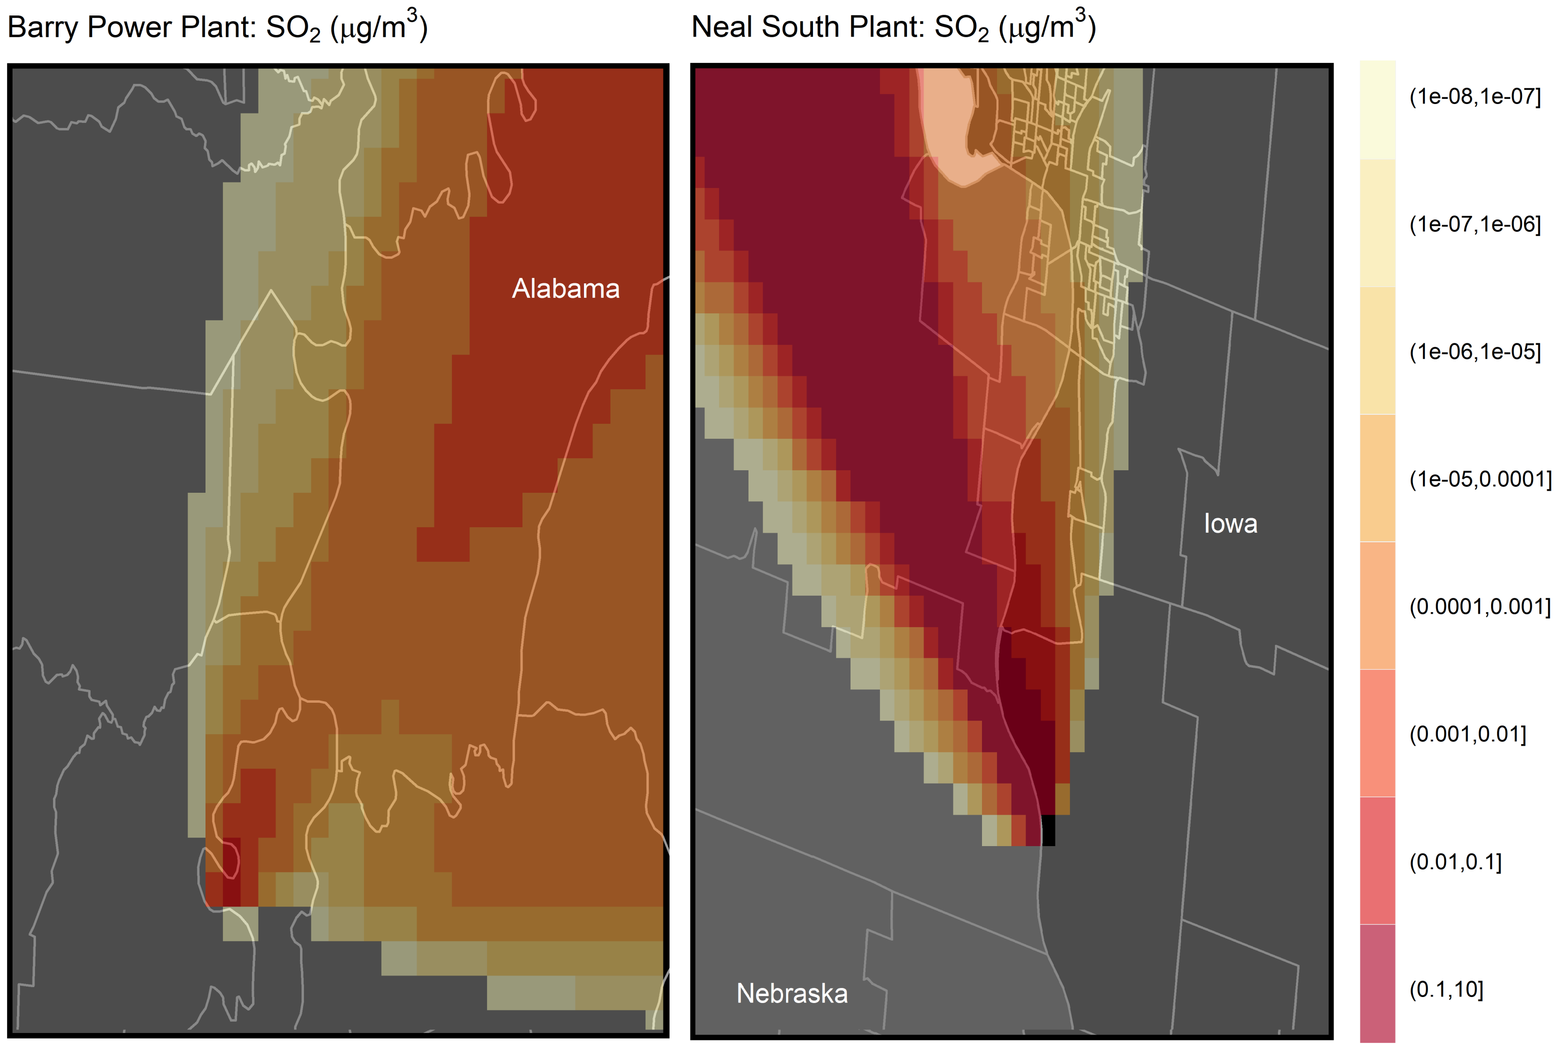

Figure 2.4: *SO_2 dispersion computed with* GAUSSMOD *is plotted over a 50,000 m^2 area around two example power plants.*

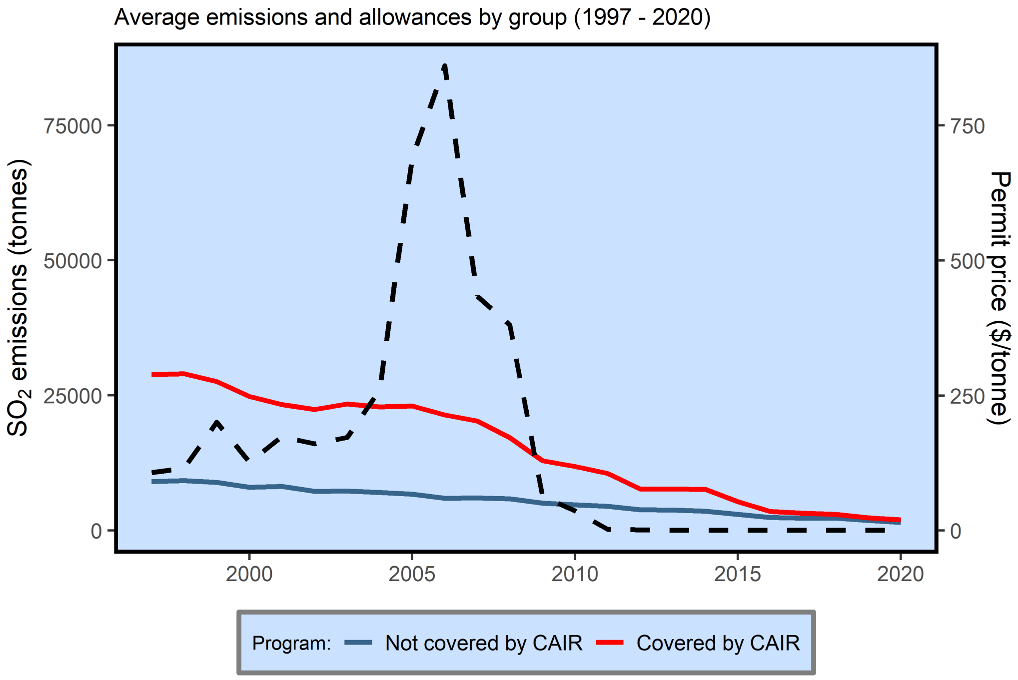

Figure 2.5: *Solid lines denote total annual SO_2 emissions across plants in CAIR states (red) and the control group (blue), and the dashed line the market price for permits. CAIR was announced in 2005.*

These groups produce the second difference in DD. All coal-fired power plants within the CAIR region are treated at the same time and in the absence of staggered treatment (Goodman-Bacon, 2021), I use the canonical two-way fixed effects difference-in-differences model with a panel of plants i and years t:

$$e_{it}^k = \alpha + \beta_1 G_i + \beta_2 CAIR_t + \beta_{DD}(G_i \times CAIR_t) + \boldsymbol{\beta X_{it}} + \epsilon_{it} \tag{2.11}$$

where the index k for outcome e denotes a) total SO_2 emissions, b) cross-border SO_2 emissions, and c) CO_2 emissions as a robustness check. G_i is a dummy variable taking the value 1 if plant i is covered by CAIR, and zero otherwise. $CAIR_t$ is a dummy variable taking the value 1 when year t is in the post-CAIR years and zero otherwise. $\mathbf{X_{it}}$ is a vector of covariates. As noted in Schmalensee and Stavins (2013), initial allocation of annual allowances to firms under the Acid Rain Program was based on heat input.

Greater heat input is therefore expected to be associated with higher emissions. Similarly I control for number of permits held by the firm, where high emitters are expected to hold more permits. Further control variables are net electricity generation, total operation time across a plant's generators, and desulphurisation technology (Bostian et al., 2022) that vary across plants. Examination of the raw data (figure 2.3) shows that a log-log specification in sulphur, heat input, generation and operating time produces the best linear model fit.

β_{DD} is the double-difference estimator and the coefficient of interest. It is the difference in average outcome in the treatment group before and after treatment, minus the difference in average outcome in the control group before and after treatment. It can be interpreted as the average treatment effect on CAIR states if, without the policy, the outcome would have evolved in parallel in the treatment- and control groups. This is the parallel trends assumption (Donald & Lang, 2007) which I will discuss in detail shortly. If β_{DD} is significantly different from zero, hypothesis I is rejected.

To test hypothesis II, model (2.11) is extended in equation (2.12) with a triple differences model (Kellogg & Wolff, 2008) where the DD variable is interacted with a dummy variable C_{it} indicating if the maximum cross-border SO_2 from plant i in year t exceeds 1% of NAAQS, or 0.75ppb. This is the screening threshold to identify states with sources that may contribute significantly to air quality problems in downwind states (Shouse, 2018; U.S. EPA, 2019). I do this to test for heterogeneous treatment effects between plants that contribute meaningfully to downwind cross-border pollution and those that do not, following similar experimental designs in e.g. Berck et al. (2016) (heterogeneous tax rates) and Dubos-Paillard et al. (2019) (flood risk). The share of treated plants in the sample of cross-border polluters is 79% versus 65% among plants that do not contribute to cross-border pollution.

$$
\begin{aligned}
e_{it} = \alpha + \beta_1 G_i + \beta_2 CAIR_t + \beta_3 C_{it} + \beta_4 (G_i \times C_{it}) + \\
\beta_5 (CAIR_t \times C_{it}) + \beta_{DD} (G_i \times CAIR_t) + \\
\beta_{DDD} (G_i \times CAIR_t \times C_{it}) + \boldsymbol{\beta X_{it}} + \epsilon_{it}
\end{aligned}
\tag{2.12}
$$

In the triple differences (DDD) setup, following the reasoning in Gruber (1994), I compare the double difference among plants that are interstate polluters (max cross-border $SO_2 > 0.75$ppb) against the double difference among plants that are not. The coefficient of interest β_{DDD} tells us the difference in the treatment effect between cross-border polluters and others. An estimate of β_{DDD} statistically different from zero rejects hypothesis II. Theory established in section 2.3 predicts a $\beta_{DDD} > 0$ due to moral hazard. The identifying assumption of this DDD estimator is fairly weak: I have previously established that there is no change in policy between $C_{it} = 1$ and $C_{it} = 0$ due to the insufficiency of CAIR to penalize cross-border pollution. Like the double difference setup, it also requires that there be no contemporaneous shock that affects the relative outcomes of the treatment group in the same state-years as the law.

Addressing selection bias and parallel trends

Figure 2.6 maps the power plants in my data broken down by average emission rates between 1997 and 2005, before the CAIR announcement. It also shows whether a given plant transported SO_2 concentrations across a state border in an average year during this period. Plants are coloured to reflect their average SO_2 emission rate over the pre CAIR period, with the largest emitters shown in red and the smallest in green. Plants that transport SO_2 into a neighbouring state are plotted as either diamonds (border distance under $1,000$ meters) or triangles (over $1,000$ meters). Figure 2.6 displays several low-emission plants as cross-border polluters, showing that location plays a role. The majority of cross-border polluting plants are located in states that would be covered by CAIR, as are those with the highest overall emission rates. This is unsurprising as the CAIR region sought to address SO_2 pollution from the worst emitters. Although Heckman et al. (1996) recommend that the two-by-two treatment group and time interaction is robust to selection bias, the double- and triple difference estimators only recover the true causal effect of the policy of interest when there are not concomitant (simultaneously occurring) trends that differentially affect the treatment and control groups (Wooldridge, 2007).

I perform a robustness test following the difference-in-differences approach in Jia et al. (2021), who compare treated observations only to "matched" controls that have similar characteristics. Results from this approach using propensity score matching are reported in Appendix A. In this case, concomitant (simultaneously occurring) treatment effects could arise from policies and economic trends that differentially (dis)incentivises pollution between CAIR states and outside. To test for concomitance bias, I also estimate a variant of equations (2.11) and (2.12) with CO_2 emissions as the outcome variable. CO_2 emissions result from the same coal burning process as do SO_2 emissions and are perfectly correlated absent any abatement. However, the two regions did not regulate CO_2 emissions differently. The concomitance hypothesis can be more confidently rejected if no CAIR-related treatment effect can be observed for carbon emissions.

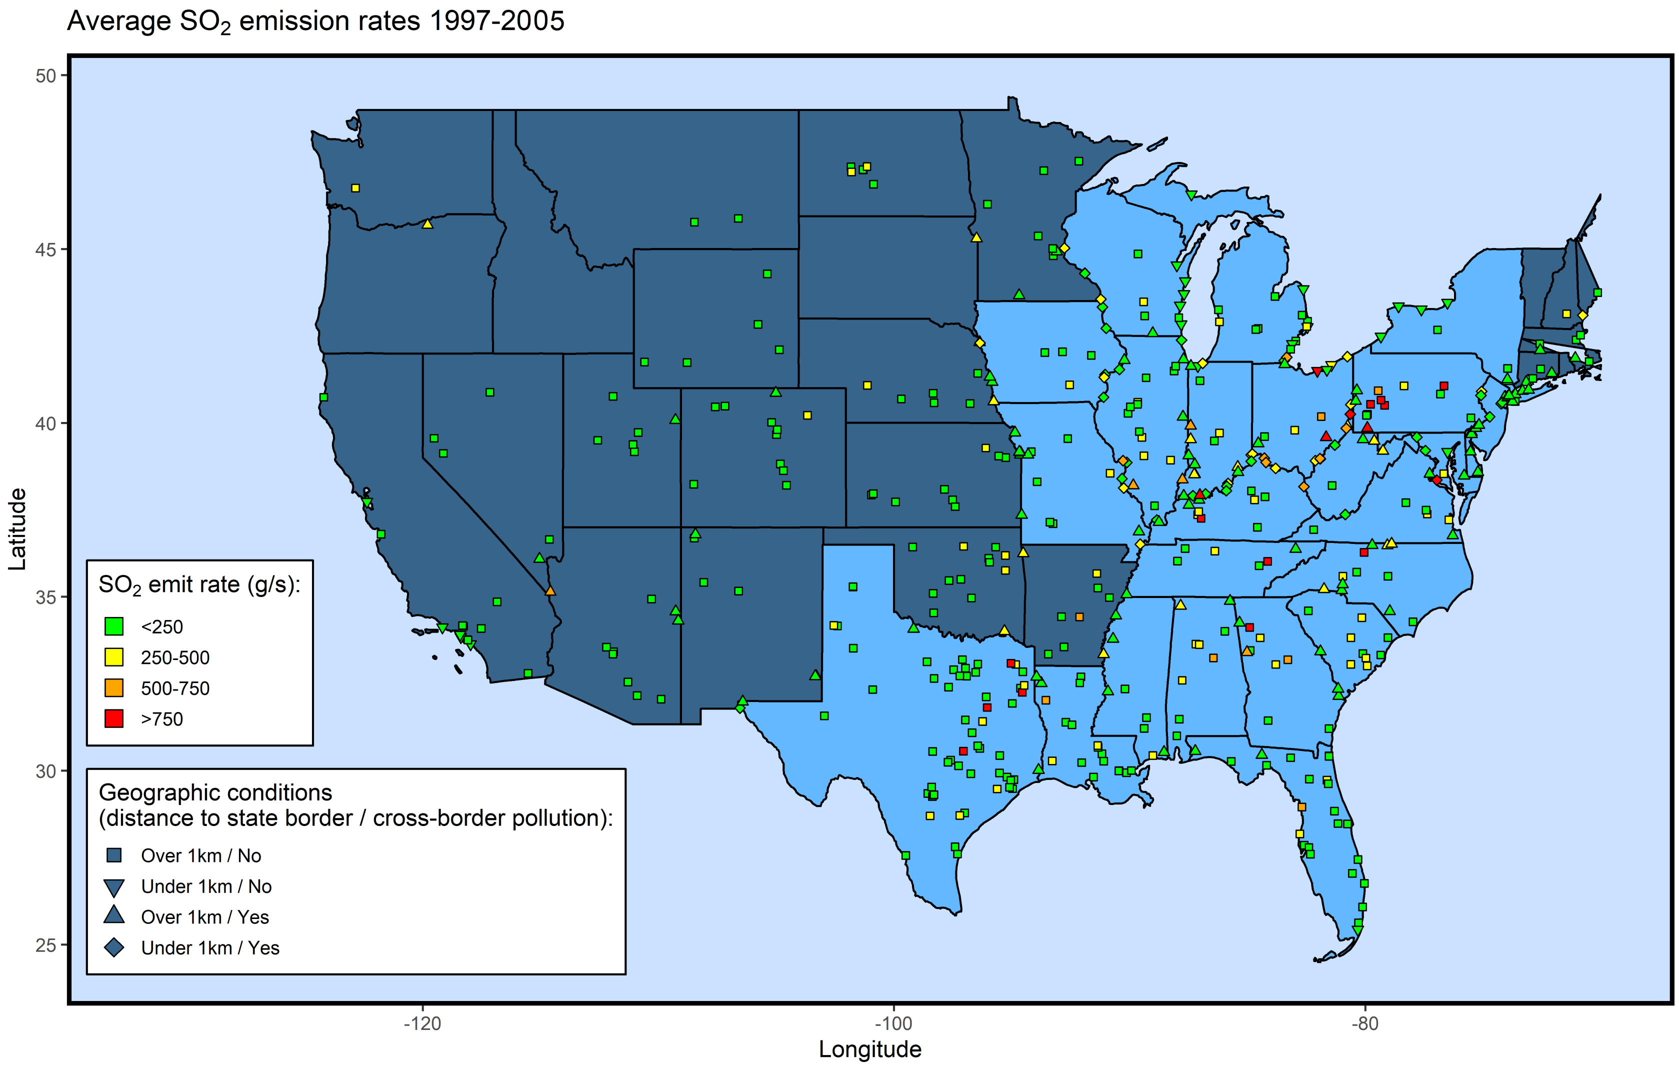

Figure 2.6: *Selection bias in assignment to the treatment group pre-2005*

2.6 Results

Event studies (figure 2.7) on the three main outcome variables (sulphur, cross-border sulphur, and carbon emissions) show that any observable pre trends are not statistically significant. Zero (or parallel) pre trends suggest that emissions in states that would be covered by CAIR were not on a different trajectory before 2005. These results support the null hypothesis that in the counterfactual state of the world (i.e. in the absence of CAIR), emissions in the two sets of states would not have evolved differently. While no definitive proof of the counterfactual exists, event studies showing zero parallel pre trends have often been used to support the hypothesis, including Barreca et al. (2021) and Fowlie et al. (2018). Figure 2.7 indicates a clear negative treatment effect for overall sulphur emissions, which suggests benefits on top of the Acid Rain Program reductions acknowledged in Chay and Greenstone (2003a) just before CAIR was announced, and more recently in Barreca et al. (2021). Moving on to carbon emissions, the event study shows no significant treatment effect from CAIR. While lagged means trend downward following the CAIR announcement, they never fall outside the 95% confidence interval around the null. This provides more convincing evidence that there were not other trends that differentiated abatement behaviour by firms in the CAIR region from others.

Finally, figure 2.7 shows the event study for our primary outcome of interest, which is denoted by a dummy variable indicating whether cross-border sulphur calculated with GAUSSMOD exceeds 1% of the NAAQS. The event study again shows a negative but less pronounced treatment effect from CAIR, where the announcement lowers the average probability that a treated plant transports at least 0.75 ppb to another state.

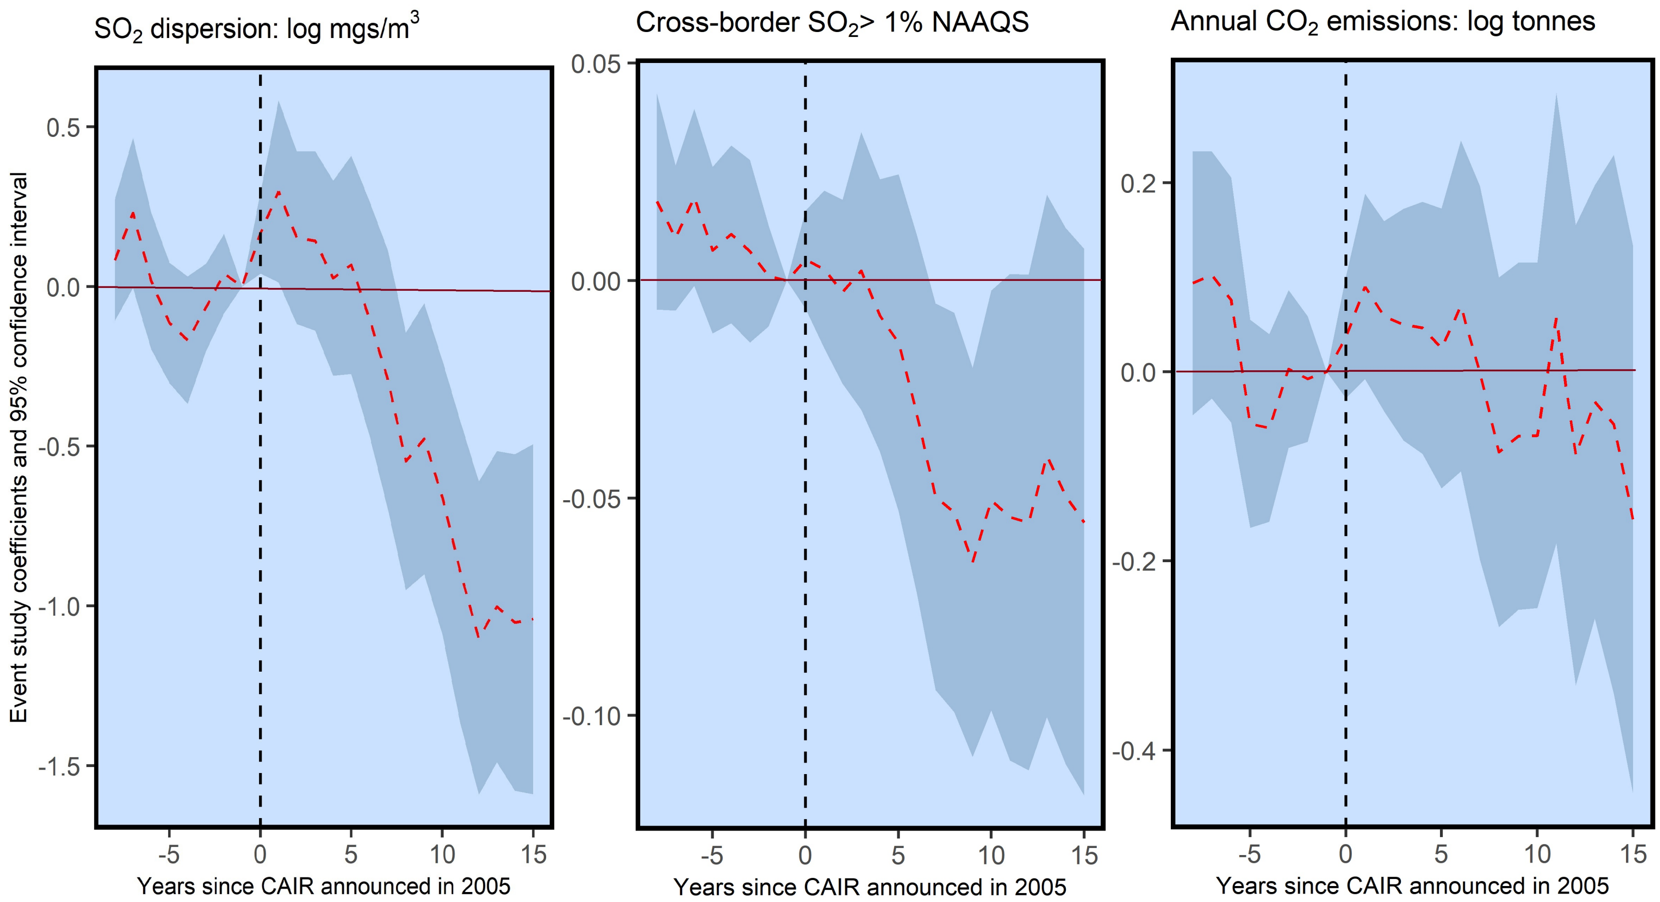

Figure 2.7: *Event studies for total annual SO_2 (left), cross-border SO_2 (middle) and CO_2 (right) with a treatment time at 2005*

Table 2.3 displays the regression estimates for the outcomes k in model (2.11). Unless otherwise specified the models are estimated using the Stata 17 `xtdidregress` command with heteroskedasticity robust standard errors clustered at the plant level (Pinzon, 2021) but the `fixest` package (Bergé et al., 2018) in `R 4.3` provides equivalent results. sulphur, carbon, heat input, generation and operation time are log-transformed to better fit the linear model ($R^2 \approx 0.25$) versus the original data ($R^2 \approx 0.1$). As suggested by the event studies, model (1) results in a significant difference-in-differences estimate of -0.24 interpreted as a $\approx 24\%$ reduction of sulphur emissions in CAIR states as a result of the policy. When the parallel trends assumption holds, the difference-in-differences estimator can be approximated as the ATT (Kahn-Lang & Lang, 2020) and is widely used for program evaluation.

Model (2) shows equation (2.11) with carbon emissions as the outcome. This model was estimated to evaluate the risk of concomitant treatment effects confounding the hypothesized causal effect of CAIR. The DD estimate for model (2) is -0.002 and is not statistically significant. The announcement of CAIR does not appear to have had affected pollutants not regulated by CAIR itself. Models (3) and (4) are the main equations of interest. The outcome in model (3) is average annual cross-border SO_2 ($\mu g/m^3$). The DD estimate is -0.02 and statistically significant. The result is that CAIR caused on average a $0.02\mu g/m^3$ reduction in cross-border SO_2 but should be cautiously interpreted. Recent research, e.g. Boulton and Williford (2018), has raised concerns about OLS with so-called semicontinuous outcomes where the data contains a large proportion of zeros. Unlike zeros resulting from censoring (Tobin, 1958), cross-border sulphur is highly skewed toward zero simply because many plants do not produce any cross-border pollution.

Binary logit or linear probability models (Buntin & Zaslavsky, 2004) have been proposed as solutions. Because logit coefficients are less easily interpreted, and the drawbacks of LPM are irrelevant in a difference-in-differences setting (prediction is not an objective), model (4) estimates the coefficients from model (3) with LPM. Its binary outcome takes the value one if the average cross-border SO_2 concentration from a plant i in year t exceeds 0.75 ppb, or 1% of the NAAQS (U.S. EPA, 2019), zero otherwise. The treatment effect is -0.03 and significant. The interpretation is that CAIR caused a 3% reduction in cross-border SO^2.

2.6.1 Heterogeneous treatment effects

Table 2.4 shows three specifications of the triple differences model designed to test hypothesis II. The triple difference estimator in model (1) is the regression coefficient for $(G_i \times CAIR_t \times C_{it}^{>0.75ppb})$ and is positive at 0.23. It suggests that the treatment effect from CAIR on average SO_2 emissions was 23% smaller among plants that transported at least 1% of the NAAQS (0.75 ppb) across state boundaries. This result supports rejection of hypothesis II, as the reduction in emissions following the implementation of CAIR was less pronounced among plants that transport a meaningful amount of SO_2 across state lines.

Sensitivity analysis: Distance to state border

Models (2) and (3) instead estimate the heterogeneous treatment effects among plants located at less than 10 and 20 kilometers from a state border, respectively. A plant's proximity to a state border is strongly but not perfectly correlated with the likelihood of producing cross-border SO_2 (0.41). It is plausible that moral hazard incentives arise not primarily from the cross-border emissions themselves but from the proximity to another state. For example, polluters may be unaware of their cross-border contribution, which a monitoring system attached to the flue stack cannot estimate, and use distance to borders as a proxy.

In table 2.4 I therefore also report DDD estimates for these two groups. For plants within 10 kilometers from a border, the DDD coefficient for SO_2 is positive and statistically significant at 0.33. Irrespective of cross-border SO_2 emissions, the abatement effect from CAIR was less pronounced among plants within 10km from the border. However, for model (3) the DDD coefficient is null. This heterogeneous treatment effect (potentially from moral hazard) does not appear to extend as much beyond 10km. These estimates arise from data further illustrated in figure 2.8, showing a smaller CAIR-associated treatment effect for plants closer to a state border. This is not due to plants close to the border starting off from a higher base rate of emissions. The correlation between emissions and border proximity is only -0.003 in the pre-CAIR period. Similarly, the post-CAIR reduction in average SO_2 emissions is lower among treated plants that transport more than 50% of their emissions across state lines, and the divergence with the control group diminishes.

The indicators of proximity to a state border do not account for wind patterns. Data on prevailing winds over a typical year is available from weather stations across the continental United States (see figure 2.1). I divide the data into increments of $15°$ and select the most frequently occurring increment (the mode) at the location of each power plant. A $360°$ direction denotes wind from north to south, $270°$ means west to east, $180°$ means south to north, and so on. I calculate the distance from each power plant to the state border in the direction of prevailing wind. Table 2.5 reproduces triple-difference models (2) and (3) in table 2.4 with the one difference that proximity to a state border is the downwind distance and not the nearest distance. I find that the treatment heterogeneity on the basis of downwind proximity to the state border is larger (ca 75%), given a regression coefficient of 0.56 (Halvorsen & Palmquist, 1980) added to the logarithm of emissions) when proximity is measured as downwind distance. This is compared to a heterogeneity of approximately 40% when nearest border distance is used.

Table 2.3: *Regression results*

	Continuous outcome			LPM
	(1)	(2)	(3)	(4)
Outcome	log(sulphur)	log(carbon)	cross-border SO_2	$> 0.75ppb$
DD	-0.24	-0.002	-0.02	-0.03
	$(0.05)^{***}$	(0.007)	$(0.002)^{***}$	$(0.007)^{***}$
log(Heat Input)	0.65	0.34	-0.01	0.014
	$(0.16)^{***}$	$(0.15)^{**}$	$(0.005)^{**}$	(0.017)
log(Operation Time)	0.64	0.08	0.008	0.004
	$(0.12)^{***}$	$(0.03)^{**}$	$(0.004)^{**}$	(0.016)
log(Permits)	0.04	-0.00	0.001	0.0014
	$(0.006)^{***}$	0.00	(0.001)	$(0.001)^{*}$
sulphur Control (%)	-0.58	0.16	-0.03	-0.07
	$(0.07)^{***}$	$(0.014)^{***}$	$(0.007)^{***}$	$(0.014)^{***}$
log(sulphur)		0.06	0.009	0.034
		$(0.004)^{***}$	$(0.001)^{***}$	$(0.002)^{***}$
R^2	0.83	0.99	0.70	0.84
within-R^2	0.25	0.94	0.06	0.10
Plant FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Obs.	$8,452$	$8,452$	$8,452$	$8,452$
Significance: $p < 0.01 : ***, p < 0.05 : **, p < 0.1 : *$				

Table 2.4: *Heterogeneous treatment effects*

	(1)	(2)	(3)
	Cross-border SO_2 $> 1\%$ NAAQS	Distance to border < 10 km	Distance to border < 20 km
Outcome variable	ln(sulphur)	ln(sulphur)	ln(sulphur)
Treatment effect for $C_{it} = 0$	-0.24	-0.31	-0.20
	$(0.06)^{***}$	$(0.06)^{***}$	$(0.06)^{***}$
Treatment effect for $C_{it} = 1$	-0.01	0.02	-0.25
Treatment heterogeneity	0.23	0.33	-0.05
	$(0.11)^{**}$	$(0.12)^{**}$	(0.11)
R^2	0.84	0.83	0.83
Within-R^2	0.29	0.24	0.24
Plant FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Obs.	$8,452$	$8,452$	$8,452$
Significance: $p < 0.01 : ***, p < 0.05 : **, p < 0.1 : *$			

C is a dummy variable indicating if plant i (1) contributes more than 1% of NAAQS across state borders, (2) is within 10 km from a state border, or (3) within 20 km from a border.

Table 2.5: *Heterogeneous treatment effects II*

	(1)	(2)
	Distance to border $<$ 10 km downwind	Distance to border $<$ 20 km downwind
Outcome variable	ln(sulphur)	ln(sulphur)
Treatment effect for $C_{it} = 0$	-0.33	-0.31
	$(0.06)^{***}$	$(0.06)^{***}$
Treatment effect for $C_{it} = 1$	0.23	0.05
Treatment heterogeneity	0.56	0.36
	$(0.14)^{***}$	$(0.13)^{***}$
R^2	0.83	0.83
Within-R^2	0.25	0.25
Plant FE	Yes	Yes
Year FE	Yes	Yes
Obs.	$8,452$	$8,452$
Significance: $p < 0.01 : ***, p < 0.05 : **, p < 0.1 : *$		

C is a dummy variable indicating if plant i *(1) is within 10 km downwind from a state border, or (2) within 20 km downwind from a border.*

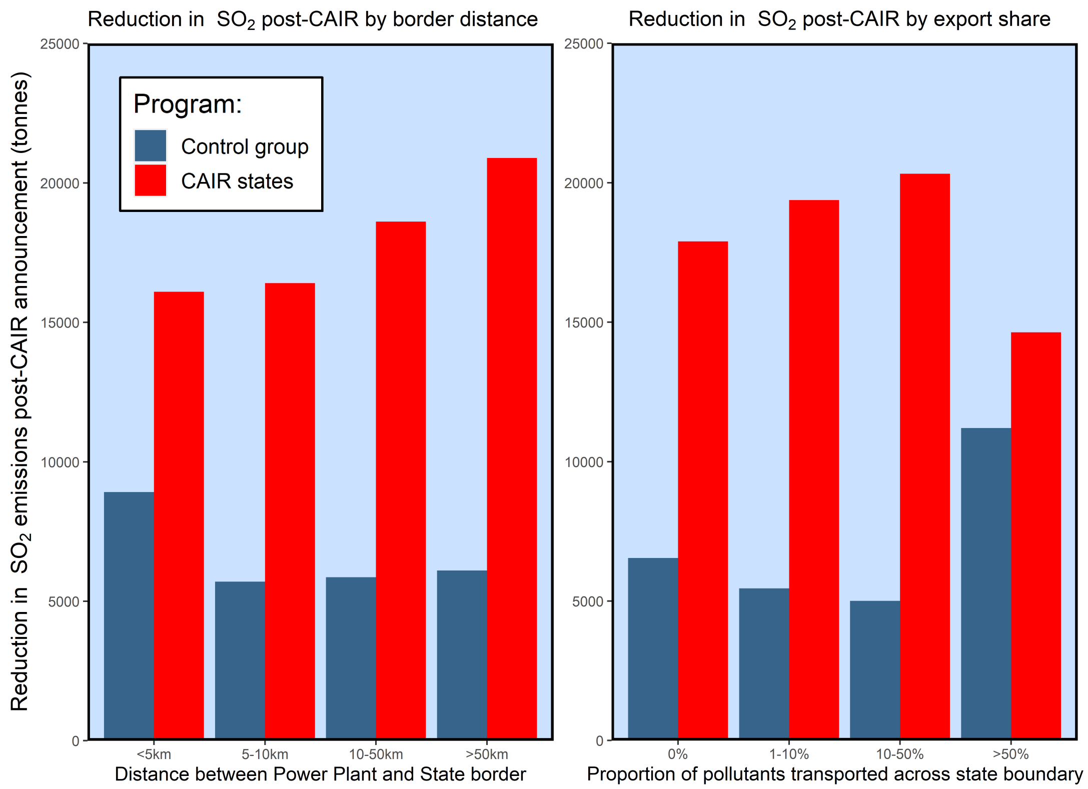

Figure 2.8: *Post-CAIR reduction in average SO_2 (%) emissions by distance to state border (left) and average cross-border pollution shares (right)*

2.7 Discussion and conclusion

While the Clean Air Interstate Rule was a regional program, its cap-and-trade mechanism was not spatially targeted. Following a U.S. court ruling against the Environmental Protection Agency in 2008, CAIR was vacated partly on the grounds that its design did not adequately protect downwind states against cross-border pollution. *North Carolina v. EPA* held that that the CAIR trading program went beyond the mandate of the Clean Air Act because the regional program did not address sources from one specific state contributing to nonattainment in another specific state.

EPA designed CAIR to eliminate pollution from out-of-state sources as a group, as summarized in Kruse (2009): "Pollution would be reduced regionally, but any state could buy enough credits to escape the requirement to reduce its impacton other states". However, if cross-border pollution primarily depends on overall emission rates, modelling CAIR on the successful Acid Rain Program may not have been a significant problem in practice. In this article I have evaluated this hypothesis and provided evidence against the argument that CAIR was ineffective at reducing interstate pollution. Using a novel combination of atmospheric dispersion modelling and difference-in-differences analysis, I support previous findings that CAIR was indeed successful in reducing overall sulphur emissions from covered sources (20-30%) due to a temporary rise in the price of permits, but also report a reduction in cross-border sulphur concentrations and the number of sources that transported sulphur across state lines. CAIR caused an average 2.3-3.7% reduction in the risk of exceeding 1% of NAAQS in a downwind state.

I support previous evidence (Glasgow & Zhao, 2017; Heo et al., 2023) that cross-border emissions are partly driven by geographic factors, most importantly the distance of the source from a state border, and also annual weather trends as I discover that there are plants several kilometers from a state border, yet contribute to downwind sulphur pollution in another state.

I add to this literature by quantifying cross-border pollution using a custom Gaussian dispersion model and showing that concentrations are universally below the national air quality standards (NAAQS) set by the EPA, although states around the former coal-mining belt of Kentucky, Indiana, Ohio, and West Virginia (see figure 2.6) share many high-emission sources along their borders.

By computing the contribution of cross-border pollution from each plant using GAUSSMOD I uncover that moral hazard may have de-fanged the effectiveness of CAIR for certain plants. The reduction in overall annual sulphur emissions caused by the CAIR announcement was weaker among affected plants that contributed more than 0.75 ppb of cross-border SO_2 concentration.

Additionally, this weaker treatment effect extends to plants within 10 kilometers of a state border, even though less than 50% contribute to cross-border non-attainment. A possible mechanism to explain this phenomenon is the way SIPs (see section 2.2) are applied. States submit SIPs to the EPA outlining their plans to achieve air quality targets *within their state* and regulations in the SIPs are generally enforced by the state. While section 126 petitions have increased over the past five years (Gerrish, 2020), states may be less motivated to regulate pollution which leaves its borders. However, my results also indicate that this moral hazard may be primarily driven by proximity to the state border, not knowledge about cross-border contributions itself.

My results provide new nuance to the arguments that led to the vacation of CAIR in the 2008 *North Carolina v. EPA* case. On the one hand, average cross-border SO_2 declined as a result of CAIR. On the other hand, the decline was considerably smaller than that of overall emissions ($2-4\%$ versus 24%). In addition, SO_2 emissions from plants that did contribute to cross-border concentrations appear less affected by CAIR, as were plants within 10 kilometers from a state border. Moral hazard can be prevented by monitoring not only emissions at the source but also cross-border transport, for example using the EPA's AERMOD dispersion model which inspired GAUSSMOD.

A trading ratio can be applied to the permit market in which a purchasing plant faces a higher (lower) price reflecting the relatively higher (lower) propensity for cross-border pollution vis-a-vis the seller Holland and Yates (2015). Acknowledging the geographic moral hazard problem is particularly important in settings where regional regulators have less incentives to collaborate. For example, Heo et al. (2023) find that trans-boundary air pollution from China significantly increases mortality and morbidity in South Korea. Even within China, Cai et al. (2016) find that provincial governments respond to pollution reduction mandates by shifting their enforcement efforts away from the most downstream county, from where pollution is directly transported into another province. A regional cap-and-trade program across East Asia or the ASEAN region would likely suffer from similar likelihood of moral hazard. A permit market with spatially explicit trading ratios based on downwind risk might help manage these concerns.

Appendix A: Matched Controls

Propensity Score Matching (PSM) is a statistical technique used in observational studies to estimate the effect of a treatment or intervention by reducing bias that arises from confounding variables. It is a common augmentation to difference-in-differences estimation. In natural experiments where assignment to the treatment group is not random (CAIR targeted states with many high-risk coal-fired plants), it is helpful to control for differences between treated and control groups that may influence the outcome. PSM works by matching power plants in the treated group with plants in the control group that have similar characteristics, as determined by the propensity score (Jia et al., 2021). The propensity score is the predicted probability of belonging to the CAIR group, which is estimated in equation (2.13):

$$\ln \frac{Pr(G_i = 1)}{Pr(G_i = 1) - 1} = \beta_1 \times BorderDistance + \beta_2 \times HeatInput + \beta_3 \times \delta_{i,S_i} \tag{2.13}$$

By implementing PSM within a DD framework, it is possible to further control for time-invariant confounding variables and ensure that the estimated treatment effect is more robust. In this case, the distance between the power plant and the state border, the base heat input of the plant's generators, and the proportion of SO_2 emissions δ that remain within the home state S_i of firm i.

Matching of control plants to treated plants was done on pre-CAIR observations from 2004. These variables most strongly predicted assignment into the treatment group using a generalized linear probability model. Propensity score matching was performed using the `MatchIt` package in `R`. The package attempts to match plants in the treatment group with controls based on their similarity. As not all treated plants could be matched to a suitably similar control, the sample in table

2.6 is a smaller balanced panel of 188 plants across 24 years. The results directionally agree with those reported in section 2.6.

Table 2.6: *Regression results*

	Continuous outcome			LPM
	(1)	(2)	(3)	(4)
Outcome	log(sulphur)	log(carbon)	cross-border SO_2	$> 0.75ppb$
DD	-0.18	-0.002	-0.05	-0.16
	$(0.05)^{***}$	(0.015)	$(0.003)^{***}$	$(0.01)^{***}$
log(Heat Input)	0.30	0.43	-0.01	0.07
	$(0.12)^{***}$	$(0.11)^{***}$	$(0.005)^{**}$	$(0.02)^{***}$
log(Operation Time)	1.03	0.12	0.008	0.004
	$(0.14)^{***}$	$(0.06)^{*}$	$(0.004)^{**}$	(0.016)
log(Permits)	0.05	-0.00	0.001	0.0014
	$(0.008)^{***}$	0.00	(0.001)	$(0.001)^{*}$
sulphur Control (%)	-0.53	0.15	-0.03	-0.07
	$(0.09)^{***}$	$(0.018)^{***}$	$(0.007)^{***}$	$(0.014)^{***}$
log(sulphur)		0.07	0.009	0.07
		$(0.006)^{***}$	$(0.001)^{***}$	$(0.004)^{***}$
R^2	0.86	0.97	0.73	0.85
$within - R^2$	0.35	0.81	0.15	0.21
Plant FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Obs.	$4,290$	$4,290$	$4,290$	$4,290$
Significance: $p < 0.01 : ***, p < 0.05 : **, p < 0.1 : *$				

Appendix B: GAUSSMOD validation

Quality assurance for the GAUSSMOD model was carried out by testing that its results agreed with theoretically predicted SO_2 dispersion given a set of input parameters. Figure 2.9 displays average monthly dispersion around a randomly selected power plant from within the sample. Legends show the average (median) wind direction measured at the most proximate weather station. Wind direction is presented in degrees from true north, such that e.g. a $0/360°$ direction is north-to-south, $90°$ is east-to-west, $180°$ is south-to-north, etc. Figure 2.9 shows that SO_2 disperse in the direction of prevailing wind, which supports correct implementation of geometry modules within GAUSSMOD. Figure 2.9 also shows that prevailing wind patterns are relatively stable over the year. This reduces uncertainty around the annual estimates of cross-border pollution made in this chapter.

The theoretical predictions from a correct implementation of GAUSSMOD were also tested by regressing simulated SO_2 at sites of nearby EPA monitoring stations (within 20 kilometers) on variables that are expected to drive dispersion. The regression results are shown in table 2.7. The outcome in models (1) and (2) is the simulated SO_2 at the EPA monitor site, excluding and including plant fixed effects respectively. The outcome in models (3) and (4) is monthly average SO_2 measured at the monitor. The EPA monitor measure is not directly comparable to the GAUSSMOD measure because the former records ambient SO_2 from all sources. The coefficient for temperature gradient (between the exit flue gas and ambient air) is negative in the GAUSSMOD models. This is expected, as gasses with large temperature gradients rise faster. The coefficient for smoke stack height is similarly expected to be negative for the same reason. A positive coefficient for wind speed is expected for monitor sites located downwind of the power plant.

The downwind dummy is positive for the GAUSSMOD models, corroborating the visuals in figure 2.9 showing that the pollutant is transported in the wind direction. The plant fixed-effects captures variation in plant coordinate accuracy. At small distances, relatively minor geolocation errors may produce inaccurate angles between the wind vector and the plant-monitor site vector.

Table 2.7: *Regression results*

	(1)	(2)	(3)	(4)
Outcome	GAUSSMOD SO_2		EPA SO_2	
Temp gradient (°C)	-0.003	-0.004	-0.008	0.064
	(0.001)***	(0.008)	(0.001)***	(0.033)**
Wind speed (m/s)	0.049	-0.046	0.878	0.016
	(0.015)***	(0.045)	(0.059)***	(0.146)
Stack height (m)	-0.001		-0.005	
	(0.0003)***		(0.0009)***	
Downwind dummy	1.424	0.438	0.18	-0.58
	(0.542)***	(0.151)***	(0.405)	(0.175)***
R^2	0.255	0.515	0.424	0.781
$within - R^2$	0.235	0.138	0.356	0.084
Plant FE	No	Yes	No	Yes
Month FE	Yes	Yes	Yes	Yes
Obs.	228	228	228	228
Significance: $p < 0.01 : ***, p < 0.05 : **, p < 0.1 : *$				

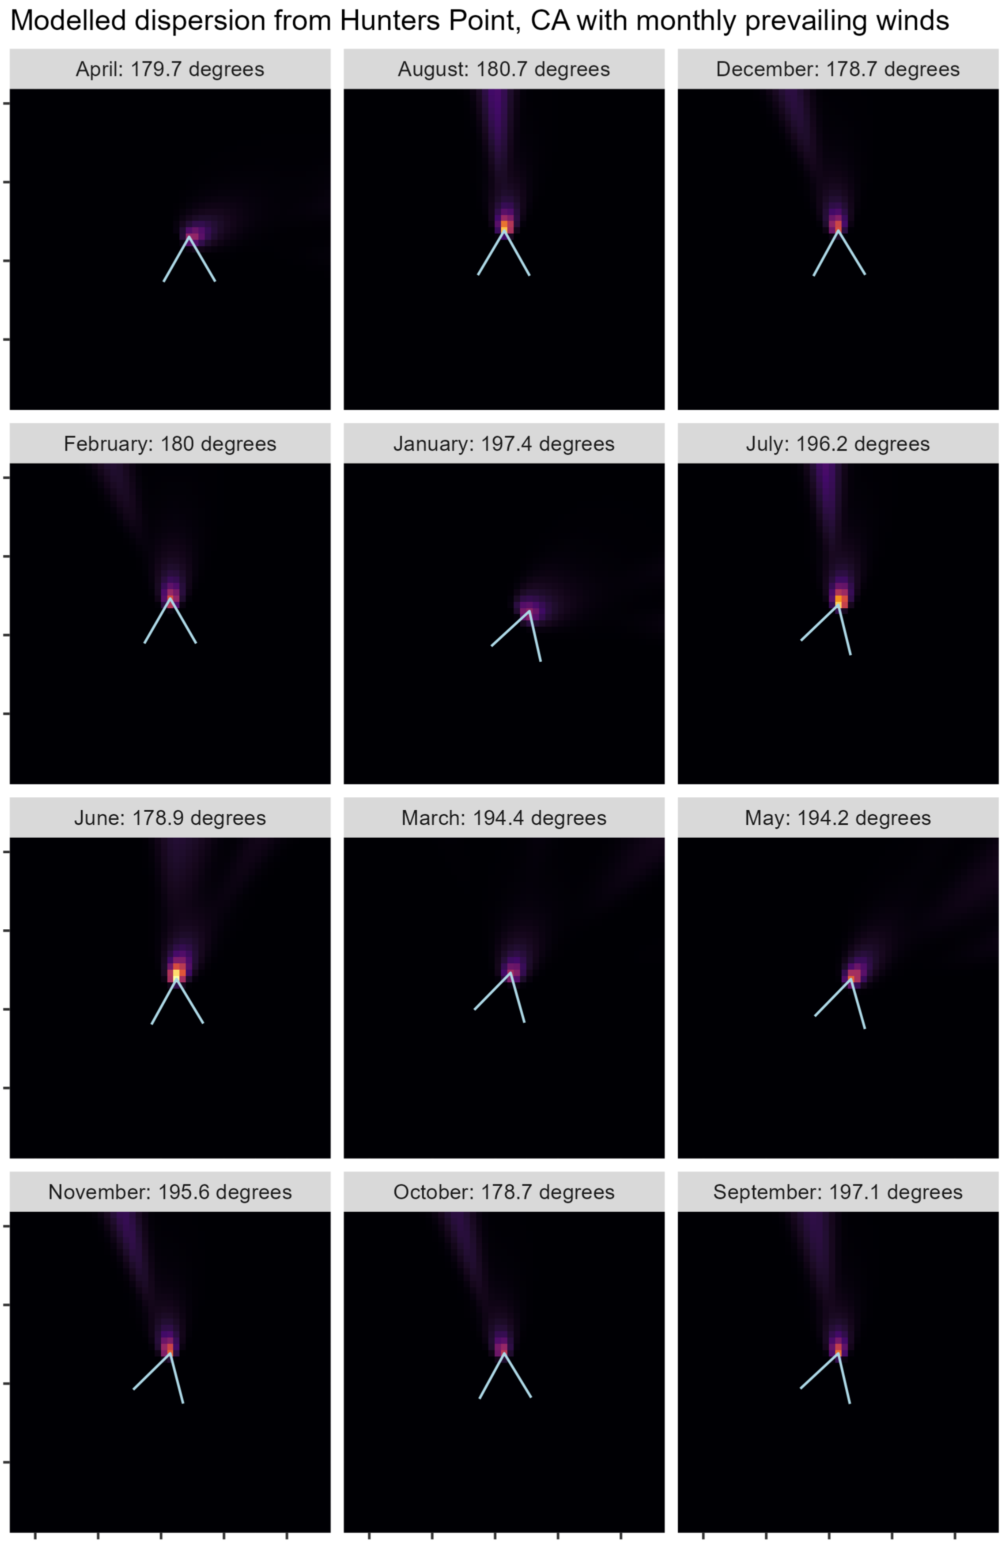

Figure 2.9: *Monthly simulated SO_2 using GAUSSMOD from the Hunters Point power plant in San Francisco, California. The average monthly wind direction is added as a blue arrow overlay.*

Chapter 3

Evaluating environmental land management using hypothetical choice experiments

3.1 Introduction

This chapter provides the necessary policy background which motivates the research presented in chapters 4 and 5. Each chapter is concerned with a particular class of environmental regulation and subsidy mechanism: Environmental land management (ELM) schemes. Such schemes compensate farmers to engage in land management actions that drive environmental outcomes such as reduced pollutant runoff (Kampas et al., 2013), flood mitigation (Holstead et al., 2017; Kenyon, 2007; Reaney, 2022), or increased biodiversity (Image et al., 2022, 2023).

Both chapters explore a common set of simulated spatial configurations for hypothetical ELM actions within a real English agricultural landscape. Chapter 4 evaluates the hypothetical actions by their potential to reduce surface water runoff and catchment flooding. It proposes a market for trade in ELM contracts and studies the impact of transaction costs. Chapter 5 evaluates their impact in terms of providing pollination services by creating insect habitats. It adds to the analysis by adding a coordination bonus to the hypothetical ELM offering (Banerjee et al., 2017), where farmers can receive additional compensation to improve habitat connectivity by collaborating with a neighbour.

Across the two chapters, the research questions outlined in chapter 1 are analysed empirically using data collected from farmers in England using survey methods. The rest of this chapter is structured as follows: The first section sets out the background and current state of ELM schemes in England. The second section describes the farmer survey, sets out the sampling strategy, describes the sample, and evaluates its representativeness for the farming population in the region.

The third section sets out three hypothetical discrete choice experiments (DCEs) that surveyed farmers were invited to participate in. Finally, the fourth section guides the reader through three approaches to analyse DCE results and motivates the modelling choices. This chapter does not report any results. Hypothesis tests and environmental modelling for cost-effectiveness analysis in terms of flood risk reduction and pollination services are presented in chapters 4 and 5, respectively.

3.2 Environmental Land Management in England

In the 1980s, the UK government began offering ELM schemes in response to growing environmental concerns, particularly around the use of pesticides and their impact on biodiversity. Soon, these schemes would be directed by the EU's Common Agricultural Policy (CAP). Agricultural- and environmental policy in the UK is devolved to its four nations and following Brexit, each nation is developing new agricultural policy and payment schemes (Clements et al., 2021). Farmers and land managers in England can enrol in ELM schemes offered through Defra that pay farmers ”to deliver, alongside food production, significant and important outcomes for the climate and environment that can only be delivered by farmers and other land managers in the wider countryside” (Defra, 2022). The ambition of Defra is to increase participation to 70,000 land managers by 2028, covering 70% of farmed land and 70% of all farms, although Clements et al. (2021) observe that the National Audit Office has raised doubts about the likelihood of timely completion.

During the period of data collection for this research (summer of 2022), two separate schemes have been on offered to eligible land managers in England. The Sustainable Farming Initiative (SFI) seeks to provide payment for simple projects that are possible for the majority of land managers to take on with minimal guidance, while the Countryside Stewardship (CS) scheme focuses on more targeted interventions relating to specific habitats and features that can be done alongside food production (Defra, 2022). SFI contracts last for three years, and tenant farmers do not need landowner permission to enrol. The durations of CS contracts vary depending on the intervention. The SFI and CS schemes are individual commitments, although a separate Landscape Recovery scheme involves groups of land managers and farmers working together to deliver a range of environmental benefits across farmland and rural landscapes (Defra, 2022).

Enrolment into the schemes worked as follows; a) a land manager selects the land parcels they would like to enter into the scheme using digital maps from the Rural Payments Service under Defra, b) they authorise a Defra agent to submit an application, c) site visits by the regulator are planned to assess how the environmental aims are met under the options in the agreement, and d) payments are made according to rates shown in figure 3.1. Farmers can participate in and receive payments from both the SFI and the CS schemes, although compensation may not be paid twice for the same action on the same land parcel via different schemes.

Minimum durations for actions in both schemes vary; contracts involving small-scale interventions normally last for five years, while more comprehensive actions such as planting trees last for ten years. It is important that schemes are designed to produce the desired environmental outcomes, as the most important interventions put in place at this time are 'locked-in' for up to a decade.

A final important aspect of the policy background is the Basic Payment Scheme (BPS) which is a general government grant to farms available both for productive and retired land enrolled in ELM. Land is eligible for the BPS if it is agricultural land (arable, permanent grassland or permanent crops), used primarily for an agricultural activity for the whole of the relevant calendar year. Defra aims to phase out this payment in favour of action-based payments such as the SFI and the CS schemes. Tyllianakis et al. (2023) suggest that the credibility of the government's position on the BPS is a reason that participants in their survey displayed a strong preference to enrol in an ELM scheme.

Source: Defra (2023), "*Countryside Stewardship: get funding to protect and improve the land you manage*", accessed at https://www.gov.uk/guidance/countryside-stewardship-get-funding-to-protect-and-improve-the-land-you-manage on 21-07-2023

Figure 3.1: *Overview of two actions offered through the Countryside Stewardship scheme*

There already exists academic evidence on the drivers of ELM participation. Mamine, Minviel, et al. (2020) summarise 79 DCE studies (including 33 from Europe) on ELM uptake among farmers. The expected income from the contract was a significantly positive predictor of uptake in 15 out of the 18 experiments that tested for it. The positive effect from an up-front payment upon signing the contract is similarly conclusive. Eight out of eleven studies also show a significantly negative effect of clauses involving collective commitments (such as coordinated placement of natural features (Kuhfuss et al., 2016)) from a pool of farmers on the same contract. This could reflect both coordination costs and reputational costs from deviating from the collective (Franks, 2011). UK farmers are aware that some environmental externalities like flood risk is not driven by practices on individual farms (Holstead et al., 2017) and have shown high endorsement in principle of higher pay for greater effort, rather than external circumstances (Loft et al., 2020). Perceived inequity can threaten participation.

The evidence provides guidance on how to effectively promote ELM schemes; a) compensation payments should be significant enough to change farmers' traditions and inertia, and offers must therefore be targeted where the environmental benefits are clear, b) information and advise should be offered, c) policy needs to internalise externalities from farm actions, and d) analysis of the distributional effects for planned schemes should recognise the equity concerns that farmers may perceive. These factors were considered in the design of the hypothetical ELM schemes.

3.3 Simulations of ELM schemes

Hypothetical landscapes were simulated by changing the land use class of individual groups of pixels in the UKCEH land cover data from agricultural land to ELM features. These features come in two types, inspired by actions funded via the Countryside Stewardship scheme: a) Natural regeneration, which involves a reversal from agriculturally productive land use into unimproved permanent grassland. On former cropland or fallow, grasses and/or flowers are to be sown and left alone. b) Broadleaf trees, including fruit trees recommended in (Image et al., 2023), which involves planting, fencing, and maintaining trees.

The landscapes were representative of the Eden catchment in the north-west of England. The Eden catchment is a largely agricultural landscape shaped by upland peat and fells feeding a fertile sandstone valley and floodplain. Farming (predominantly sheep, dairy and grassland grazing, with limited arable) dominates land use and local identity. That farming both supports the local economy and creates environmental pressures (nutrients, sediment, flood risk), which local partnerships and stewardship schemes are actively addressing (Eden River Trust, 2025).

Natural features were placed across samples from the catchment in four different spatial configurations. Examples of these configurations are shown in figure 3.2. The first variant (upper left) is corridors along field edges, where a field is defined as a contiguous patch classed as either cropland (cereals, soybeans, etc.), fallow, or grassland used for grazing. The second variant (upper right) is in-field corridors, in which the features are placed in straight lines across the fields. Such in-field corridors are more disruptive to farming operations, as they take more productive land out of production and obstruct thoroughfare with tractors and other machinery. The third variant (lower left) is in-field "isles", disconnected patches, 10 to 20 meters wide, distributed evenly across the field. More permeable than the in-field corridors, the isles can nonetheless cover larger fields, while retiring significantly less land. Finally, the fourth variant (lower right) shows a larger contiguous patch placed randomly in the landscape.

The total area of natural features in each simulation is governed by adjusting the gap between isles and corridors. The contiguous patch is drawn to match the combined area of the in-field and field edge corridors. The in-field isles will cover a smaller area than the corridors and patch at each gap. Larger gaps between natural features mean less farmland taken out of production and less need for coordination between farmers, at the expense of fewer habitats and less connectivity. In this chapter gaps between 200 meters and 1,500 meters are simulated.

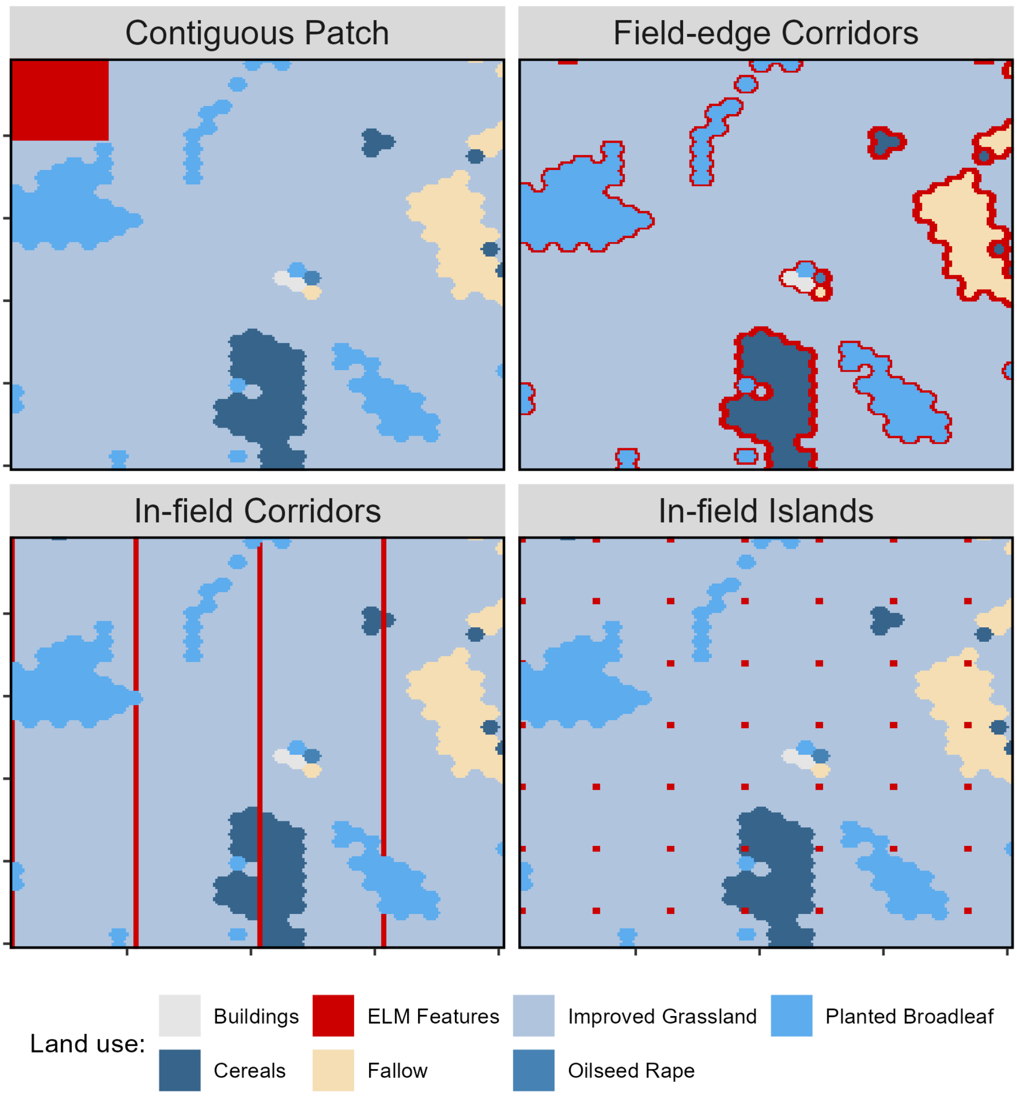

Figure 3.2: *Spatial configuration of NFM features: Upper left shows the contiguous patch; upper right shows field-edge corridors; lower left shows in-field corridors; lower right shows in-field islands*

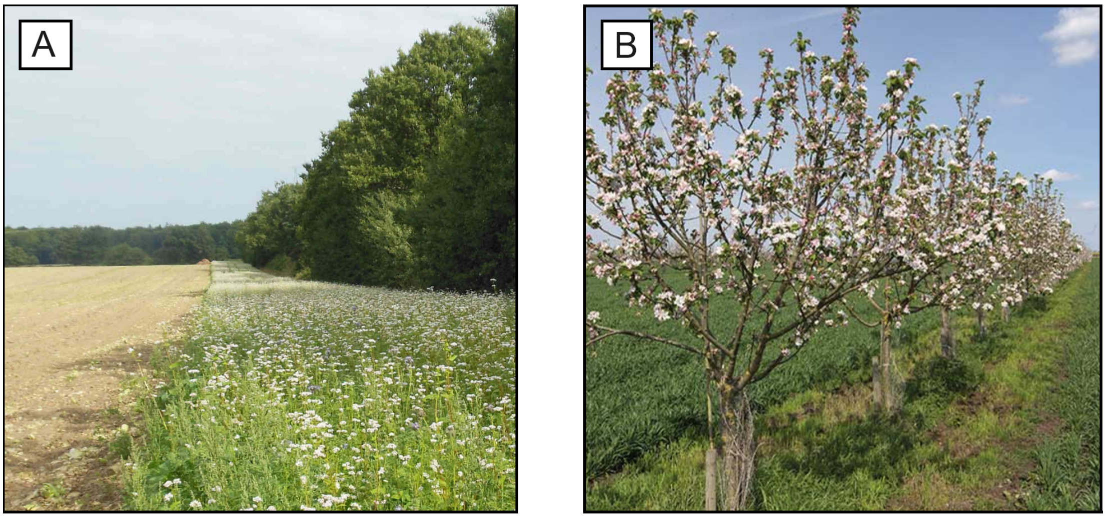

Figure 3.3: *Photographs of natural regeneration along field edges (A) and in-field rows of flowering fruit trees (B) in an agricultural landscape in the UK (Image et al., 2023)*

3.4 Survey and sample characteristics

Respondents for the farmer survey were recruited from several counties in the north of England and initially contacted via mail. Addresses of probable farmers were hand-collected from council authorities by searching for and recording entries that included the word *farm*. Access to electoral registers containing the addresses was granted for research purposes. The author personally visited the Eden District Council in Penrith and the Durham County Council in the city of Durham. A research assistant was hired to visit the Northumberland County Council in Morpeth and the North Yorkshire Council. Invitations to participate in the survey were sent out to approximately $3,100$ unique addresses with a brief written introduction to the research project, the authors' contact information and an offer of a £50 payment to compensate them for their time if electing to complete the experiment. Funding to cover compensation was gratefully received via a joint grant from the UK Natural Environment Research Council and the Economic and

Social Research Council, as well as a Durham University Seedcorn grant. This was a promised payment, not to be paid until the individual had completed the survey.

Slonim et al. (2013) document a so-called opt-in bias in participation in economic experiments where individuals with more leisure time, a greater interest in economics and science, and that are more pro-social than average are more likely to participate in experiments. Cash payments for participation can reduce the opt-in bias along intrinsic motivators by introducing a competing extrinsic incentive (Groves et al., 2000). This leaves open channels for bias in terms of differences in the economic value of time between individuals. However, these differences are easier to observe and control for using a battery of socioeconomic control questions. The payment is also more generous (ca five times the UK minimum wage for a $30 - 45$ minute experiment) compared with the $\$30$ (twice the local minimum wage) offered by Slonim et al. (2013) for a similar commitment. It has been suggested that a higher payment reduces opt-in bias by incentivising a broader segment of the population (Slonim et al., 2013).

Invitations to participate were mailed out in three rounds covering the separate segments of the sampled geographies approximately three weeks apart. In addition, reminders were sent out following each round to farms that had not responded to the initial mail-out. Interested individuals contacted the researchers to receive a link to an online survey. The questionnaire was created using the surveying software Qualtrics (Qualtrics, 2020).

While most surveys were completed remotely online, 36 surveys were also administered in person to include farmers who were unfamiliar with web-based survey participation. These were either conducted in focus groups or individually at the respondent's home. In-person surveys were more costly but could reach a wider set of respondents and clarify any ambiguities in the survey presentation (Johnston et al., 2017).

717 persons responded to the survey, and 494 persons completed it. A further 67 respondents were dropped due to failure to answer necessary socio-demographic control questions, leaving a sample of 427 individuals. Farm sizes range widely between 2 and 2000 hectares with Northumberland and County Durham hosting the highest concentration of large farms. The average farm size is 233 hectares. 26% of respondents were female, and the average age in the sample was 54 years. 76% of respondents said that farming was their primary source of income, and 55% were currently enrolled in an ELM scheme. Summary statistics are shown in table 3.1.

Postcodes for the farms were collected from the respondents. I used the UK Office for National Statistics (ONS) postcode directory to extract latitude and longitudes (Reid et al., 2017) in order to approximately geolocate the farms. Figure 3.4 shows the sampling area and land endowments of farms in the final dataset. The distribution of farm sizes defined as hectares of productive land has long tails with extreme values at both top and bottom ends. The smallest stated farm size is one hectare ($10,000$ m^2) and the largest is 5000 hectares, which puts it among the largest in the UK (Lowenberg-DeBoer et al., 2019). Respondents are comparably older, with a median age of 57, compared to 40-41 years for the overall UK population and mostly male (73.6% of the sample). However, the sample's demographics are relatively representative of UK farmers; 70% of UK farm holders were above 55 years of age and 84% were male (UK Department for Environment, Food and Rural Affairs, 2022).

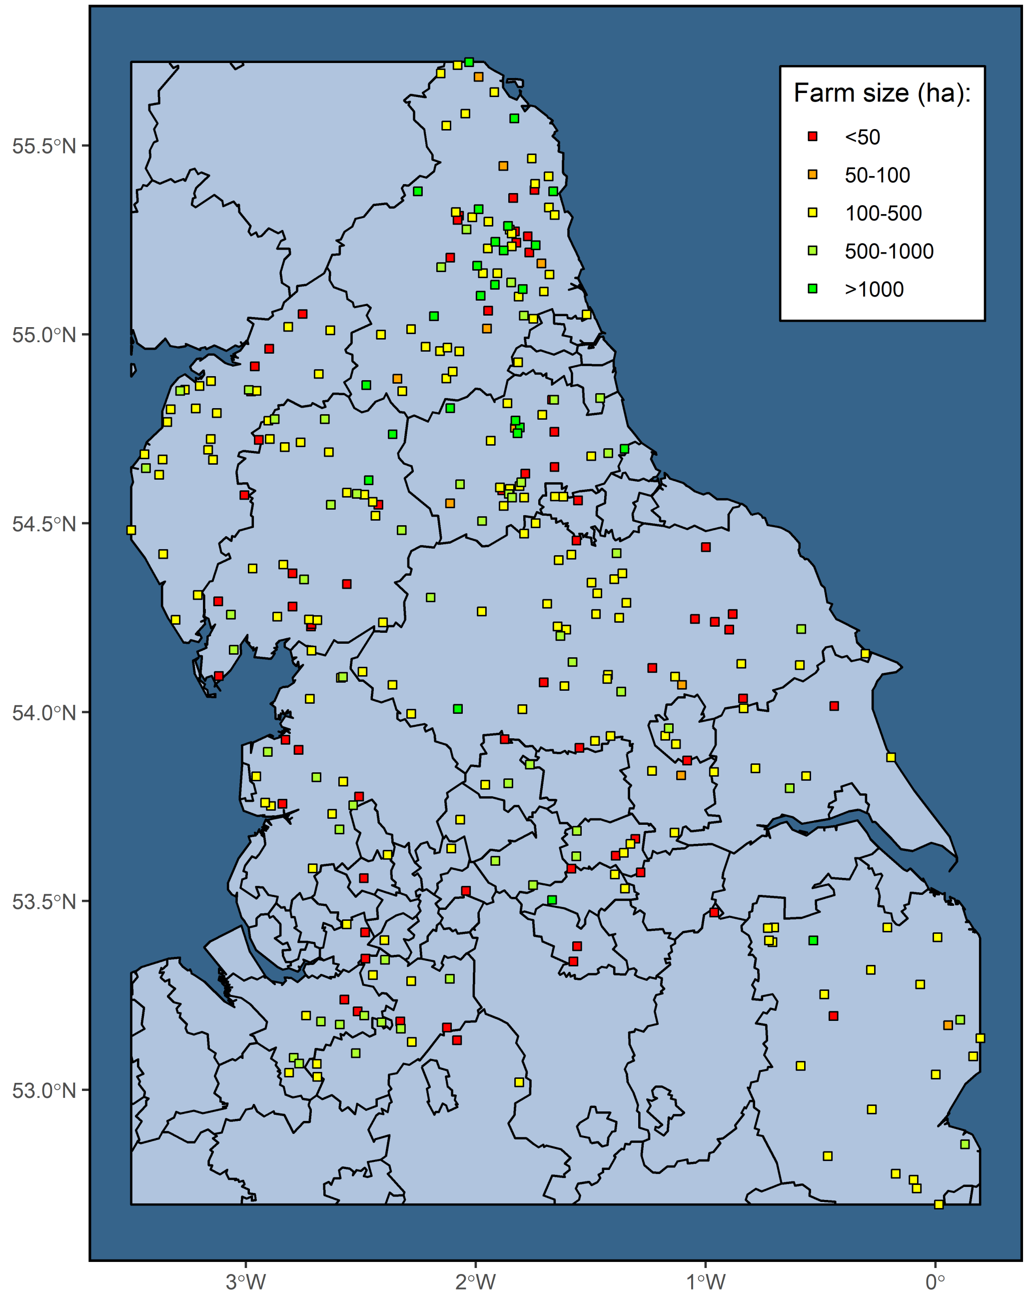

Figure 3.4: *Sampling area and respondent farm land endowment in hectares*

73% of respondents had received at least 11 years of formal education and 33.5% held a university degree. 53.5% of respondents were enrolled in an ELM scheme at the time of the survey. This allowed me to control for familiarity with similar schemes. When asked to state their own level of involvement with their local farming community, only 15.9% of respondents rated their involvement as quite strong or very strong. Hurley et al. (2022) identify low social capital as one barrier to involve segments of the farming community in the design and uptake of ELM schemes. Individuals with low social capital can be isolated from their peers and government, making it less likely for their behaviour to be influenced by others. Almost one fifth (18.7%) of participants in the sample rated their community involvement as very weak compared to their peers.

Answers to the survey lends support to the thesis presented in Hurley et al. (2022). Respondents who rated their social involvement as weak were significantly less likely to currently be enrolled in a Defra scheme compared with their more socially connected peers. Figure 3.5 shows that the proportion of ELM uptake increases with the self-rated community involvement. In addition to lower costs, one advantage of the online survey is that participants can complete it at home, which increases the likelihood of reaching socially isolated farmers. Still, it must be assumed that some selection bias remains and that this isolated demographic is under represented in the sample.

A positive correlation between community engagement and ELM participation may be interpreted as indicative of a link between social connections and prosocial attitudes more broadly. ELM projects are meant to provide public goods such as habitat conservation and flood management.

Figure 3.6 breaks down the distributions of expressed concern about flooding in local agricultural catchments by levels of community engagement. There does not appear to be a strong overall correlation, although the proportion of respondents stating that they are very concerned about catchment flooding is higher in the highly socially connected group.

Table 3.1: *Summary statistics*

	Minimum	Median	Mean	Maximum
Age	19	57	54	91
Farm size (ha)	1.0	130.0	146.3	5,000
Farm tenure (years)	1.0	33.0	31.8	73.0
No. tracts of land	1.0	1.0	2.5	40.0
% grazing	0	23.5	39.5	100
No. farm neighbours	0	4.0	5.1	30.0
Highest education		% of respondents		
	GCSEs (11 years)	17.6%		
	A-levels (13 years)	21.9%		
	UG Degree	27%		
	PG Degree	6.5%		
	None of the above	27%		
NFM Priority		% of respondents		
	Low	30.7%		
	Medium	31.6%		
	High	32.3%		
	Missing data	5.4%		
ELM participation		53.5%		
Sharing farm equipment		45%		
Women		26.6%		
Community involvement		% of respondents		
	Very weak	18.7%		
	Quite weak	24.2%		
	About average	41.1%		
	Quite strong	12.7%		
	Very strong	3.2%		

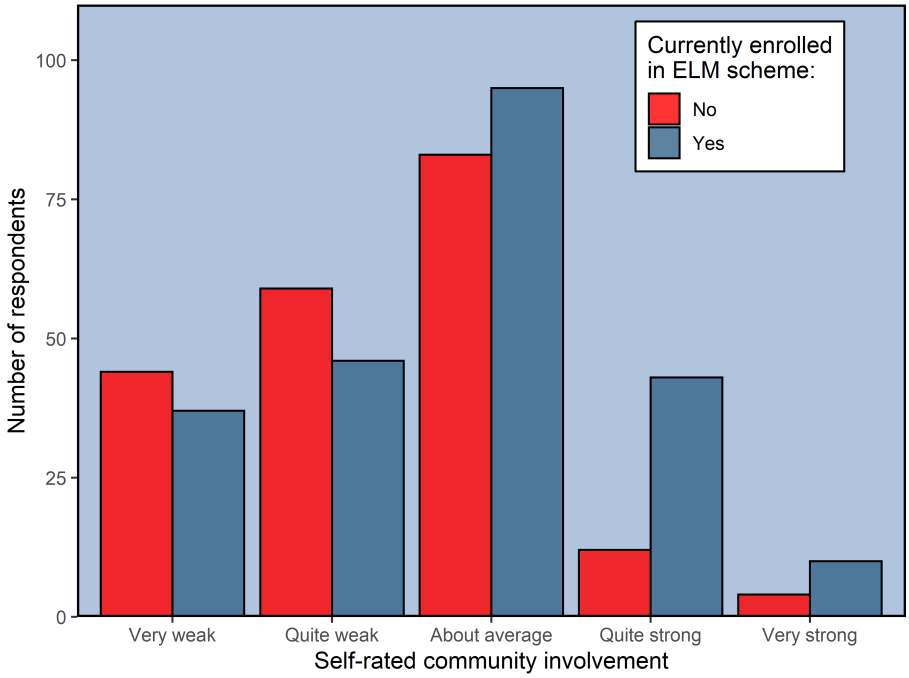

Figure 3.5: *Distribution of ELM enrolment by respondents' community involvement rating on a Likert scale*

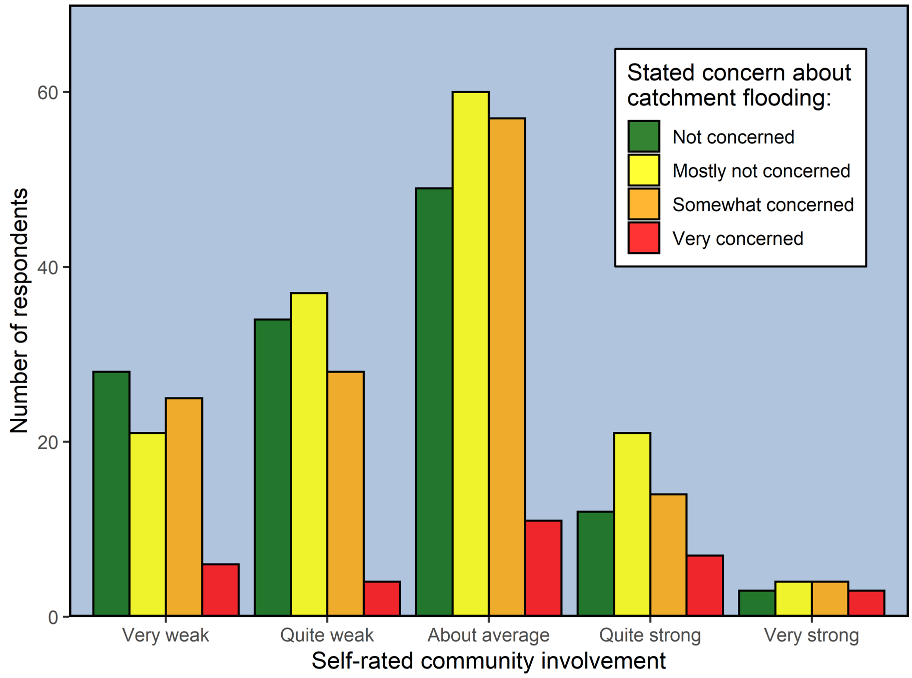

Figure 3.6: *Distribution of stated concern about catchment flooding by respondents' community involvement rating on a Likert scale*

3.4.1 Regional representativeness

Table 3.2 illustrates to what extent the surveyed farmers were representative of the farming population in the region. A common theme across all regions is that farms in the sample managed more land than the regional average. Respondents in the North West and in Yorkshire and the Humber managed on average almost twice as much land as the average farm in the region. In addition, farmers who completed the survey typically owned a greater proportion of the land they managed, compared to the population. The proportion of land used for animal grazing versus cereals and other crops was however broadly representative. The cause of the over-representation of large farms could not be isolated but was likely driven by two factors, individually or in combination. First is a biased sampling frame resulting from posting invitations to those addresses in the electoral rolls that contained the word *farm*. Some small farms may not be advertised as such. Second is non-response bias arising due to the smallest farms lacking experience with ELM schemes. Olsen (2009) compare outcomes from identical choice experiments based on samples recruited online and via post. The authors note that inequities in technology literacy and computer access may give rise to similar problems with unrepresentative samples in online surveys. Olsen (2009) found that observed demographic differences do not translate into significant differences in WTP estimates. Resampling was not attempted due to resource constraints. However, the effects of socio-economic variables are evaluated by using models that account for taste heterogeneity. These are discussed in detail in section 3.6.

Table 3.2: *Sample representativeness by region*

	Sample mean	Population mean (2022)
	North East	
Farm size (ha)	201	146
Land ownership (%)	67	55
Grazing (%)	45	46
Cereals (%)	19	29
	North West	
Farm size (ha)	160	77
Land ownership (%)	74	59
Grazing (%)	43	62
Cereals (%)	8	8
	Yorkshire and the Humber	
Farm size (ha)	171	93
Land ownership (%)	68	63
Grazing (%)	16	35
Cereals (%)	37	32
	South East	
Farm size (ha)	145	87
Land ownership (%)	82	73
Grazing (%)	30	30
Cereals (%)	25	31

3.5 Discrete Choice Experiments

The review on factors affecting ELM uptake by Mamine, Minviel, et al. (2020) focuses on studies using hypothetical discrete choice experiments (DCEs). This is a survey method in which respondents are asked to choose their preferred ELM scheme from a number of options, sometimes including a status-quo alternative (Johnston et al., 2017). Each scheme is associated with a set of characteristics, or *attributes*, that differentiate it from the other options. Inference about farmers' perceived costs can be drawn from observing their choices.

Participants completed three hypothetical DCEs. The DCEs were designed using HTML and CSS code and administered by PC or tablet directly following the socioeconomic and demographic questions. The first DCE involved eight randomly ordered choice tasks presenting the farmers with a hypothetical choice between simple action-based payments. The aim of this DCE was to measure the marginal cost of retiring land parcels of various sizes and qualities to create natural features. The projects were framed to respondents as contributing to natural flood management (NFM) by reducing surface run-off. The second DCE assumed a minimum required amount of natural features for each farm and opens up trading in contracts between farmers. The purpose of the second DCE was to evaluate barriers to market making. First, by exploring farmers' willingness to engage on either side of the trade in ELM contracts for cash. Second, by estimating the impact of transaction costs as a barrier to trading. The third DCE introduced a voluntary bonus payment contingent on collaborating with neighbour(s) to strategically connect natural features across farm boundaries.

To ensure that farmers found the hypothetical schemes to be understandable and realistic (Johnston et al., 2017), they were presented so as to resemble existing schemes offered by the UK Department for Environment, Forestry and Agriculture (Defra). The interventions included are a) planted broadleaf trees, currently offered under the UK Countryside Stewardship scheme for £350 per hectare, and b) natural regeneration, offered under the Countryside Stewardship scheme as arable reversion to grassland for £326 per hectare (Defra, 2022).

Table 3.3: *Countryside Stewardship Scheme capital payments*

Intervention	Requirements	Payment
Natural Regeneration	• Eligible land cultivated for at least 2 years • Sow wild grasses and flowers • Habitat options linked where possible • Free advice and training from Catchment Sensitive Farming	£326/ha
Planted Broadleaf	• Keep planted trees free from competing vegetation • Maintain fences, tree shelters or spiral guards • Replace any dead trees • Free advice and training from Catchment Sensitive Farming	£350/ha

Notes: Payment criteria are subject to changes, please see Defra (2022) for updates

3.5.1 DCE I: Individual Payment

DCE I was made up of eight randomly ordered choice tasks. Table 3.4 shows the attributes and levels in the first choice experiment. The NFM features in the hypothetical scheme were allowed to vary between two types: First, by increasing surface roughness via natural regeneration. Second, by planted broadleaf trees. Planting and maintenance of trees is more expensive than natural regeneration which largely involves retiring farmland from production to rewild.

The *type* attribute therefore serves as a proxy for the cost of natural feature creation. The effect on utility from switching from natural regeneration to planted trees is therefore expected to be negative. These types of features were chosen to mirror previously cited examples of natural flood management (Forbes et al., 2015) which is the topic of chapter 4. At the same time, they were designed to closely resemble real ELM schemes that respondents would be familiar with. This reduces hypothetical bias (Johnston et al., 2017). The *location* attribute states where on the farm the NFM features would be created, and could vary between three location categories, each implying a different management cost. These were locations mid-field, on the field border, and a river edge. The *land quality* attribute defines the quality of land to be set aside for NFM features and varied between rough grazing (low quality) and prime grazing or high-yield crops (high quality).

These attributes represents variation in the opportunity cost of taking this land out of production in favour of NFM, and correlated with the factor productivity of agricultural land in the model. It predicts that compared to an alternative with low quality land, shifting to an alternative citing high quality land will result in a decline in utility. The *area* attribute denotes how much land is to be set aside for NFM.

Table 3.4: *DCE I: Attributes and levels*

Attribute	Levels
Type: *The type of natural feature*	Natural Regeneration, Planted Broadleaf Trees
Location: *Where the feature is placed on the farm*	1) Mid-field, 2) Field boundary, 3) River edge
Land quality: *Suitability of land for agriculture*	1) Rough grazing, wet, steep, rocky etc., 2) Prime grazing land or high yielding crops
Area: *Amount of land set aside*	1) 1/20 hectare (500m^2), 2) 1/10 hectare (1000m^2)
Payment: *Annual payment*	£200, £300, £400, £500

3.5.2 DCE II: Trade in payment-for-NFM contracts

Prior to the second choice experiment respondents were given an information brief which asked them to keep in mind the ELM schemes presented in DCE I. This setting was now changed in two ways: First, respondents were asked to assume that the policy requires that farms enrol a minimum share of their land into ELM projects. Farmers are compensated for these enforced projects per the same mechanism as in the prior, voluntary scheme. Secondly, these government ELM contracts are now tradable between farmers. The land on each farm is given a score based on its potential to generate significant runoff during heavy rains. A higher score means that natural flood management (NFM) is more impactful. Trading ratios based on the relative scores between farms allow farmers of high NFM priority land to take over the NFM requirements of a lower priority farm while receiving a multiple of the contractual payment. This multiple is proportional to the trading ratio. Similarly, low priority farms can benefit from buying out of the NFM requirement for a proportionally lower payment given its trading ratio. The information brief also included a visual guide to the tradable contracts shown here in figure 3.7.

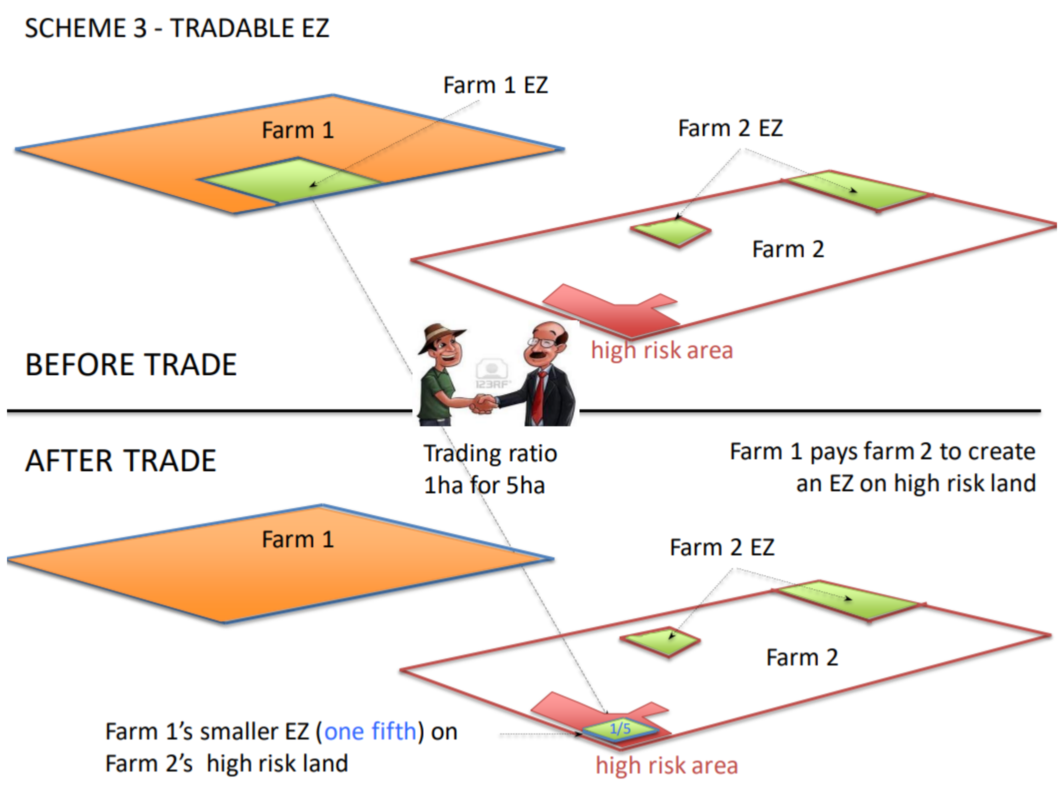

Figure 3.7: *Stylised illustration of how farmers of high- and low risk land can benefit from trading in natural features [referred to here as environmental zones (EZ)]*

Attributes and levels for the second choice experiment are shown in table 3.5. The trading ratio attribute is discrete and has three levels. In the set of six choice tasks, the trading ratios are greater than 1: 5 (the respondent can set aside one fifth of the stipulated area taken over from the low risk farm for the full payment), 10 (one tenth of the stipulated area) and 20 (one twentieth of the stipulated area). The transaction fee has two levels and varies between 5% and 10% of the total payment and is paid by the respondent. The base payment attribute works in the same way as in DCE I. However, the choice cards also show respondents the per-hectare payment they can receive given the trading ratio, which is the base payment multiplied by the ratio. I guide the reader through an example below:

Willingness-to-accept (WTA) example: A respondent (called farmer A) is asked to imagine a hypothetical scenario where they are required to create NFM feature on in total one (1) hectare of agricultural land. In return, they receive a £2,000 per year payment from the government, intended to compensate for lost agricultural output, NFM creation, fencing, maintenance, etc. In this example, A is farming land classed as high risk due to the runoff generation potential at the site. A may take over the 1 hectare NFM obligation of another farmer, B via the trading market. However, because A's land is ten times as suited to NFM projects compared to B's land, *A may take over B's obligation at a trading ratio of 10.* This means that A will receive in total £4,000 per year in exchange for creating NFM features on 1.1 hectares. This area results from adding one tenth of B's 1 hectare obligation to A's original 1 hectare obligation. Due to the runoff generation potential which is ten times higher on A's land, the risk reduction after trading is equivalent to A and B creating one hectare of NFM each. To arrange the trade, A is also required to pay a percentage fee on the value of the trade. In this case, a 5% transaction fee adds a one-time £100 cost.

This WTA (payment in exchange for additional NFM obligations) scenario is followed by a 'willingness-to-pay' (WTP) scenario. Here, respondents were asked to put themselves in the position of a farmer buying themselves out of their NFM obligation. In practice, this involves relinquishing in full or in part their government payment for NFM. In these choice tasks, available trading ratios are set to less than 1: $1/5$, $1/10$, and $1/20$. This means that their trading counterparty needs to set aside proportionally less land when assuming their NFM obligation. It was explained to respondents that a smaller trading ratio would incentivise other farmers to assume their NFM obligation.

Willingness-to-pay (WTP) example: Respondents are asked to imagine a hypothetical scenario where they (farmer A) are required to create in total one hectare of NFM in exchange for £2,000 per year as described in the previous example. A has the choice to buy out of their NFM obligations by transferring the government payment to another farmer, B, who manages land more suited to NFM. In the case of a $1/10$ trading ratio, A receives no money from the government but does no longer have to create any NFM. B has to create additional NFM proportional to the trading ratio, i.e. one tenth of the nominal one hectare amount. To arrange the trade, A is also required to pay a percentage fee on the value of the trade. In this case, a 5% transaction fee adds a one-time £100 cost.

Table 3.5: *DCE II: Attributes and levels*

Attribute	Levels
Trading ratio (WTA, r): *The factor by which respondents can increase their per-hectare payment for NFM by trading*	5, 10, 20
Trading ratio (WTP, $1/r$): *The ratio by which the respondent can reduce their expected per-hectare cost to get out of their NFM obligations by trading*	$1/5$, $1/10$, $1/20$
Transaction fee: *A percentage of the base payment borne by the respondent*	5%, 10%
Payment: *Annual payment, received in the WTA setting and paid in the WTP setting*	£200, £300, £400, £500

3.5.3 DCE III: Voluntary coordination bonus

Respondents were presented with a short information brief describing the scheme aimed at creating natural habitats, explaining how improving connectivity can provide ecosystem services. The projects on offer remained retiring land for natural regeneration and planting broadleaved flowering trees. Participants in this hypothetical scheme received an annual payment for every 100 meters of ecological corridors placed to connect natural features. The annual payments on offer ranged between £200 to £500 per 100 meters of corridor created. Further, participants received an additional one-off bonus payment for coordinating with a neighbouring farm to connect features across their combined land. The bonus scales linearly with the number of neighbours and is allocated equally between them. The one-time coordination bonus ranges between £100 and £400 per neighbour the respondent connects natural features with. If the respondent does not coordinate with anyone, the bonus payment is always zero. If they coordinate with at least one neighbour, the payment to each coordinating farmer is multiplied by their total number (including the respondent). This is done to compensate participants for the added coordination costs of connecting features. The hypothetical contracts specify a minimum required width for the corridors of either 10 or 20 meters. Attributes and levels are summarised in table 3.6.

3.6 Choice modelling

The theoretical foundation for DCEs is hedonic consumer theory (Lancaster, 1966), in which goods or services can be broken down into attributes, each contributing differently to an individual's utility from consuming that good or service. The respondent's choices are assumed to be determined by their trade-offs between the

Table 3.6: *Discrete choice attributes and levels*

Attribute	Levels
Type: *The corridor feature*	Natural Regeneration, Planted Broadleaf Trees
Width: *The required width of corridors*	10 meters, 20 meters
Coordination: *The number of connected farms*	None, One, Two
Bonus: *One-time bonus payment per connected farm*	£100, £200, £300, £400
Payment: *Annual payment per 100m of corridor*	£200, £300, £400, £500

attributes, and the respondent is expected to choose the alternative that maximises their net utility. By modelling a farmers' utility as a function of e.g. payment, location and contract duration, researchers can understand the contribution of each attribute to the likelihood of uptake. Specifically, the ability to estimate the value of attributes at the margin and the possibility of testing for internal consistency[1] (Hanley et al., 1998; Holmes & Adamowicz, 2003) are presented as key advantages of DCEs.

3.6.1 Random utility foundations: Multinomial logit

Sampled farmers ($q = 1, ..., Q$) can choose between J discrete alternatives. Each choice is characterised by a set of attributes ($k = 1, ..., K$) that are assumed to influence respondents' utility. In this case, farmers were asked to choose from among two hypothetical schemes and one opt-out alternative. Alternatives in a choice task are distinguished by the levels of their respective attributes (Welling et al., 2022). The indirect utility farmer q derives from the scheme in alternative j in choice task t is expressed as follows:

[1]For example, DCEs can be designed to test that respondents display consistent preferences across multiple choice tasks.

$$U_{qjt} = \boldsymbol{\beta}' \mathbf{x}_{qjt} + \epsilon_{qjt} \tag{3.1}$$

where $\boldsymbol{x_{qjt}}$ denotes the attributes of NFM scheme in this case, and $\boldsymbol{\beta}$ is a vector of parameters associated with attributes representing the respondents' taste variation. The error term of the utility function ε follows an independently and identically distributed type I extreme value distribution (McFadden, 1974; Scarpa et al., 2008). In theory, β describes how an attribute k contributes to the farmer's utility and its sign tells us whether an increase in a continuous attribute or a shift from a categorical baseline increases or decreases utility for farmer q. Under i.i.d. assumptions, the closed-form expression for the probability that farmer q chooses alternative i in choice task t is given by:

$$P_{qit} = \frac{\exp\left(\boldsymbol{\beta}' \mathbf{x}_{qjt}\right)}{\sum_{j \in C_q} \exp\left(\boldsymbol{\beta}' \mathbf{x}_{qjt}\right)} \tag{3.2}$$

where C_q denotes the choice set available to individual q. The above specification is known as the multinomial logit (MNL) and has important limitations. First, the independence of irrelevant alternatives property implies that the relative odds of choosing between two alternatives are unaffected by the presence or attributes of other alternatives (Hausman & McFadden, 1984). The inclusion of an opt-out alternative illustrates the challenge: If in a binary choice, farmers have equal probability of choosing an NFM scheme (50%) and to opt-out (50%), IIA implies that the odds ($0.5/0.5 = 1$) must be maintained if a second NFM scheme is offered. Assume that the additional scheme is so similar to the original that they are equally likely to be chosen. In such a case, the only way to maintain the original odds of opting for scheme A versus scheme B would be if A is chosen with a probability $1/3$, B with probability $1/3$, and the opt-out alternative with probability $1/3$ (McFadden, 1974).

A consequence is that the regulator could theoretically increase NFM uptake to 100% simply by offering more marginally differentiated schemes. In reality, studies have shown that some farmers would habitually opt out, if given the choice (Hurley et al., 2022). Secondly, the MNL assumes that all individuals share the same parameter vector $\boldsymbol{\beta}$, which may be unrealistic when preferences vary systematically across the population. For example, Hurley et al., 2022; Kenyon, 2007 characterise some farmers as hard-to-reach or low-trust.

3.6.2 Directional hypothesis testing by farmer segments: Latent classes

To address unobserved preference heterogeneity, the latent class model (LC) provides a flexible extension of the MNL framework. The key idea is that the population is segmented into a finite number of classes (or segments), each characterised by its own parameter vector. Individuals are not directly observed to belong to a particular class; instead, class membership is treated as a latent (unobserved) variable. Formally, suppose there are S latent classes. For an individual q in class s, the choice probability of selecting alternative i in task t follows an MNL structure:

$$P_{qit|s} = \frac{\exp\left(\boldsymbol{\beta}_s' \mathbf{x}_{qit}\right)}{\sum_{j \in C_q} \exp\left(\boldsymbol{\beta}_s' \mathbf{x}_{qjt}\right)} \tag{3.3}$$

Instead of assuming a uniform preference structure across all respondents, LCMs allow for variation in preferences by segmenting farmers into different latent classes based on their responses (Greene & Hensher, 2003). This approach helps to identify and understand the different farmer segments that may exhibit diverse decision-making patterns, which can be helpful for designing targeted policy interventions (Tyllianakis et al., 2023). The probability $\pi_{q,s}$ of individual q belonging to class s is defined as:

$$\pi_{q,s} = e^{\delta_s + g(\gamma_s, z_s)} \Big/ \sum_{l \neq s}^{S} e^{\delta_l + g(\gamma_l, z_l)} \tag{3.4}$$

where the class allocation parameters δ and γ for one class are set to zero (Greene & Hensher, 2003; Hess & Palma, 2019). The latent class estimations allow parameter estimates to vary among the (latent) classes, thus accounting for heterogeneous preferences among respondents. Following earlier research applying DCEs to evaluate ELM schemes in the UK (Garrod et al., 2012; Ruto & Garrod, 2009; Tyllianakis et al., 2023), latent class models were estimated to test hypotheses in chapters 4 and 5 that suppose an inequality. The models were therefore estimated in preference space, as specified above. The taste parameter for attribute k, β_k was interpreted as the shift in probability of alternative j associated with a shift in k. Posterior class probabilities were recovered for each respondent that indicate their likelihood of belonging to each class.

3.6.3 Willingness-to-accept distributions: Mixed logit

A key objective of the chapter was to obtain cost estimates for the hypothetical NFM schemes. Such estimates may guide policymakers towards schemes that deliver environmental outcomes cost-effectively. To make realistic predictions about farmer uptake for a given hypothetical scheme, it is helpful to show the estimated cost as a distribution over the sample. Distributions with a long upper tail caution that a portion of farmers will not be reached with any realistic payment. The mixed logit generalises the MNL by allowing random taste variation across individuals and relaxing the IIA assumption. Utility is specified as:

$$U_{qjt} = \boldsymbol{\beta}'_q x_{qjt} + \epsilon_{qjt} \tag{3.5}$$

where $\boldsymbol{\beta}_s$ is the vector of individual-specific taste parameters for respondent q.

Since β varies across individuals, the mixed logit probability integrates over the distribution of β, θ:

$$P_{qit} = \int \frac{\exp(\beta' x_{qit})}{\exp(\beta' x_{qjt})} f(\beta|\theta) d\beta \tag{3.6}$$

If the compensation farmers receive for participating in a scheme is π, the taste parameter associated with the payment attribute is β_π. Respondent q's willingness-to-accept a shift in attribute k can be expressed as the ratio between β_k and β_π (Scarpa et al., 2008; Train & Weeks, 2005; Welling et al., 2022). The WTA estimate is assigned the opposite sign of the taste parameter β_k because a greater payment is needed for farmers to tolerate greater dis-utility from shifts in k[2].

$$WTA = -\frac{\beta_k}{\beta \pi} \tag{3.7}$$

For these reasons two models were used for each of the three DCEs. First, a latent class model was estimated. These estimates were used to identify discrete preference heterogeneity within the farmer sample. Taste parameters estimated the latent class models were also used to test hypotheses that posited inequalities. Second, a mixed logit model with individual-specific taste parameters was estimated in willingness-to-pay space. This was done to identify the distribution of monetary values associated with each scheme. These values were used conduct cost-effectiveness analyses of the hypothetical schemes in terms of flood risk reduction (4) and pollination services (5).

[2]The subscript k differentiates the taste parameters in a mixed multinomial logit model β_k from the Cobb-Douglas output elasticity for agricultural land β.

3.7 DCE design and power analysis

Once DCE attributes and levels have been decided, the choice tasks were finalised. This process is known as the DCE design and involves the pairing of alternatives into choice tasks in such a way that the information they reveal about respondents' preferences is maximised (Rose et al., 2008). To intuitively see the importance of DCE design, imagine a choice between two ELM schemes that are identical but for the payment attribute. In this choice, the respondent does not need to consider any other attribute but the base payment, which clearly has a positive parameter. In other words, a rational respondent will always choose the higher paying alternative in such a choice task. Observing this choice adds nothing to the information about the value of other attributes, such as the area of land to retire or the land quality. The number of choices that can be observed is limited by the ability to recruit DCE participants and the number of choice tasks respondents can typically complete before experiencing fatigue, at which point some respondents may exhibit inconsistent preferences (Campbell et al., 2015). It is therefore important to derive the most information from the available observations.

The metric to compare the information yielded from different designs is called efficiency. An experimental design is more efficient than another design if it produces data that enables estimation of parameters with lower standard errors. The information yield of a given design can be estimated given assumptions about respondents' tastes.

Recall from equation (3.2) that the likelihood of choosing a given alternative depends on the ratio of the utility derived from that alternative over the utilities derived from remaining alternatives. Taking the second derivate of the log-likelihood with respect to the taste parameters multiplied by the number of respondents Q produces the Fisher information matrix (Rose et al., 2008):

$$I(\beta) = Q \times \frac{\partial^2 L(\beta)}{\partial\beta\partial\beta'} \tag{3.8}$$

The Fisher matrix is also known as the curvature matrix because its values are largest at the point of maximum curvature, or the peak, of the log-likelihood function. This is the vector of attribute-specific taste parameters β_k where the likelihood of observing the choice in the data is optimised. The Fisher matrix is of further econometric importance as its inverse is the asymptotic variance covariance (AVC) matrix, including the scaling of $1/Q$. This means that the impact of sample size Q on the design can readily be investigated. The asymptotic standard errors are the roots of the diagonal of the AVC matrix, therefore these standard errors decrease with a rate of $1/\sqrt{Q}$ of the sample size.

The efficiencies of different designs were compared by plotting the standard errors as a function of the sample size. I chose the design that achieved the smallest standard errors at a given sample size. Two so-called D-efficient designs were compared, where the determinant of the AVC matrix was minimised by drawing taste parameters β_k from a distribution. One was a naive Bayesian D-efficient design where the parameters are drawn from normal distributions centered around a prior of zero. It is called naive because setting priors to zero means that no assumptions are made about the taste parameters. The second was a uniform D-efficient design where priors were drawn from uniform distributions that are either positive or negative given my assumptions about the directionality of the taste parameters (ChoiceMetrics, 2012). Chosen cutoffs for the assumed uniform distributions of each parameter across DCE I and DCE II are displayed in table 3.7, based on economic intuition as well as previous research featuring similar attributes (Tyllianakis et al., 2023). The economic rationale for the limits is discussed further in section 4.3 of chapter 4. The assumed distribution limits for DCE III are similarly displayed in table 3.8. The economic rationale for these distributions is discussed further in section 5.3 of chapter 5. All choice cards were designed using the `Ngene` software (ChoiceMetrics, 2012).

Table 3.7: *DCE I & II: Uniform distributions for taste parameters*

Taste parameter	Motivation	Distribution Limiyd
Natural regeneration	Preferred to planted trees due to shorter time horizons and lower costs	$0.01 - 0.5$
Field edge placement	Preferred to in-field placement due to less disruption to production	$0.01 - 0.5$
River edge placement	Preferred to in-field placement due to less disruption to production	$0.51 - 0.1$
Low-quality land	Preferred to high-quality land due to lower opportunity cost of creating NFM	$0.01 - 0.5$
$500m^2$ for NFM	Preferred to $1000m^2$ NFM due to lower costs	$0.01 - 0.5$
10:1 trading ratio	Preferred to a 5:1 ratio due to a smaller additional NFM obligation when trading	$0.01 - 0.25$
20:1 trading ratio	Preferred to a 5:1 ratio due to a smaller additional NFM obligation when trading	$0.251 - 0.5$
5% transaction fee	Preferred to a 10% fee due to cost-minimisation	$0.01 - 0.5$
Payment (WTA)	Strictly positive due to cost-minimisation	$0.01 - 0.5$
Payment (WTP)	Strictly negative due to cost-minimisation	-0.5 to -0.01

Table 3.8: *DCE III: Uniform distributions for taste parameters*

Taste parameter	Motivation	Distribution Limits
Natural regeneration	Preferred to planted trees due to shorter time horizons and lower costs	$0.01 - 0.5$
10 meter corridor width	Preferred to a 20 meter width due to less disruption to productive land	$0.01 - 0.5$
No collaboration	Preferred to collaboration with two neighbours due to zero co-ordination costs	$0.251 - 0.5$
Collaboration ($n = 1$)	Preferred to collaboration with two neighbours due to lower co-ordination costs	$0.01 - 0.25$
Coordination bonus	Strictly positive due to cost-minimisation	$0.01 - 0.5$
Payment per 100 meters	Strictly positive due to cost-minimisation	$0.01 - 0.5$

The efficiency of the DCE designs is directly linked to the minimum sample size required to produce statistically reliable results. The more information that can be recovered about respondents' tastes using a particular design, the fewer choices need to be observed to achieve narrow standard errors. de Bekker-Grob et al. (2015) have proposed a commonly used rule of thumb for estimating the required sample size for accurate taste parameter estimates. This rule of thumb is shown in equation (3.9) below:

$$N > \left((z_{1-B} + z_{1-A}) \cdot \sqrt{\frac{\Omega}{\hat{\beta}}} \right)^2 \tag{3.9}$$

where N is the required sample size for the DCE. z_{1-B} is the z-score corresponding to the statistical power $(1-B)$, which reflects the probability of correctly reject the null hypothesis, while z_{1-A} corresponds to the significance level, the probability of a false positive. Ω is the AVC matrix and $\hat{\beta}$ is the priors for the taste parameters. A comparison of designs for DCE I is shown in figure 3.8 and another example for DCE III is displayed in figure 3.9. For every taste parameter and across designs, observe how the required sample size increases as higher cut-offs for statistical significance, α, are enforced. The dashed lines intersect the required sample sizes at the 5% and 1% significance levels, respectively. There was negligible difference between the naive and the uniform parameter samples, the standard errors diminish sharply after about 100 observations. I proceeded with the uniform design, as it requires no assumption about the standard deviation of the distributions and there is prior evidence about the expected signs of attributes from previous studies. These results indicate that the sample of 494 is very likely to yield informative estimates for each DCE.

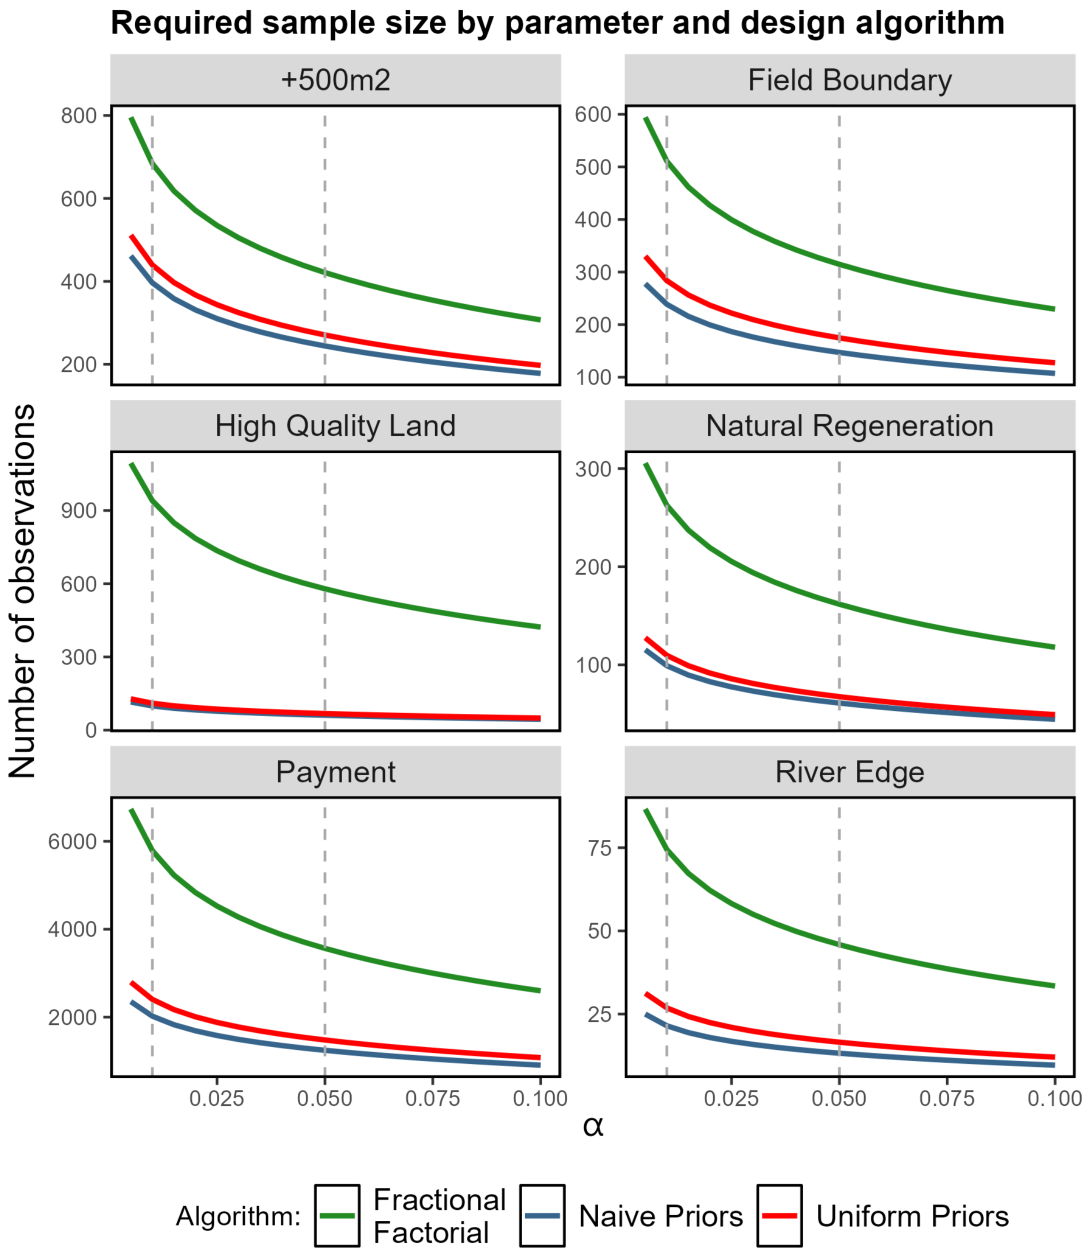

Figure 3.8: *The increase in required sample size as we enforce a lower probability of incorrectly rejecting the null hypothesis, illustrated across three different designs, including a) randomly sampled choice tasks from a factorial design, b) parameters drawn from a normal distribution all with naive means of zero and c) drawn from uniform distributions of signs motivated by theory. In each case, the number of choice tasks is eight.*

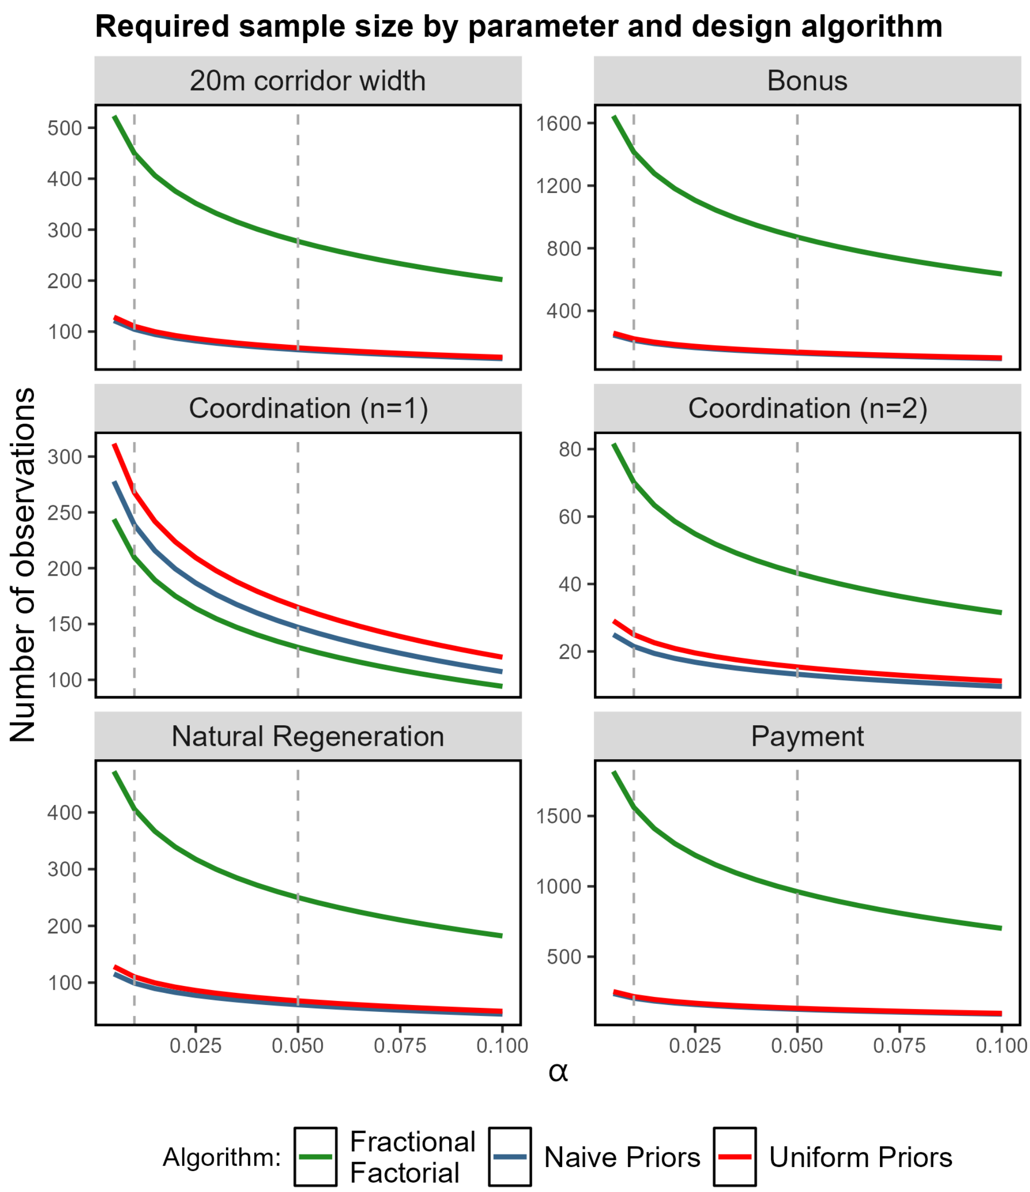

Figure 3.9: *The increase in required sample size as we enforce a lower probability of incorrectly rejecting the null hypothesis, illustrated across three different designs, including a) randomly sampled choice tasks from a factorial design, b) parameters drawn from a normal distribution all with naive means of zero and c) drawn from uniform distributions of signs motivated by theory. In each case, the number of choice tasks is eight.*

3.8 Identifying serial non-participants

Serial non-participants are those who always choose the opt-out alternative, and can be divided into two types; protesters and very high compensation requirements. Protesters are respondents who, for whatever reason, disagree with the idea of ELM schemes or environmental protection or with the hypothetical exercise of a discrete choice experiment. The latter are respondents who demand compensation higher than ever offered within the choice experiment. Under rationality assumptions, all farmers would be expected to participate in the scheme if the compensation is sufficiently high. However, not all land managers may be entirely driven by profit maximisation and non-profit-based motives (which can reflect self-interest or not) can have an important impact on a land manager's decision making. Protesters should be considered to be out of the market and should therefore be omitted from the analysis used to derive WTA estimates (Villanueva et al., 2017). Table 3.9 shows how the serial non-respondents differ from those who choose an ELM contract in at least one choice task across DCEs I through III.

Table 3.9: *Predictor averages by choice type*

	Contract A	Contract B	Opt-Out	Serial Opt-out
Age	54.7	54.6	55.8	57.7
Farm Size (ha)	176.6	174.4	173.7	169.3
% Female	25.0	25.3	26.3	22.5
% Poll. Dependence	43.5	41.6	32.6	35.0
% Primary Income	75.3	76.9	81.9	75.0
% ELM Uptake	60.2	58.9	56.2	37.5
Response Time (s)	18.7	18.9	20.6	9.2

The strongest predictors of serial non-participation are current ELM participation (37.5% of serial non-participants versus 56-60% for all other choices) and response times, i.e. for how long the respondent takes to make their choice. Serial non-participants move on considerably quicker, after only 9.2 seconds on average compared with ca 20 seconds for others. In total there are 34 serial non-responders, or 7.9% of the full sample. Serial non-participants who are currently enrolled in a real ELM scheme should likely not be regarded as protesters when it comes to ELM schemes as such, as current Defra schemes are voluntary (Defra, 2022). However, they may still protest the concept of making hypothetical choices itself. This is supported by the considerably shorter response times, indicating that these respondents do not consider each option as carefully as other respondents. Respondents who chose the opt-out alternative 100% of the time were excluded from the analysis. Latent class analysis was used to identify the respondents with merely a high propensity of opting out (Burton & Rigby, 2009).

3.9 Limitations

Although DCEs are widely used in research on ELM schemes due to a dearth of data on farmers' revealed preferences from past schemes (Mamine, Minviel, et al., 2020), evidence from DCEs should be used cautiously when implementing policy in different environments/populations (Hanley et al., 1998). Much of the concern revolves around the hypothetical nature of the choice, where respondents are not committing to participate in a real scheme (Johnston et al., 2017). The deviation in welfare estimates due to the hypothetical nature of the experiment is known as *hypothetical bias* (Haghani et al., 2021). Although hypothetical bias is an inherent limitation of all stated preference methods, a number of steps were taken to minimise its impact.

One source of hypothetical bias is unfamiliarity. Good DCE practice demands that choice attributes enter into the respondents' utility function (Carson & Groves, 2007). This can be achieved by making sure that choices are easily understood and that choice attributes are relevant, based on evidence from previous studies, current schemes, or a pilot survey (Johnston et al., 2017). The hypothetical scheme presented in DCE I was modelled after payments currently available to farmer under the Countryside Stewardship scheme (Defra, 2022). The majority of farmers participating in the DCE were enrolled in similar ELM schemes during the surveying period. Similar efforts were taken to ensure respondents' familiarity with the scheme in DCE III. Although past decades have seen only limited support for collaboration within UK ELM schemes (Jones et al., 2023), ideas are increasingly disseminated from government to the farming community. An example is the Natural England Facilitation Fund which finances farmer clusters working towards environmental goals with a facilitator (Dewally et al., 2025).

DCE II was judged to be particularly difficult for respondents to understand. This was because a) current UK ELM schemes do not support trade in contracts, b) trading ratios require a degree of numeracy to interpret properly (fractions), and c) the DCE involved scenarios where the respondent was on either side of the hypothetical trade. Efforts were taken to deal with these limitations in two ways: At the survey design stage, illustrations were added to the information brief preceding the DCE to add a visual explainer. Additional strategies to identify respondents who misunderstood the options were deployed ex-post in the estimation stage. All three DCEs were analysed using a latent class model which can identify heterogeneous preferences across classes of respondents. Overwhelming choice complexity may cause farmers to optimise only for the payment or default to the opt-out alternative (Adamowicz et al., 2014; Zhang & Adamowicz, 2011). Latent class estimates were used to separate these effects. Responses to strictly dominant alternatives (those where every attribute is more attractive than others in the same choice task) were analysed to identify respondents who made ”irrational” choices.

Another source of hypothetical bias is consequentiality (Vossler et al., 2012). Stated preferences are less informative if the respondent a) does not believe that the ELM schemes would bring environmental benefits (output consequentiality (Czajkowski et al., 2021)), or b) does not believe that their responses would have any influence over actual ELM schemes (survey consequentiality (Liu & Tian, 2021)). The evidence on the impact of consequentiality for WTA estimates to provide public goods is inconclusive, with study settings most similar to this one (farmers’ WTA to produce environmental goods) finding an insignificant effect in one case (Granado-Diaz et al., 2024) and significantly biased WTA estimates (38% higher) in another (Villanueva et al., 2025). Recent work has suggested that full-time farmers who own their land are more likely to perceive the survey as consequential and the authors attribute this finding to familiarity with ELM schemes (Villanueva et al., 2025). These characteristics are controlled for in the econometric modelling.

Chapter 4

Analysis of a hypothetical water runoff permit market with spatial targeting

4.1 Introduction

Several success stories can be found among real-world experiments with tradable emissions permits, for example in terms of air quality (Shapiro & Walker, 2018) and public health (Chay & Greenstone, 2003b), and includes schemes such as the US federal Acid Rain Program and the EU Emissions Trading Scheme. However, a large literature is devoted to the numerous ways a cap-and-trade system can theoretically fail to achieve optimal effectiveness, including transaction costs (Stavins, 1995) and heterogeneous damages (Fowlie & Muller, 2019; Montgomery, 1972). Where there are no transaction costs and the marginal damage per unit of pollutant is uniform across sources, Xepapadeas et al. (1997) shows that the social welfare-maximising regulator allocates initial allowances such that the market price for permits equals the marginal damage from pollution.

However, as early as the 1970s, Montgomery (1972) demonstrated that when the marginal damages from pollution differ between sources, uniform (or one-for-one) trading will not achieve the social optimum. Uniform trading means that all polluting firms face the same market price for permits. Differences in geography, demographics, and vulnerable ecosystems may cause marginal damages to differ across sources (Fowlie et al., 2012). In addition, moral hazard may necessitate stronger abatement incentives in some geographies than in others. In an experiment aimed at evaluating the causal effect of CAIR, a legally contentious US cap-and-trade program for SO_2 emissions, Leppert (2023) found that sources exporting pollutants outside the state where they are regulated responded less to a reduction in the cap. In China, Cai et al. (2016) similarly find that provincial governments enforce river pollution reduction mandates less forcefully in counties directly upstream of the provincial border, as water-borne pollution damages are transported to downstream provinces.

This research studies another externality where geographic differences are pronounced and important for policy design.Agricultural land use has been shown to affect flooding. Farmland runoff and subsurface drainage may act as pathways, causing flooding in downstream receptor areas (Posthumus et al., 2008). Emphasis has therefore been placed on natural flood management (NFM) as an adaptation method, defined as '...the alteration, restoration or use of landscape features to reduce flood risk'. NFM is a potential benefit from environmental land management (ELM) schemes, where the government pays farmers to manage their land in specific ways. ELM schemes providing NFM involves an economic cost to farmers who may no longer use certain land for crops or grazing, which increases with the agricultural value of retired land.

Qualitative work carried out in Scotland by Holstead et al. (2017) suggests that appropriate long-term financial incentives are needed to increase uptake of ELM schemes. Incentives must be administratively simple and be joined up with other farm payments. Trading ratios have been proposed as a policy approach to geographically heterogeneous damages and incentives (Holland & Yates, 2015). In such a scheme an exchange rate is applied to the permit market such that the price faced by a source reflects its relative marginal contribution to the externality.

Theoretical findings in Fowlie and Muller (2019) comparing trading ratios to undifferentiated cap-and-trade show an average welfare gain from differentiation, although there is a welfare loss when marginal abatement costs are underestimated. In a cap-and-trade scheme, demand for permits will be higher among firms facing comparatively high marginal abatement costs. With undifferentiated trading, these firms will pollute more at every permit price the market decides. If these sources are also generating higher damages, spatial targeting may produce very high costs. Trading ratios are therefore suitable in settings where marginal damages and abatement costs are not strongly correlated. Runoff generation from agricultural land use causing flooding and diffused pollution is one such setting where the externality is mainly produced at higher altitudes while the abatement cost is higher at the more productive lower altitudes (Forbes et al., 2015).

Few empirical studies of real-world spatially differentiated permit markets so far exist (Holland & Yates, 2015), making the type of observational quasi-experimental policy evaluation from Leppert (2023) and Fowlie et al. (2012) difficult. Discrete choice experiments (DCE) featuring hypothetical schemes have been widely used to estimate likely costs and benefits when observational data is not available (Hoyos, 2010). I run two DCEs with 494 English farmers. The first experiment elicits preferences for an action-based payment for spatially targeted NFM interventions, intended to reduce flood risk. These were deliberately designed to resemble ELM schemes currently available via the UK Department of Environment, Food and Rural Affairs (Defra). The second experiment features a variation of the first where trading and trading ratios are introduced.

This is the first study of trading ratios applied to a market for NFM provision. A potential barrier to farmer participation in a market for NFM contracts is transaction costs. Compared to ELM schemes currently offered in the UK, tradable contracts would add costs by matching 'buyers' and 'sellers', communicating relative trading ratios, and facilitating transactions. Schmalensee and Stavins (2013) and Schmalensee and Stavins (2017) do not find transaction costs to be a significant barrier in emission permit markets. These results may not translate to a hypothetical market in NFM contracts. British farming is a low-margin sector and transaction costs have been identified as a barrier even in bilateral agreements between farmers and Defra (Peterson et al., 2015). This chapter uses DCEs to isolate transaction costs in a hypothetical market for spatially targeted NFM. The transaction costs associated with trading may be evaluated in context of the perceived fairness compared to a spatially targeted scheme where payments are offered only to farms in NFM priority areas. This is relevant because UK farmers are aware that runoff generation is not driven by practices on individual farms (Holstead et al., 2017). Farmers have shown high endorsement in principle of higher pay for greater effort, rather than external circumstances (Loft et al., 2020) and perceived inequity can threaten participation. Using a hypothetical DCE I present support for cost-savings from trading that are robust to transaction costs up to 10 percent.

The rest of the chapter proceeds as follows: The background section provides a review of the current state of knowledge on NFM and of the relevant case studies from the UK. The theory section presents a model of a spatially differentiated market in payment-for-NFM contracts and explore how a trading ratio applied to the farm should affect the demand for contracts. The following section on methodology describes the econometric specification to test the hypotheses that follow from the model. Next, the reader is introduced to the theory behind `SCIMAP-Flood`, a geophysical model of surface runoff which is used to identify priority areas for NFM. I integrate results from `SCIMAP-Flood` with the choice experiment to demonstrate the benefit from trading. Finally I present the results and discuss their relevance for future DCE studies.

4.2 Background

Floods are among the most economically costly natural hazards in the UK, causing significant damage to property, infrastructure and local livelihoods. For 2020, the Association of British Insurers reports £817 million in flood-related losses for the UK alone (Bates et al., 2023). Flooding is a natural process, but floodplains are also ideal for agriculture and urban development close to water resources and navigation. Consequently, development in floodplains has increased the exposure of people, property and infrastructure to floods. In many cases it is not practical, cost effective or politically feasible to relocate communities, property and economic activities away from areas prone to flooding, so measures are put in place to manage flood risk by reducing the probability of inundation and/or the negative consequences when a flood does occur (Posthumus et al., 2008).

4.2.1 Natural Flood Management

Natural flood management (NFM) seeks to restore or enhance catchment processes that have been affected by human intervention. These activities aim to reduce flood hazard, while also sustaining or enhancing other potentially significant co-benefits including enhanced ecosystem services (aquatic, riparian and terrestrial) such as greater biodiversity, improved soil and water quality (Wingfield et al., 2019). Floods can be categorised into different types, including fluvial (caused by an overflowing river), pluvial (caused by extreme rainfall independent from a body of water) or coastal. These are often analysed in isolation, where in reality, they may act in combination. This, along with the complexity of flood risk modelling and relative infrequency of significant flood events, has contributed to a lack of data and conclusive evidence on the efficacy of various natural flood management schemes (Dadson et al., 2017).

Dadson et al. (2017) review and summarise the evidence to date on NFM in the UK. They focus on projects in river catchment meant to reduce fluvial flooding. At spatial scales less than 20 km^2 they find evidence of an effect from land use on flood flows, including a reduction from upland forestry compared to a grassland baseline. Both arable and livestock agriculture have been shown to increase surface runoff at local scales. Two experiments with tree-planted plots reduced runoff by 48% and 78% respectively compared to grazed controls, although there was a high degree of variability between sites. Wingfield et al. (2019) echo the conclusion of Dadson et al. (2017) that catchment-scale evidence on NFM effectiveness is limited.

The high-level evidence base presented to policymakers by subject experts emphasise the variability of these types of NFM projects in terms of cost and effectiveness (Wilkinson et al., 2019). As shown in Table 4.1, NFM measures vary in land use requirements, engineering requirements, and cost. It is also important to choose the appropriate measure for the type of land in question (Forbes et al., 2015). The suitability of a site for the implementation of natural flood risk mitigation measures is determined by the travel time of the flood waters to the point of impact, the spatial pattern of rainfall depth, the effectiveness of the land cover in generating rapid flood flows (overland, drains and near surface flows) and the strength the of the hydrological connectivity from the landscape to the river channels (Reaney, 2022). Literature review by Dadson et al. (2017), heavily biased towards tree-planting, also finds that while typical NFM projects show effectiveness during small and moderate floods, flows were not reduced significantly during the worst flood events.

There is considerable geographic variability in the effectiveness of NFM. The UK environment agency has published digital maps (Environment Agency, 2021) to assist the prioritisation of NFM or land management changes with the aim of slowing water flows and reduce the flood risk. They have been specifically created to contribute to the spatial prioritisation of catchments within the pilot Local Nature Recovery and Landscape Recovery land management schemes for NFM interventions (Broadmeadow et al., 2023), but also within the other grant awarding schemes such as the England Woodland Creation scheme and the England Peatland restoration scheme. The NFM priority map of the sampling area in the north of England is shown in figure 4.1. It shows that there is a considerable concentration of high NFM priority areas across the region. In total, $21,000\ \mathrm{km}^2$ are classed as high-risk, $16,500\ \mathrm{km}^2$ as medium risk, and $14,000\ \mathrm{km}^2$ as low risk.

By high risk, Environment Agency (2021) refers to land where NFM projects can have the greatest impact in terms of flood reduction. From here on out, "High-risk" land should therefore not be interpreted as facing a higher likelihood of being flooded. Instead, the term refers to land which has a high potential to generate surface run-off and contribute to flooding in surrounding lowlands. By altering the land use in these areas, the risk of flooding in the river catchment can be reduced (Reaney, 2022).

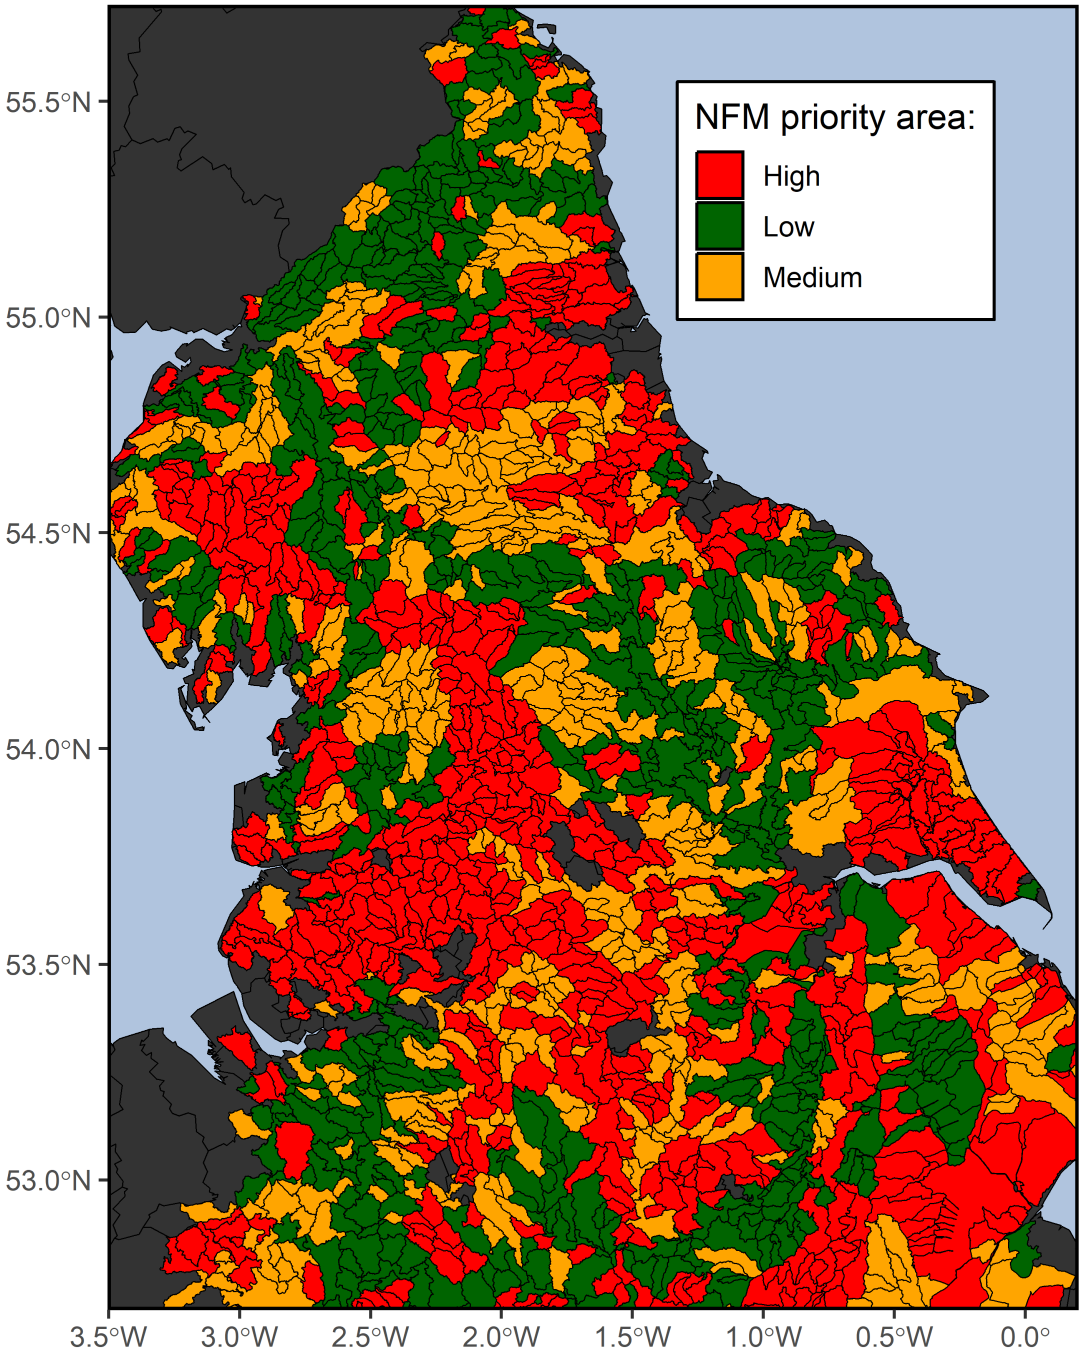

Figure 4.1: *Spatial prioritisation of catchments suitable for using Natural Flood Management in the north of England (Environment Agency, 2021). Dark gray areas signify missing data.*

In summary, recent reviews and technical briefs highlight uncertainty around spatial and flow variability (Dadson et al., 2017; Forbes et al., 2015; Wilkinson et al., 2019) and cost-effectiveness therefore requires that projects happen where surface runoff during heavy rain is most reduced. Holstead et al. (2017) present results from interviews with farming focus groups in Scotland, suggesting that appropriate long-term financial incentives are needed to increase uptake.

Incentives must be administratively simple and be joined up with other farm payments. 64% of respondents surveyed by Holstead et al. (2017) cited lack of information as a barrier to uptake, and 60% oppose it primarily on the grounds of tradition or habit. In addition to more information and simplifying uptake, financial compensation must also be large enough to incentivise deviation from the status-quo.

A focus group study run by Posthumus and Morris (2010) also indicate that UK farmers are unwilling to pay for externalities, with participants noting that NFM projects on their land would reduce flooding elsewhere. While the participant-led multi-criteria evaluation by Kenyon (2007) reveals that the views and opinions of the wider community, particularly those communities that host NFM is increasingly important. The evaluation and design of incentives for uptake are so far largely absent from the economics literature.

4.2.2 Barriers to top-down spatial targeting of NFM

Given the similarities between common NFM actions outlined in table 4.1 and projects currently eligible for payments under the Countryside Stewardship and Landscape Recovery schemes (see section 3.2 for details), it is plausible that the UK government could direct payments to target effective flood management. The NFM priority maps published by the Environment Agency (2021) communicate

Table 4.1: *Types of natural flood risk management*

NFM MEASURE	TECHNICAL INFORMATION	COST
Woodland creation	Aims to reduce local flooding (catchments smaller than 100km^2). 10% increase in conifer or broadleaved forest cover in catchment could achieve a 40mm and 25mm decrease in water yield, respectively (Forbes et al., 2015). This is the action included in this study.	Variable according to site, tree species and management. Woodland creation as part of the Countryside Stewardship scheme has been compensated at £350 per hectare per year (Defra, 2022).
Wetland creation	Small or large scale depending upon the overarching aim. The key is to create the wetland in areas where the flood reduction potential is greatest.	Depends on the extent of engineering required but likely to be moderate, with some low cost maintenance required.
Washlands and storage ponds	Areas next to a river or stream where flood water is directed at times of high flow. May use barriers such as earth bunds to intercept overland flow (together referred to as runoff attenuation features or RAFs). Suitable sites tend to be large floodplains with suitable foundations for supporting any embankments or control structures.	Extremely variable depending on scale. May require pre-work assessments and planning permission for large-scale projects.

the plausibility of *spatially targeted* schemes directed particularly at those farms that manage "high-risk" land. These are farms where the creation of NFM features would have the greatest impact in terms of reducing catchment flooding. However, there are objections to such a policy design. As evidenced in Dadson et al. (2017) and Wilkinson et al. (2019), the potential contribution of a land parcel to downstream flooding is driven as much by geography as by farming practices. A voluntary payment-for-NFM scheme targeting high-risk farms could therefore be perceived exclusionary, while a targeted command-and-control policy could be viewed as unfairly punishing.

The interviews with farmers from Scotland conducted by Holstead et al. (2017) highlight the importance of social dynamics and possible stigma. One respondent expressed worries that neighbours would judge the payment as benefiting unfairly from state benefits. Based on scepticism of Defra programs and of state intervention in certain segments of the farming community (Hall & Pretty, 2008), it is therefore prudent to consider how a hypothetical targeted NFM scheme offered to select high-risk farms would be perceived.

> I know NFM is not money for nothing, but it would be viewed as that by people. Then six months down the line they will say ”oh such and such is getting £10 000 for that [...] He has nothing on it?”
>
> Farmer 10 (Holstead et al., 2017)

Cultural barriers also exist. Farmers interviewed by Holstead et al. (2017) expressed that receiving payments for effectively retiring farmland does not align with perceptions of what it means to be a farmer. Taking land out of production to create NFM may diminish the cultural and professional significance of that land. The interviews reveal a commonly held pride in the idea of working the land and that what one puts in is what one takes out.

> Some people would see that [being involved in an NFM scheme] as a benefit, you wouldn't be doing the same amount of work and you would be getting the same return. I would say that this goes against the grain of 90% of farmers or more
>
> Farmer 13 (Holstead et al., 2017)

Farmers have shown high endorsement of principle of higher pay for greater effort, rather than external circumstances (Loft et al., 2020) and perceived inequity can threaten participation. For these reasons, a top-down government NFM scheme targeted at high-risk farms may be unpopular. Farmers targeted for the scheme may worry about neighbours' perceptions and the impact of retiring significant amounts of land on their professional self-image. This chapter therefore proposes a market for tradable NFM contracts with spatial targeting. Trading allows high-risk farms who are opposed to NFM the option to buy out of their NFM obligation. However, spatial targeting means that high-risk farms benefit more financially from engaging in NFM than does low-risk farms. This model is discussed in more detail in the following section.

4.3 Theoretical background

With the background set, I proceed with developing a model of farmers' demand for enrolling land into ELM schemes. This section does not aim to paint a complete picture of farmers' decision-making, but fills three important functions: First, setting up a stylised theoretical background to the choice experiment allows me to predict some behaviours that, if observed in the data, would add confidence that respondents have understood the survey and acted in a rational way. Second, the model guides the design of the hypothetical ELM schemes by predicting the variables that enter into the farmers' optimisation problem. Third, the model predicts the expected signs for the variables of interest that allow me to formulate and test hypotheses.

4.3.1 A base model of ELM uptake

I begin with a model allowing enrolment into the individual schemes currently available via Defra. I extend the model to a spatially targeted cap-and-trade scheme with trading ratios as defined in Holland and Yates (2015). I consider a farmer with an endowment of productive land $\bar{L}$ that can be used either for agricultural production or be enrolled in an ELM scheme. The area of land used in agriculture is denoted by L_{AG} hectares and the area used for ELM by L_{NF} to indicate natural features. The first constraint in the farmer's choices is therefore that the sum of land area used for agriculture an for ELM actions can not exceed the total land endowment:

$$L_{AG} + L_{NF} \leq \bar{L} \tag{4.1}$$

As is typical in production economics, I assume that the land endowment is fixed in the short run and that entering additional land into an ELM scheme is always a substitution from productive land. This is a simplification, as some land eligible for these schemes may be of very marginal economic value (Defra, 2022). Consider the following Cobb-Douglas production function:

$$Y = X^{\alpha} L^{\beta} \tag{4.2}$$

Agricultural output Y is a result of a two-factor Cobb-Douglas production function of land L_{AG} and other inputs X (Dawson & Lingard, 1982). Assume that returns to scale are constant such that $\alpha + \beta = 1$ and that returns to land are diminishing ($0 < \beta < 1$), because the most productive land being farmed first. Absent any spatially targeted incentives or eligibility requirements, farmers will retire marginally productive land first. That $0 < \alpha < 1$ should be clear by recognising that there is an upper limit to how much seed or grazing cattle one can pair with a unit of land. I deviate from older specifications (e.g. Ulveling and Fletcher (1970)) by excluding labour as a distinct production factor. This is a result of the study design, where I am interested especially in substitution of land between agriculture and ELM projects.

Any environmental benefits from ELM are assumed to be fully externalised. Although research on whether UK farms are profit-maximisers is lacking, its absence in the recent agricultural economics literature may itself be revealing. About U.S. agriculture, Crespi et al. (2012) writes that research on the market power of farms has been replaced by a persistent concern about food processors', handlers', and occasionally retailers' potential market power, as buyers of farm products and the impact such power might have on the future of small farms.

U.S. farm data also show only few violations of cost minimisation (Zereyesus & Featherstone, 2017; Zereyesus et al., 2021) which indicates a competitive market for farm outputs. It can be argued that some of the drivers behind this shift (entry of large, low-cost food retailers, globalisation (Saitone & Sexton, 2010)) apply also to the UK, as well as EU-wide changes such as the shift from price-based CAP subsidies (Velázquez et al., 2017). I continue on the assumption that farmers operate on a competitive market without the ability to set output prices or collude with competitors to do so. Their objective is therefore to minimise costs, subject to meeting the residual demand $\bar{Y}$ they face at market prices:

$$\begin{aligned} \text{minimise:} \quad p_X X + c_{NF} L_{NF} - \pi L_{NF} \quad \text{subject to} \\ X^{\alpha} L_{AG}^{\beta} = \bar{Y},\ L_{AG} + L_{NF} \leq \bar{L} \end{aligned} \tag{4.3}$$

where p_X and c_{NF} are the market prices of inputs and cost of creating the natural features respectively. The incentive to put land into an ELM scheme is a payment π, proportional to the amount of land L_{NF} enrolled. By solving for L_{NF} in the Lagrangian in equation (4.4) we find the demand for enrolling land into the ELM scheme. Equations (4.5) - (4.7) are the first-order conditions.

$$\begin{aligned} \mathcal{L} = p_X X + c_{NF} L_{NF} - \pi L_{NF} - \mu_1 \left(\bar{Y} - X^{\alpha} L_{AG}^{\beta}\right) - \\ \mu_2 \left(\bar{L} - L_{AG} - L_{NF}\right) \end{aligned} \tag{4.4}$$

$$\frac{\partial \mathcal{L}}{\partial X} = p_X + \mu_1 \alpha X^{\alpha - 1} L_{AG}^{\beta} = 0 \tag{4.5}$$

$$\frac{\partial \mathcal{L}}{\partial L_{AG}} = \mu_1 \beta X^{\alpha} L_{AG}^{\beta - 1} + \mu_2 = 0 \tag{4.6}$$

$$\frac{\partial \mathcal{L}}{\partial L_{NF}} = c_{NF} - \pi + \mu_2 = 0 \tag{4.7}$$

By solving for μ_1 and μ_2, substituting into the constraints and simplifying, I find the cost-minimising demand for using land in the ELM scheme:

$$L_{NF}^{*} = \bar{L} - \left(\frac{\frac{\beta}{\alpha} p_X \bar{Y}^{1/\alpha}}{c_{NF} - \pi} \right)^{\frac{\alpha}{\alpha+\beta}} \tag{4.8}$$

Differentiating equation (4.8) with respect to the payment π reveals the marginal increase in the amount of land retired for ELM as the payment increases. It depends on α and β such that the marginal increase is lower when the dependence of production on high-quality land is high. As shown in equation (4.8) the L_{NF} demand function is defined only if $\pi > c_{NF}$, when $\partial L_{NF}/\partial \pi > 0$. If $\pi r_i < c_{NF}$ the farmer takes a loss on every unit of NFM created, while the marginal product of land used for agricultural production, L_{AG}, is diminishing but strictly positive. Cost-minimising behaviour therefore results in no land used for NFM when $\pi r_i < c_{NF}$. That is, when the payment multiplied by farm i's individual trading ratio is lower than the cost of creating natural features. From here, I state the first hypothesis which would, if not rejected, support the validity of the base model:

Hypothesis I: In response to an increase in the per hectare ELM payment, a) farmers will set aside more productive land area towards ELM, and b) comparatively more will be set aside by farms where land productivity is comparatively low.

4.3.2 Trading in ELM contracts and spatially heterogeneous damages

ELM schemes in the UK are currently voluntary and can therefore most accurately be thought of as a subsidy for provision of environmental services by the agricultural sector. However, one can imagine the regulator taking a more proactive stance towards environmental goods like flood- and pollution-management, pollination services, and habitat conservation. An alternative perspective sees agricultural production generating negative externalities that can include erosion of surface roughness, soil quality, and destruction and/or fragmentation of wildlife habitats. Leppert (2023) shows a causal decline in sulphur emissions following the introduction of a cap-and-trade scheme creating a price on the environmental externality.

The theoretical literature on cap-and-trade instruments has overwhelmingly been developed with the power and industrial sectors in mind (Xepapadeas et al., 1997) as are the majority of applications in environmental policy (Chan et al., 2012; Leppert, 2023). This is partly a natural consequence of these industries contributing large shares of economy-wide emissions, but also partly a matter of convenience, as the point-source emission sources lend themselves to regulating emissions directly.

These instruments would seem less suited to regulate agriculture (Spicer et al., 2021). However, advances in modelling spatial data at high resolutions, for example runoff generation (Pearson et al., 2022; Reaney, 2022) and habitat fragmentation (Häussler et al., 2017), allow regulators to treat these problems as closer to point-source externalities.

Consider a social bad arising from agricultural land use, such as nutrient run-off (Griffin & Bromley, 1982; Kling, 2011) or elevated downstream flood risk due to land management (Dadson et al., 2017). Assume that agricultural land can be taken out of production and used for rewilding or for natural flood management, to reduce damage via some function $F(L_{NF})$, where once again L_{NF} denotes the area of land devoted to natural features. Aggregating this across Q farms in a catchment, adapting the notation of Holland and Yates (2015) to the land use case, the total benefit B from natural features can be defined as:

$$B(L_{NF}) = F\left(\sum_{q=1}^{n} \delta_q L_{NF}\right) \tag{4.9}$$

The coefficient δ_q represents the contribution to the externality from land use change at farm i. Two familiar special cases of regular damage functions are a) uniformly mixed pollution, in which $\delta_q = 1$ for every q, and b) constant marginal benefits, in which F is linear. In practice, neither of these cases may be relevant for policymaking.

Like in a traditional cap-and-trade regime, the regulator has set the allowance (or in this case, the minimum required area for NFM features) $\tilde{L}_{NF}$ and the marginal abatement costs across affected firms determine buyers and sellers (Montgomery, 1972). One-for-one trading allows farms that are willing to set aside more land for NFM to take over the obligations of another farm in exchange for payment. However, Fowlie and Muller (2019) observe that such a scheme does not effectively target producers of the largest externalities.

Following Holland and Yates (2015) we propose differentiated trading ratios as a solution to spatially heterogeneous damages to incentivise high-priority farms to take up the NFM obligations of low-priority farms. Farm q's objective is to minimise the following cost function:

$$p_X X + c_{NF} L_{NF} + \pi\left(\tilde{L}_{NF} - r_q L_{NF}\right) \tag{4.10}$$

The trading ratio r_q determines how much less ($r_q > 1$) or more ($r_q < 1$) land farm q would need to retire to take over the NF obligation of another farm. A trading ratio of 1 implies one-for-one trading (Holland & Yates, 2015). Setting up the Lagrangian:

$$\begin{aligned}\mathcal{L} = p_X X + c_{NF} L_{NF} + \pi \left(\tilde{L}_{NF} - r_q L_{NF} \right) - \\ \mu_1 \left(\bar{Y} - X^{\alpha} L_{AG}^{\beta} \right) - \\ \mu_2 \left(\bar{L} - L_{AG} - L_{NF} \right)\end{aligned} \tag{4.11}$$

Once again, differentiating with respect to X, L_{AG}, and L_{NF} yields the first-order conditions:

$$[X]: \quad p_X + \mu_1 \alpha X^{\alpha-1} {L_{AG}}^{\beta} = 0 \tag{4.12}$$

$$[L_{AG}]: \quad \mu_1 \beta X^{\alpha} {L_{AG}}^{\beta-1} + \mu_2 = 0 \tag{4.13}$$

$$[L_{NF}]: \quad c_{NF} - \pi r_q + \mu_2 = 0 \tag{4.14}$$

Rearranging (4.12) to solve for μ_1 and substituting into (4.13) lets me solve for μ_2. A function for L_{AG} can be written by entering μ_2 expressed as inputs, outputs and prices into equation (4.14) Like in the base model, the demand function for land to be retired for NFM projects is then derived.

$$L_{NF}^{*} = \bar{L} - \left(\frac{\frac{\beta}{\alpha} p_X \bar{Y}^{1/\alpha}}{c_{NF} - \pi r_q} \right)^{\frac{\alpha}{\alpha+\beta}} \tag{4.15}$$

Equation (4.15) can be substituted into the benefits function (4.9). For two farms 1 and 2 we can express r_1 and r_2 as r_1 and $1/r_1$, which the regulator selects to maximize:

$$B(L_{NF}) = \delta_1 \left[\bar{L} - \left(\frac{\frac{\beta}{\alpha}p_X\bar{Y}^{1/\alpha}}{c_{NF}-\pi r_1}\right)^{\frac{\alpha}{\alpha+\beta}}\right] + \delta_2 \left[\bar{L} - \left(\frac{\frac{\beta}{\alpha}p_X\bar{Y}^{1/\alpha}}{c_{NF}-\pi 1/r_1}\right)^{\frac{\alpha}{\alpha+\beta}}\right] \tag{4.16}$$

Maximising B with respect to r_1 when land endowments, residual demand, and prices are normalised, we see that only when marginal damages are uniform, i.e. $\delta_1/\delta_2 = 1$, is the optimal trading ratio one-for-one trading. To incentivise greater NFM uptake among high-risk farms when marginal damages are spatially differentiated, trading ratios should reflect the relative flood generation risk between the farms. As with the voluntary scheme introduced in the previous section, the cost-minimising area to be set aside for NFM decreases with the cost to implement NFM features c_{NF}. More interesting is the marginal change in demand for NFM given an increase in the annual per hectare payment π:

$$\frac{\partial L^*_{NF}}{\partial \pi} = -\left(\frac{\alpha}{\alpha+\beta}\right) r_q \frac{\left(\frac{\beta/\alpha p_x \bar{Y}^{1/\alpha}}{c_{NF}-\pi r_q}\right)^{\frac{\alpha}{\alpha+\beta}}}{c_{NF}-\pi r_q} \tag{4.17}$$

Once again, the marginal demand for creating NFM is increasing with the payment when $\pi r_q > c_{NF}$. In a scheme with minimum catchment-wide NFM and tradable contracts, a cost-minimising farmer q whose NFM creation costs (including opportunity costs) exceed the payment ($r_q\pi < c_{NF}$) would buy out of their NFM requirement (i.e. paying another farmer to create it). This also happens if the payment is less than the farm's opportunity cost of taking land out of production to create NFM. In this way, the trading ratio r_q governs whether farm q will take over the NFM contracts of another farm or pay to absolve itself of its current NFM obligation. Figure 4.2 shows NFM demand curves for a set of different trading ratios and sizes of β when the NFM creation cost is negligible.

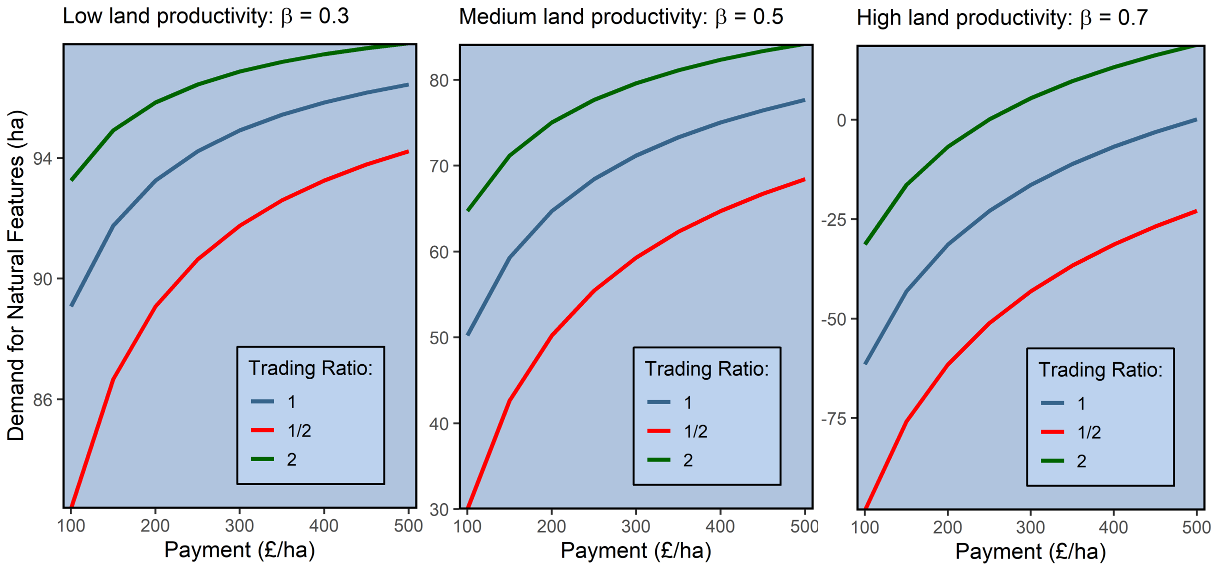

Figure 4.2: *Illustrative demand curves for NFM for a 100 ha farm with a residual demand of 500 units of agricultural output Y. Assume that the cost of creating natural is only opportunity cost. A negative demand (as in the panel showing demand curves for farmers of high land productivity) means that the farmer will want to buy out of NFM contracts.*

If the null of hypothesis I is rejected, lending credibility to the base model, the following hypothesis can be added with the aim of validating the extension to the model involving trading. Hypothesis II tests whether farmers a) understand the concept of a market for NFM contracts with spatial targeting and b) confirms that they would act rationally within it.

Hypothesis II: Increases in the trading ratio offered to a farmer (e.g. from 1/5 to 5 to 10) lower the minimum payment she is willing to accept to create additional NFM features.

Hypothesis II recognises from the demand function for NFM that when $(\pi r_q - c_{NF}) > 0$, farmers wish to create more NFM on their own land as the annual payment π increases. This is because a farmer for whom the risk-adjusted government payment, $r_q\pi$, exceeds the cost of creating NFM would be accepting payment to take over additional NFM obligations from other farmers. Therefore, an increase in the risk-adjusted payment encourages her to create more NFM. The NFM demand function also predicts that under the trading regime, an increase in the risk-adjusted payment will lower a farmer's demand for NFM while $(\pi r_q - c_{NF}) < 0$. Perhaps less intuitive, this result arises because while the farmer will always seek to get out of her NFM obligation while the cost exceeds the payment, other farmers who wish to take on more NFM contracts via the trading market benefit from a higher payment. Increased demand for NFM by farmers of high-risk land (high r) improves the opportunity for others (low r) to buy out of their NFM obligations. Similarly, steep gradients in the trading ratios facilitates trading and increases the market-clearing concentration of NFM among the most high-risk farms.

4.3.3 Transaction costs

Trading is likely to involve transaction costs arising from matching farmers, facilitating transactions, and setting up intermediaries to do so (Nguyen et al., 2025). Transaction costs can incorporated into the model as a percentage of the payment. Irrespective of whether it falls on the farmer buying out of an NFM obligation, on the farmer taking up additional obligations, or is shared between the two, a higher transaction cost is expected to result in fewer trades. I add transaction costs τ to the farm's cost function as a percentage of the total cash value of any trade. The corresponding adaptation to equation (4.10) then becomes:

$$p_X X + c_{NF} L_{NF} + T\pi \left(\tilde{L}_{NF} - r_q L_{NF} \right) \tag{4.18}$$

where T is equal to $(1+\tau)$ when the farmer is buying out of their NFM obligation (net expenditure is positive) and is equal to $(1-\tau)$ when the farmer is accepting payment to take up additional NFM (net expenditure is negative). Solving the Lagrangian for cost function (4.18) in the same manner as in section 4.3.2, I derive the marginal demand function for NFM, L_{NF}, in terms of the payment rate π shown in equation (4.19). See the appendix 6 for a step-by-step derivation of the marginal demand function from the Lagrangian.

$$\frac{\partial L^*_{NF}}{\partial \pi} = -\left(\frac{\alpha}{\alpha+\beta}\right) r_q \frac{\left(\frac{\beta/\alpha p_x \bar{Y}^{1/\alpha}}{c_{NF}-T\pi r_q}\right)^{\frac{\alpha}{\alpha+\beta}}}{c_{NF}-T\pi r_q} \tag{4.19}$$

The function shows that L_{NF} demanded by farm q is growing with π when the payment πr_q exceeds the cost of creating NFM, c_{NF}. In this case, $T = (1-\tau)$ and so the rate of NFM creation declines with the transaction cost τ. The opposite is true when $\pi r_q < c_{NF}$, when τ reduces the rate of farmers buying out of their NFM obligations.

Hypothesis III: Increases in the transaction cost faced by the farmer results in a reduction in trade volume irrespective of whether the farmer's demand for NFM contracts is positive or negative.

4.4 Econometric modelling

Each hypothesis was tested using discrete choice modelling with latent classes as described in section 3.5 of chapter 3. Hypotheses I and II are tested based on the results from DCE I. Hypothesis III is tested using the WTA and the WTP scenarios that are part of DCE II. Following Boxall and Adamowicz (2002), the number of classes was decided based on minimising the Bayesian Information Criterion (BIC). Models with $2-4$ classes were estimated, but with no more than two classes did the models converge. The BICs for the two-class model were consistently lower than the BIC for the base MNL model. Accordingly, models with two latent classes were estimated for each DCE.

4.4.1 DCE I

Table 4.2 reminds readers of the choice attributes and levels of DCE I. The corresponding variable notation from the theory section 4.3 has been added next to the attribute names to bridge the gap in notation between the economic model and the econometric model. Both the type of natural feature and its location were assumed to be drivers of the cost to farmers of creating it, c_{NF}.

Table 4.2: *DCE I: Attributes and levels*

Attribute	Levels
Type (c_{NF}): *The type of NFM feature*	Natural Regeneration, Planted Broadleaf Trees
Location (c_{NF}): *Where the NFM feature is placed on the farm*	1) Mid-field, 2) Field boundary, 3) River edge
Land quality (proxy for β): *Suitability of land for agriculture*	1) Rough grazing, wet, steep, rocky etc., 2) Prime grazing land or high yielding crops
Area (L_{NF}): *Amount of land set aside for NFM*	1) 1/20 hectare (500m^2), 2) 1/10 hectare (1000m^2)
Payment (π): *Annual payment*	£200, £300, £400, £500

Hypotheses I-II were tested by estimating taste parameters for individual attributes in DCE I. This was done by estimating the latent class model with the utility from option (ELM scheme) i specified as follows:

$$
\begin{aligned}
U_{s,i} = ASC_{i,s} + ASC_{i,s} \times FEMALE + ASC_{i,s} \times GRAZING + \\
\beta_{TREES,s} \times TREES + \beta_{RIVEREDGE,s} \times RIVEREDGE + \\
\beta_{FIELDEDGE,s} \times FIELDEDGE + \beta_{QUALITY,s} \times QUALITY + \\
\beta_{AREA_{1000m^2},s} \times AREA_{1000m^2} + \beta_{PAYMENT,s} \times PAYMENT + \\
\lambda_L + \delta_s
\end{aligned}
\tag{4.20}
$$

Equation (4.20) models the utility that farmers in class s derive from choosing option i. The attributes are described in table 3.4. The alternative-specific constant, $ASC_{i,s}$, is interacted with a dummy variable indicating whether the respondent is female and with the proportion of land the respondent uses for grazing. Other interactions, including educational attainment and current enrolment in ELM schemes, were tested and found insignificant.λ_L represents the land endowment elasticity and is an estimation of how the sensitivity to larger ELM features, $\beta_{AREA_{1000M^2}}$, varies with respondents' land endowment. δ_s is an offset describing, on average, to what extent the utility of class s is different from that of class 1.

Testing the first hypothesis seeks to confirm the validity of the base model of farmers as cost-minimisers. Implicit in the assumption about the functional form of agricultural production is that the marginal productivity of land is strictly positive and diminishing, i.e. $0 < \beta < 1$. It follows that the required payment to accept a 1/10th hectare feature over a smaller 1/20th hectare feature is lower when the initial area of the farmer's productive land, L_{AG}, is high. Hess and Palma (2019) and Axhausen et al. (2008) illustrate how the income elasticity can be computed by estimating $\beta_\pi \times (Y/\bar{Y})^{\lambda_Y}$ where Y is income and λ is the elasticity.

The estimate of λ_Y gives the elasticity of the sensitivity to price with respect to changes in Y. With negative elasticity, the (absolute) sensitivity decreases with increases in Y, with the opposite applying in the case of positive elasticities. Finally, the rate of the interaction is determined by the absolute elasticity, where a value of 0 indicates a lack of interaction. I similarly estimate how farmers' land endowment affect their demand for land retired for NFM: $\beta_{L_{NF}} \times (L/\bar{L})^{\lambda_L}$. Rejecting the null requires that $\beta_\pi > 0, \beta_{L_{NF}} < 0$, and $\lambda_{L_{NF}} < 0$. This implies that the dis-utility from larger NFM feature size decreases with the farm size. Hypothesis I was stated as the following null and alternative hypotheses:

$$\text{H0: a) } \beta_{PAYMENT} \leq 0 \leq \beta_{AREA}$$
$$\text{H0: b) } \lambda_L = 0$$
$$\text{H1: a) } \beta_{AREA} < 0 < \beta_{PAYMENT}$$
$$\text{H1: b) } \lambda_L < 0$$

$\beta_{PAYMENT} > 0$ means that farmers prefer a higher payment for creating NFM. $\beta_{AREA} < 0$ implies that farmers would prefer less land for NFM, holding the payment constant. These inequalities are necessary conditions for cost-minimising behaviour. Rejecting the null lends credibility to the the theoretical model. $\lambda_L < 0$ means that farmers managing large land areas are less sensitive to the land area set aside for NFM. Such a result supports the assumption that the land factor productivity is positive but diminishing. In other words, that the marginal productivity of land as an agricultural input factor is decreasing with the land area put to use.

Alternative hypothesis H1 a) is a joint inequality, which can be evaluated using simulated draws from a joint distribution. When β_{AREA} and $\beta_{PAYMENT}$ are estimates from a maximum likelihood estimation:

$$N \to \infty, \theta_{ML} \sim \mathcal{N}(\theta, \Omega) \tag{4.21}$$

The variance-covariance matrix Ω for β_{AREA} and $\beta_{PAYMENT}$ was extracted from the latent class logit model. When the sample N is large enough, it is possible to sample R times from the asymptotic normal distribution using the vector of taste parameters as the mean and Ω obtained from the model. After R draws, the cases for the inequality of interest were counted. The statistic to report for the probability of $H1$ a) can be computed as:

$$\frac{\sum_{r=1}^{r=R} 1(\beta^r_{PAYMENT} > 0 > \beta^r_{AREA})}{R} \tag{4.22}$$

The resulting fraction represents the proportion of cases from $10,000$ simulated draws where the null hypothesis is true, and can be compared against the significance threshold which is 5%. Failure to reject H0 a) suggests that surveyed farmers do not display cost-minimising behaviour. Failure to reject H0 b) would suggest that large farms are not less sensitive to increases in the land area set aside for NFM. For example, because farmers perceive that the costs of creating and maintaining NFM features dwarf the opportunity cost of agricultural land.

4.4.2 DCE II

Individual latent class models are specified and estimated for the WTA and WTP scenarios. The attributes are identical, except for the trading ratios and the payment, as per table 4.3. As described in section 4.3, the trading ratio r governs the rate of exchange between farms in terms of the land area required to meet the conditional payment for NFM. For the farmer taking on additional NFM obligations, a trading ratio above one means that they will have to retire proportionally less land. For the farmer buying-out of their NFM obligation, the corresponding ratio below one means that they will pay less. From the government's perspective, this is all motivated by a higher marginal flood risk reduction if NFM is created at the high-risk farm. The constrained cost-minimisation described in section 4.3 supposed a continuous demand function for NFM where demand could be positive (farmer is a net taker of additional NFM obligations) or negative (farmer is a net buyer-out). The one departure from here on out is that the econometric models are estimated based on six choice tasks each from discrete WTA and WTP scenarios in DCE II. In the WTA scenario, trading ratios are always above one. In the WTP scenario, trading ratios are always less than one.

Table 4.3: *DCE II: Attributes and levels*

Attribute	Levels
Trading ratio (WTA, r): *The factor by which respondents can increase their per-hectare payment for NFM by trading*	5, 10, 20
Trading ratio (WTP, $1/r$): *The ratio by which the respondent can reduce their expected per-hectare cost to get out of their NFM obligations by trading*	$1/5$, $1/10$, $1/20$
Transaction fee (τ): *A percentage of the base payment borne by the respondent*	5%, 10%
Payment (π): *Annual payment, received in the WTA setting and paid in the WTP setting*	£200, £300, £400, £500

As in DCE I, models featuring $2-4$ classes were estimated, and convergence was achieved only with 2 classes. The model is presented in equation (4.23):

$$\begin{aligned} U_{s,i} = ASC_{i,s} + \beta_{r=10:1,s} \times (r = 10:1) + \beta_{r=20:1,s} \times (r = 20:1) + \\ \beta_{FEE,s} \times FEE + \beta_{PAYMENT,s} \times PAYMENT + \\ \delta_s \end{aligned} \tag{4.23}$$

The estimates of the taste parameters for trading ratios of 20:1 and 10:1, respectively, are measuring the preference relative to the reference level 5:1. To test hypothesis II, the following alternative and null hypotheses were stated:

$$\text{H0: } \beta_{r=20:1} = \beta_{r=10:1} = 0$$
$$\text{H1: } \beta_{r=20:1} > \beta_{r=10:1} > 0$$

Similar to hypothesis I, H1 is a joint inequality. Accordingly, the same procedure for testing is followed.$10,000$ draws from the bivariate normal distribution that satisfy the H0 are summed up. The proportion is then compared against the significance level of the test, which is again 5%.

Hypothesis III posits that the transaction cost (FEE or τ) reduces the volume of trade. Rejecting the null requires that $\beta_\tau < 0$ in each scenario such that a higher transaction cost τ requires a higher payment rate to facilitate trading. For both trading scenarios, the alternative and null hypotheses are stated as follows:

$$\text{H0: } \beta_{\tau=10\%,s} = 0$$
$$\text{H1: } \beta_{\tau=10\%,s} < 0$$

4.5 Estimating trading ratios and runoff reduction

As discussed in previous sections, the contribution of this chapter goes beyond adopting a theory of spatially targeted cap-and-trade to NFM and testing its predictions in a choice experiment with active farmers. This part of the work can be thought of as the costing portion in a cost-benefit analysis of the hypothetical trading regime for NFM in England. In this chapter, I also aim to make explicit the benefits portion of the analysis. I first estimate the reduction in water runoff generation risk attributable to a set of variants of the NFM schemes featured in the choice experiments. I do this by comparing runoff generation potential δ across real agricultural landscapes and simulated landscapes with NFM features of different types and placements.

The SCIMAP-Flood model (Reaney, 2022) is a spatially distributed tool designed to identify critical source areas for floodwaters within a catchment, thereby aiding in the prioritisation of natural flood risk management (NFM) interventions. The model has been validated in the Eden catchment. By analysing spatial patterns of rainfall, land cover, and topography, SCIMAP-Flood determines locations where mitigation measures, such as storage ponds, flow-slowing debris dams, and land-use changes, would be most effective in attenuating flood peaks. The output map from SCIMAP-Flood combines relative scores to each of the flood hazard driving factors and then combines these to give a point scale assessment of the potential value of slowing flows at that location for decreasing flood generation (Reaney, 2022).

A key feature of SCIMAP-Flood is its ability to handle uncertainties in input data, particularly variations in rainfall patterns and land cover information. This probabilistic approach enables the model to provide not only potential sites for NFM interventions but also the confidence levels associated with these predictions. Such information is crucial for decision-makers aiming to implement effective and reliable flood mitigation strategies. The model operates by assessing the hydrological connectivity within a catchment, identifying areas where surface runoff is likely to contribute significantly to flood events.

By targeting these critical source areas, SCIMAP-Flood facilitates the strategic placement of NFM measures, enhancing their overall effectiveness in flood risk reduction. SCIMAP-Flood has been applied in various contexts, including catchments in the UK and Nepal (Pearson et al., 2022), demonstrating its adaptability to different environmental conditions. Its development was initiated following Storm Desmond in 2015, with the aim of introducing innovative approaches to catchment-based flood hazard management. The model has since undergone testing and refinement, incorporating feedback from diverse applications to improve its accuracy and reliability (Reaney, 2022). Figure 4.3 illustrates how `SCIMAP-Flood` uses input data (shown in blue) to produce an assessment of the relative flood risk at every parcel of land in the catchment. The model takes the following inputs:

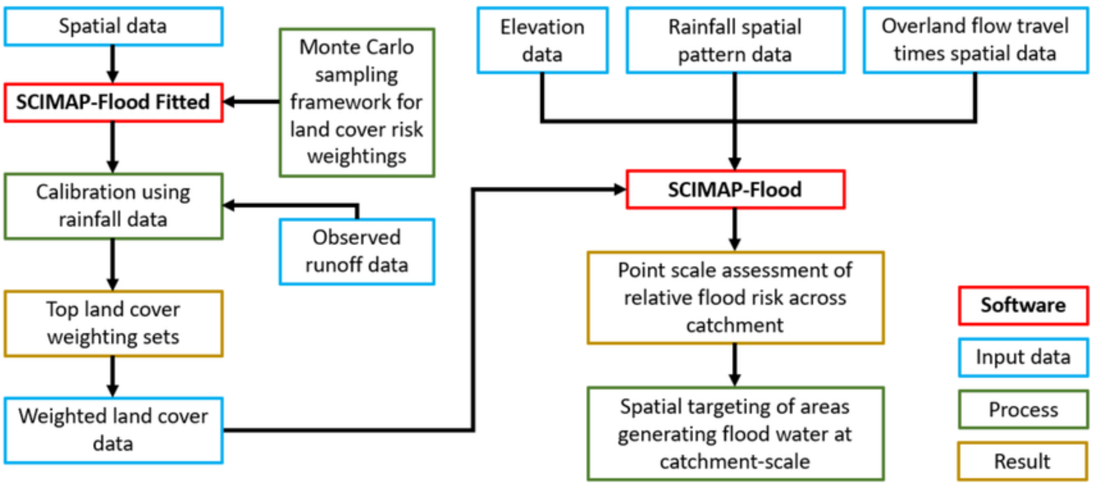

Figure 4.3: *Diagram illustrating the process of executing* `SCIMAP-Flood`*, from Pearson et al., 2022*

Land cover data: Different rural land uses contribute differently to runoff through variation in soil permeability (Pattison & Lane, 2012) and surface roughness (Reaney, 2022). The former governs the capacity of soil to absorb water at scale and speed. Permeability can be impacted for example by ploughing of fields, where heavy machinery causes wheel tracks to be compacted. The latter refers to the capacity of different land cover classes to inhibit water flows. For example, natural regeneration of inactive farmland increases roughness compared to bare-ground fallow (Niehoff et al., 2002).

This research uses the 2022 edition of the 10m^2 resolution land cover maps made available by the UK Centre for Ecology and Hydrology (UKCEH). The dataset features 21 unique land cover classes, and is created by combining many classified images into a single map of the whole country. A random forest supervised learning classifier is used to estimate the likelihood of each type of land cover. The land cover type with the highest likelihood is selected as the most probable. In the dataset, the first layer contains numbers that represent the most likely UKCEH land cover type. The second layer shows the probability for that classification, giving an idea of how confident the result is.

Unlike earlier UKCEH datasets, the 10m^2 pixel data has not been simplified by merging it with the UKCEH Land Parcel Spatial Framework. This means that it keeps detailed features like narrow strips or small areas of habitat that are too small to be shown in the 0.5-hectare minimum mapping size used in the UKCEH Land Parcel Spatial Framework (Marston et al., 2023). Following Reaney (2022) the UKCEH land cover classes have been converted to *runoff weights* between zero and one, where a higher weight signifies greater runoff generation potential. For example, the weight for unimproved grassland is 0.15 while the weight for arable land is 0.8. The distribution of runoff weights is mapped in figure 4.4. It shows that the flood generation potential is greatest in the north of the catchment, around the town of Carlisle. Urbanisation contributes to flood risk due to clearing of vegetation, paving of road, etc. (Niehoff et al., 2002).

Elevation: Higher elevations tend to experience more intense and rapid runoff due to steeper slopes, which accelerate the flow of water, reducing infiltration and increasing the risk of downstream flooding. In contrast, low-lying areas often serve as collection points for runoff, making them more prone to water accumulation and potential flood events. I used a digital terrain model (DTM) made available by the UK Environment Agency in 2022 with a $10m^2$ resolution. The DTM is derived from a combination of the Agency's Time Stamped archive and National LIDAR Programme surveys, which have been merged and re-sampled to give the best possible coverage. Where repeat surveys have been undertaken the newest, best resolution data is used. Where data was resampled a bilinear interpolation was used before being merged (UK Environment Agency, n.d.). The elevation is mapped in figure 4.5. The river Eden is visible as it flows north through the catchment toward its mouth at Solway. Following Reaney (2022), slope rasters were created from the DTM using mapping software SAGA.

Hydrological connectivity: The slope and land cover rasters, along with mapping of the river network supplying the experimental catchment, are used to compute the *hydrological connectivity* of the catchment.[1] A map of hydrological connectivity (figure 4.6) shows the paths water will traverse a landscape. It is a measure of the ease by which a volume of water is able to move from one point to another (PEARSON et al., 2016).

Rainfall patterns: Pattison and Lane (2012) observed that regional rainfall circulation patters have an impact on the kinds of storms that have resulted in severe flooding in the past. Precipitation maps were selected from the CEH Gridded Estimations of Areal Rainfall, GEAR, dataset (Tanguy et al., 2021). This dataset comprises daily rainfall estimates based on the observed rain gauges presented in

[1]Hydrological connectivity, although etymologically related, is distinct from *habitat- or ecological connectivity* which will feature in a later chapter. When referring to hydrological connectivity, I will make this explicit throughout.

$1km^2$ resolution. Rainfall records from six days with the heaviest precipitation in 2019 were used.

SCIMAP-Flood was run across the Eden catchment which is a largely rural. flood-prone area in the north-west of England, and is also the home of several of the participants in our choice experiments. The mapping of runoff weights (figure 4.4) shows the urbanised areas in the catchment light up as major runoff generation hotspots. This results from the land use change from natural vegetation to a built up environment, including paving over roads, inherent in urbanisation. I also illustrate the correlation between the runoff weights, elevation, and connectivity. Connectivity is higher where slopes are steeper.

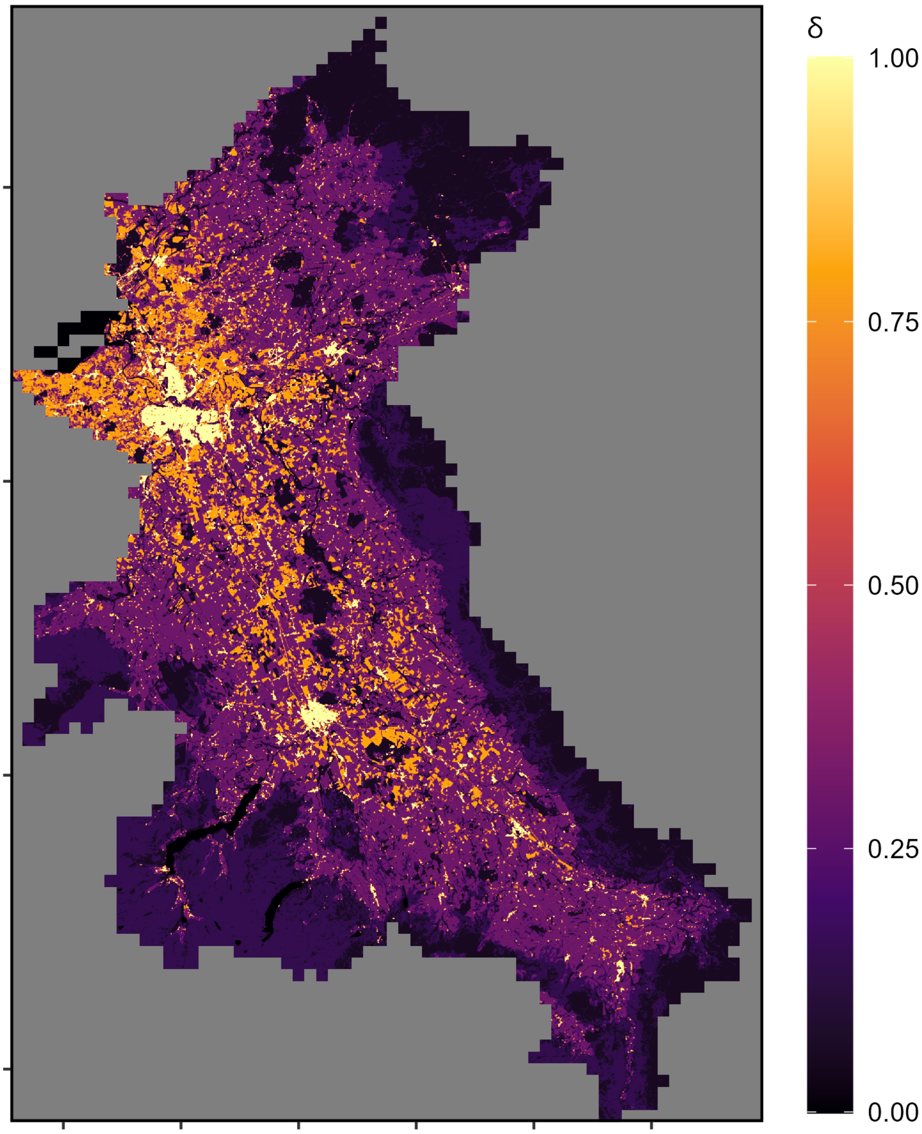

Figure 4.4: *Surface water runoff weights, δ, indicating the relative flood risk driven by geography and land use.*

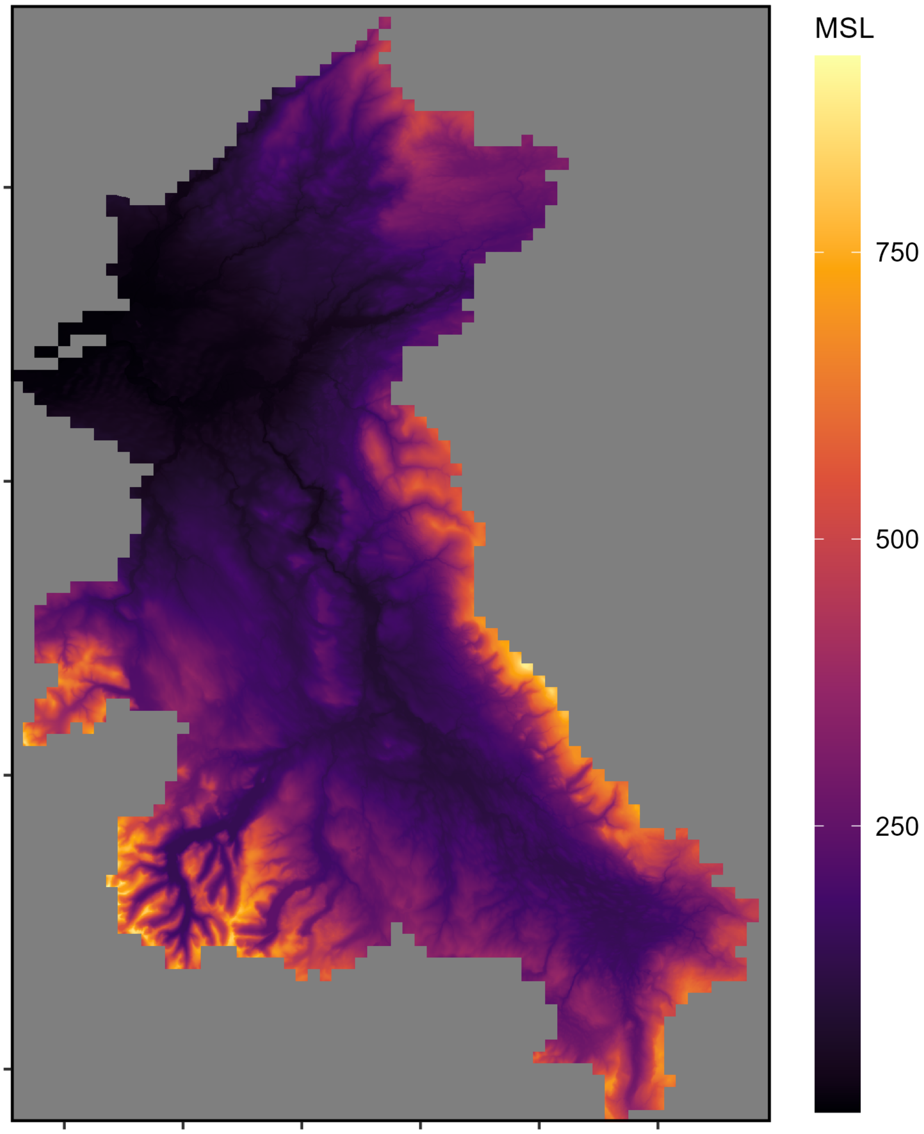

Figure 4.5: *Elevation from a digital terrain model of the UK (meters above sea levels) (UK Environment Agency, n.d.)*

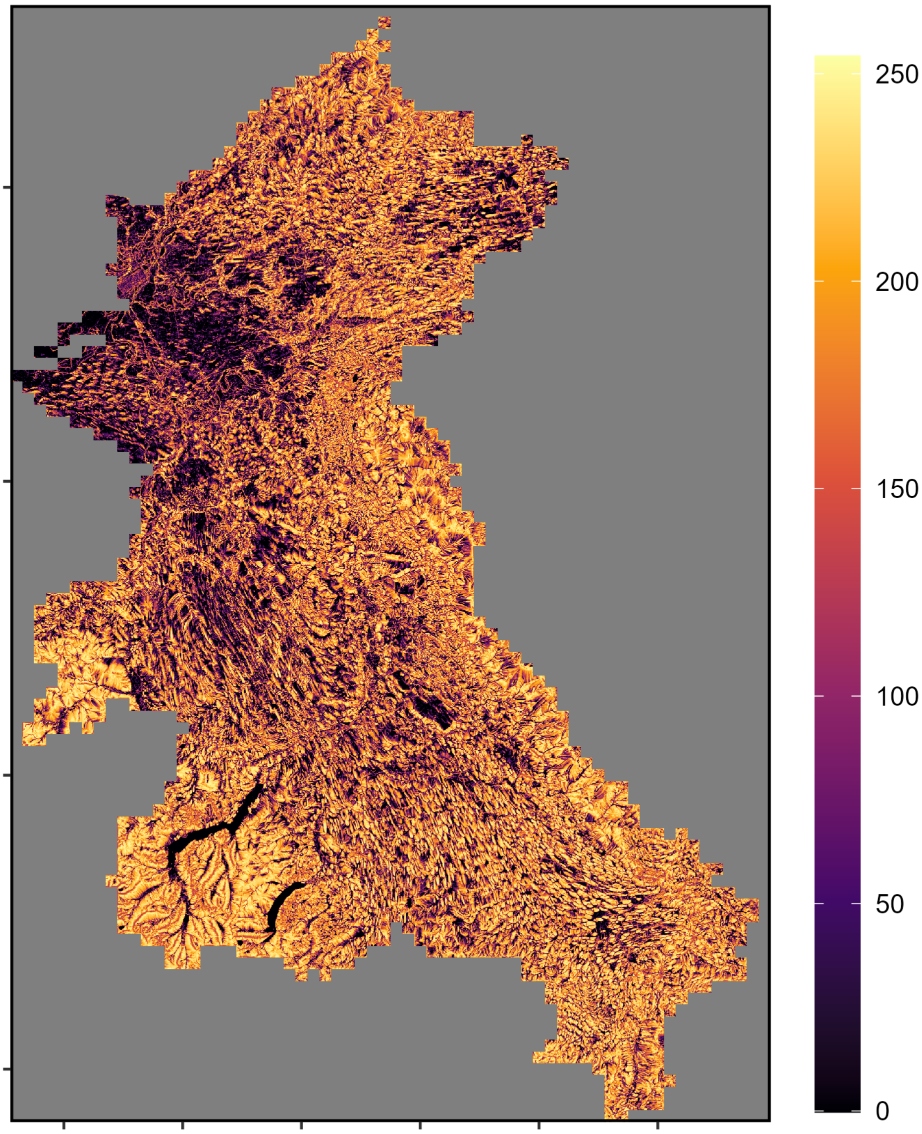

Figure 4.6: *Hydrological connectivity (Reaney, 2022) which describes the ease with which water from one location in the landscape can move to another*

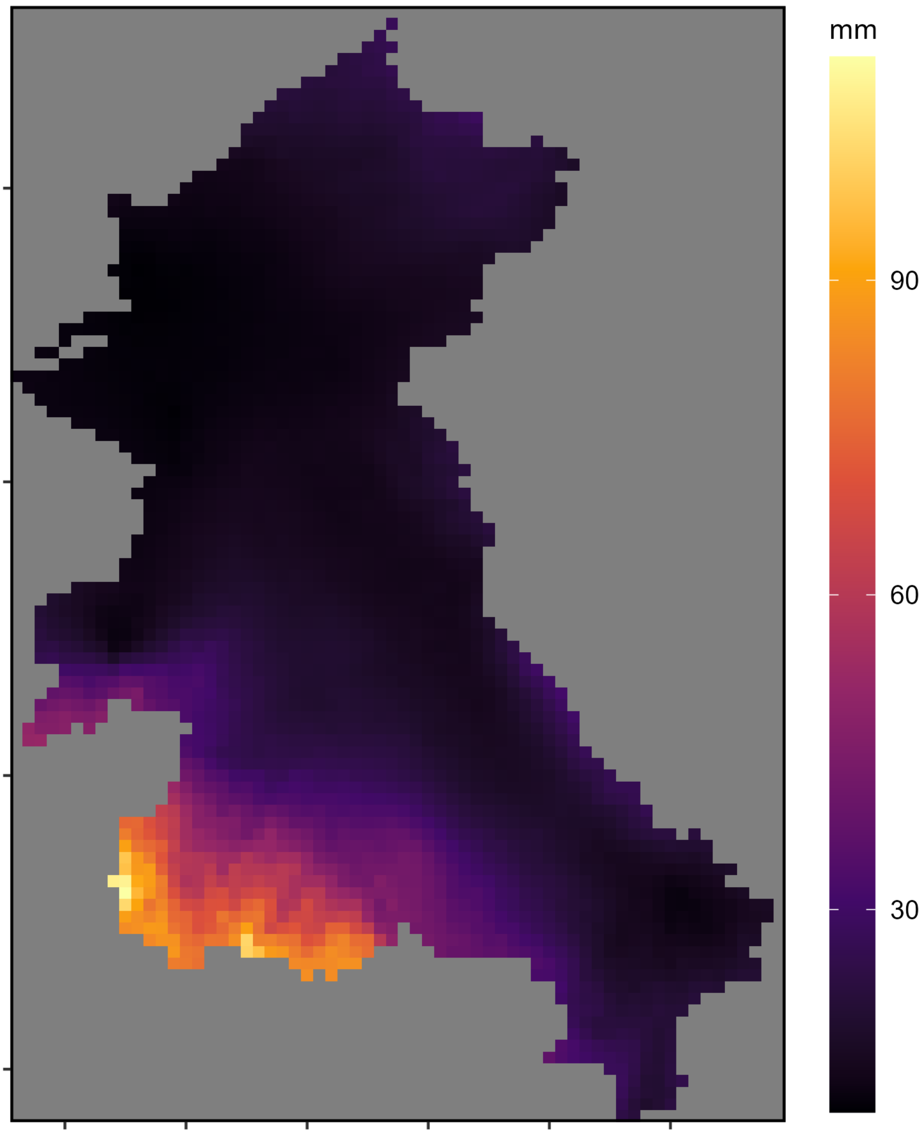

Figure 4.7: *Rainfall patterns (mm) over the Eden, recorded on 10 December 2019, one of the heaviest rains of that year (Tanguy et al., 2021)*

4.6 Results

I begin by reporting the results from the first choice experiment, DCE I, which aims to test hypothesis I as well as provide some general information about the perceived barriers to ELM participation within the sample. Two models are presented using the results from DCE I. In the first model I am primarily concerned with predictors of NFM uptake. I seek to benchmark the experimental evidence against the qualitative insights from Hurley et al. (2022) and Holstead et al. (2017) as well as previous DCE studies (Tyllianakis et al., 2023) about what demographic and economic characteristics present barriers to ELM enrolment. Tyllianakis et al. (2023) report that their class of older, full-time farmers without extensive experience with ELM schemes nonetheless exhibit a strong aversion to being left out of new ELM schemes and choose not not opt out of the hypothetical schemes. Farmers' preference for enrolling in the schemes are correlated with previous or current experience with NFM schemes, farmers' ages and pro-social attitudes. A suitable model for this purpose is a latent class model, also used in Tyllianakis et al. (2023), which allows me to group respondents into distinct classes based on their preferences. Each respondent has a posterior conditional probability of belonging to each class. I assign respondents to the class where their conditional probability is at least 80%. R code for this procedure is included in the appendix. I then illustrate how demographic and psychological differences predict class membership.

The second model is a mixed logit (MMNL) model, where I allow inter-individual taste heterogeneity and estimate taste parameters specific to each respondent. For each class of respondents, the mixed model allows me to establish how the amounts of money respondents ascribe to attributes of the NFM schemes are distributed in the sample. Narrow distributions give policymakers a good idea of the required payment for a particular scheme, while distributions with a high variance indicate that a ”one-size-fits-all” design may be infeasible. Parameters are drawn 1,000 times using a Modified Latin Hypercube Sampling (MLHS) algorithm which has been shown to outperform alternative Halton draws for MMNL models (Hess et al., 2006). The initial normal distributions from which taste parameters are drawn have means set to zero. The payment parameter is drawn from a uniform distribution between 0 and 1. The MMNL models were estimated using the Apollo package in `R (v4.1.3)` (Hess & Palma, 2019).

I report results from DCE II in the same way, displaying results from a latent class model and augmented by a mixture model in order to understand the distribution of tastes within the sampled farmers. I test Hypothesis III, which claims that farmers who do not believe that their land contributes to flood risk are less likely to value an increase in the trading ratio. I leverage results from previous research on consequentiality (Lloyd-Smith & Adamowicz, 2018) to predict that this would happen because these farmers doubt that a policymaker would ever assign their land a high trading ratio. I test the hypothesis by interacting (Block et al., 2024) the trading ratio attribute with respondents' stated concern about flooding in the catchment. The regression coefficient for this interaction represents the difference in preference for higher trading ratios between farmers expressing belief in the flood risk of their land and those who do not.

Then, NFM schemes from the experiments are compared in terms of reduction in runoff generation potential using `SCIMAP-Flood`. I incorporate estimates of required payments per scheme from the choice experiment to conduct a cost-benefit analysis, comparing reductions in flood risk per amount spent on compensation to farmers. This allows me to evaluate the cost-effectiveness of the schemes and make policy recommendations. Finally, enable trading between simulated "high-risk" and "low-risk" farms to illustrate improvements in cost-effectiveness per my theoretical model.

The alternative-specific constants for ELM schemes A and B are both negative and significant compared against the constant for the opt-out, or status quo (SQ), alternative. In agreement with Tyllianakis et al. (2023) this means that respondents display a statistically significant preference for opting into the schemes.

4.6.1 DCE I: Barriers to enrolment into NFM schemes

Table 4.4 shows the results from the latent class model. This research finds that farmers in the sample can be grouped into two distinct classes that are significantly different. Class I make up 73% of the sample while Class II make up 27%. The standard errors and resulting statistical significance must be be read with that difference in class size in mind. The notable differences in taste parameters between the two classes are the alternative-specific constants for the NFM schemes. $\text{ASC}_{SchemeA}$ and $\text{ASC}_{SchemeB}$ measure respondents' preference for enrolling in the NFM scheme compared against opting out. Positive and statistically significant taste parameters mean that, all other attributes being equal, respondents prefer the feature in question over the reference level. Conversely, negative and significant values mean that respondents prefer the reference level. Statistically insignificant taste parameters mean that respondents are indifferent between the two.

On average, members of Class I prefer enrolment into any available NFM scheme over opting out, while Class II prefer opting out. In other words, before attributes of the schemes are considered, Class II can be characterised as NFM *sceptics*. Among members of Class I, the preference for enrolling in the scheme is also higher among women. There is no statistically significant gender difference in Class II. Across both classes, farmers who use a larger proportion of their productive land for grazing (as opposed to cereals, soybeans, horticulture, etc) have a stronger preference for enrolment in the scheme. I hypothesise that this is because the types of NFM features involved in the schemes are less disruptive to grazing. For example, natural regeneration may involve only fencing off the protected areas. Farm size is not a significant predictor of scheme enrolment in either class.

Taste parameters for planted trees are negative for both classes, although the difference compared to natural regeneration is only statistically significant for Class I. This means that respondents would prefer to maintain natural regeneration features rather than planting trees. This result is in line with expectations. Similarly in line with expectations, I find that placing the NFM features either along a river edge or along the field boundary is each preferred over in-field features. However only the river edge parameter is significant within Class II. Retiring good quality land (e.g. prime grazing) is only moderately worse than lower quality land according to the respondents, only significant at the 10% level in Class I. The interaction between size of the NFM features and the size of the farm is insignificant, suggesting that small farms are not more unwilling to increase the area devoted to NFM.

Figure 4.8 shows how respondents in the two latent classes differ along a number of key characteristics. There is a major difference in the propensity to choose the opt out alternative, with Class II (27% of respondents) much more likely to decline enrolment in either available scheme. Farmers of small land areas are more likely to display preferences of Class II. So are those who are not already enrolled in an ELM scheme, and those who do not collaborate with neighbours in farming activities. Taking these differences into account, Class I is called the the ”high engagement” class, and Class II the ”low engagement” class. I choose this naming

Table 4.4: *DCE I: Preferences for NFM schemes*

Attribute	Taste Parameters		Reference level
	Class I	Class II	
$ASC_{SchemeA}$	1.98	-1.88	ASC_{Optout}
	(0.21)***	(0.27)***	
$ASC_{SchemeB}$	1.85	-1.98	ASC_{Optout}
	(0.21)***	(0.27)***	
Trees	-0.28	-0.19	Natural Regeneration
	(0.05)***	(0.13)	
River Edge	0.61	0.81	In-field
	(0.06)***	(0.14)***	
Field Boundary	0.66	0.11	In-field
	(0.08)***	(0.17)	
Good Quality Land	-0.05	-0.10	Poor Quality
	(0.05)*	(0.11)	
$1000m^2$	-0.26	-0.28	$500m^2$
	(0.05)***	(0.11)**	
Payment	2.09	3.26	
	(0.22)***	(0.54)***	
ASC_{Scheme} × Female	1.21	-0.13	
	(0.53)**	(0.19)	
ASC_{Scheme} × % Grazing	0.01	0.05	
	(0.005)**	(0.002)**	
Elasticity of Land	0.15	-0.27	
	(0.25)	(0.34)	
Summary of class allocation for model: Class 1 (73%) and Class 2 (27%)			
Adj. R^2 vs observed shares: 0.21, BIC: 4750, MNL BIC: 5721			

convention because not only are members of Class II less likely to engage with the hypothetical schemes in the choice experiment, they are also modestly less likely to engage with real ELM schemes, nor engaging with neighbouring farmers such as sharing farming equipment. Differences in educational attainment are less clear-cut between the classes. While low engagement farmers are more likely to state that their highest qualification is some sort of vocational certification, they are also more likely to have attained a postgraduate degree.

Test of Hypothesis I: A positive and significant taste parameter for the payment and a negative and significant taste parameter for larger NFM features partially rejects the null for Hypothesis I. The proportion of draws satisfying the null hypothesis is 2.2% for the low engagement class and 0% for the high engagement class. Each are below the significance cutoff at 5%. This suggests that the factor productivity of land $\beta > 0$. However the elasticity between the size of NFM features and the farm's land endowment is not different from zero at any significance level. This means that regression analysis is not enough to confidently confirm that $1 > \beta$, i.e. whether there are diminishing returns to land inputs. I attribute these results either to $\beta \to 1$, or omitted variables. For example, respondents reporting large land endowment is correlated with reporting farming as their primary source of income (0.2) which may dissuade them against more NFM.

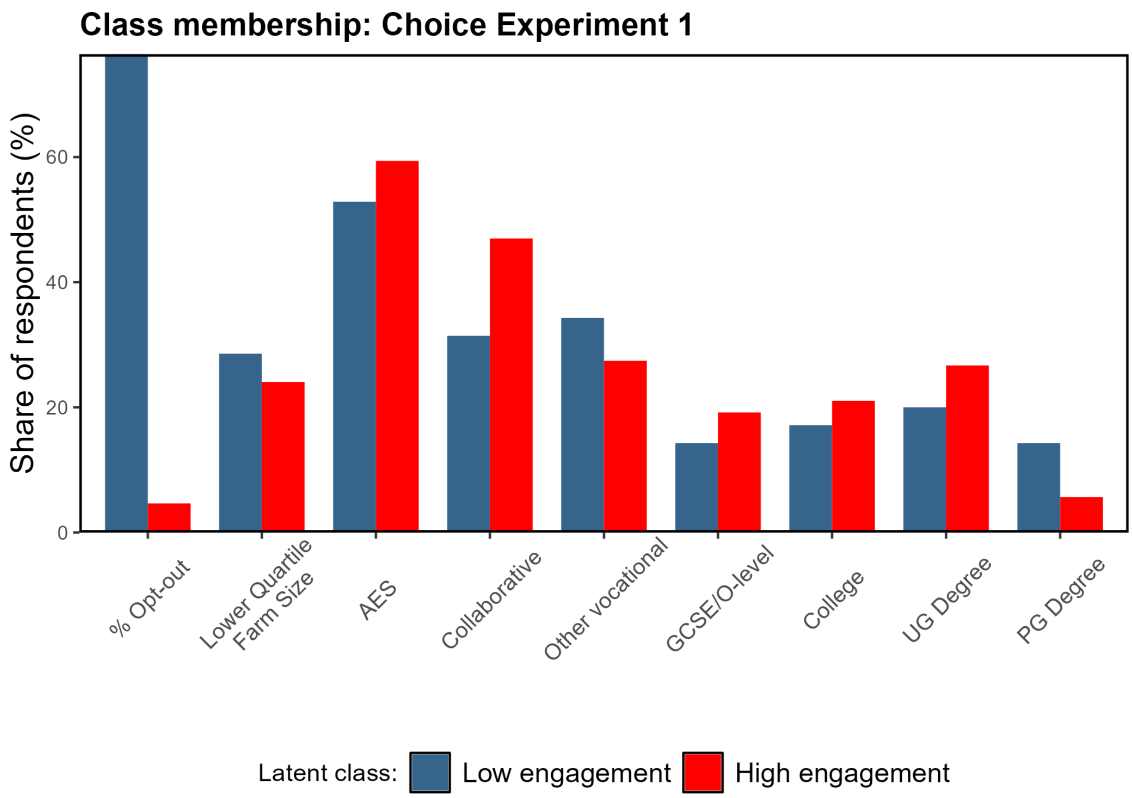

Figure 4.8: *Socio-demographic and behavioural predictors of latent class membership in choice experiment 1*

4.6.2 DCE II: Farmers' willingness to engage in trading

Next I report the results from the second choice experiment, which consists of two sets of choice tasks. In the first set, respondents are asked to imagine that they are in a position to receive extra payment by taking over the NFM obligation of other farms. The trading ratios being offered within this set are 5, 10, or 20. In the second set, respondents are instead asked to imagine that they can pay to get out of their NFM obligation. The trading ratios being offered within this set are once again 5, 10, and 20. In my theoretical model, these correspond to $1/5$, $1/10$ and $1/20$ but have been explained only verbally to respondents.

The latent class results from the first set are displayed in table 4.5. I once again identify two distinct classes, with a moderately stronger split than in the first choice experiment, 86% and 14% of the sample respectively. Members of Class I display a significant preference for engaging in trading. Class II is on average indifferent between engaging in trading and opting out. Class I displays a positive and significant preference for being offered a trading ratio of 10 over a ratio of 5. The taste parameter among members of Class II, however, is not statistically distinguishable from zero. As would be expected expect, the preference for a trading ratio of 20 is greater still among Class I, preferred over both ratios of 5 and 10. Class II defies expectations, as the taste parameter is negative. However, it is only significant at the 10% level. Members of Class I have strong preferences for lower transaction fees and higher payments, which is in line with cost minimising behaviour. However the preference for lower fees is weak among Class II and only statistically significant at the 10% level.

Figure 4.9 shows how respondents in the two latent classes differ along a number of key characteristics. It shows that the predictors of class membership in the second choice experiment are identical to the first. Once again, Members of Class I

Table 4.5: *DCE II: Willingness to accept*

Attribute	Taste Parameters		Reference level
	Class I	Class II	
$ASC_{SchemeA}$	2.13	-1.69	ASC_{Optout}
	(0.27)***	(0.57)***	
$ASC_{SchemeB}$	1.95	-2.05	ASC_{Optout}
	(0.38)***	(0.34)	
Trading Ratio = 10	0.23	0.02	Trading Ratio = 5
	(0.08)***	(0.31)	
Trading Ratio = 20	1.08	0.78	Trading Ratio = 5
	(0.08)***	(0.32)*	
Transaction Fee (%)	-0.07	-0.06	
	(0.01)***	(0.04)*	
Payment	3.43	3.84	
	(0.21)***	(1.15)***	
Summary of class allocation for model: Class 1 (86%) and Class 2 (14%)			
Adj. R^2 vs observed shares: 0.19, BIC: 3041			

are much less likely to choose the opt-out alternative and not engage in trading. Compared to choice experiment 1, there is a greater difference between the classes in terms of current enrolment into real agri-environment schemes, with 60% of Class I participating compared to 40% in Class II. I keep to the naming convention of calling Class I the "high engagement" class, and Class II the "low engagement" class. Differences in educational attainment are once again ambiguous. While low engagement farmers are more likely to state that their highest qualification is some sort of vocational certification, they are also more likely to have attained a post-graduate degree.

Moving now to the second set of choice tasks, where respondents are asked to consider an offer to transfer their NFM obligation to other farmers in exchange for their government NFM payment. In this case, a lower trading ratio for the respondent means a higher ratio for the trading counterparty, for whom the NFM obligation taken on will be proportionally smaller. Table 4.6 shows the latent class results after these respondents have been dropped from the analysis. Consistent

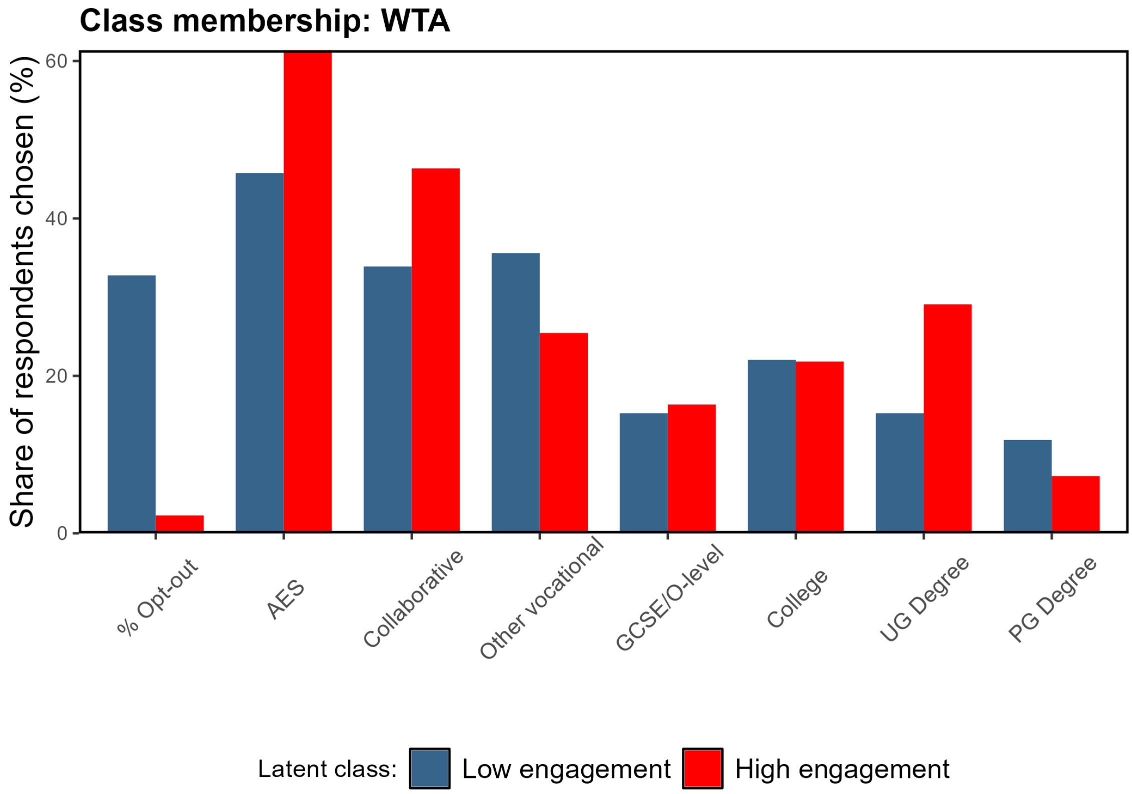

Figure 4.9: *Socio-demographic and behavioural predictors of latent class membership in choice experiment 2: WTA*

with the willingness-to-accept case, Class I displays a preference for engaging in trading, while Class II is indifferent. Also consistent is that Class I displays cost-minimising and transitive preferences (Loomes et al., 1991) for a higher trading ratio, while Class II does not display a significant enough preference for either ratios of 10 or 20 over 5. Both the payment- and transaction fee taste parameters are significant and negative within Class I, in expectation with the theory. Only the payment parameter is statistically significant and negative within Class II.

Test of Hypotheses II and III: I can reject the null for Hypothesis II across the willingness-to-accept and willingness-to-pay scenarios as regards farmers in the high engagement class. There is a consistent and significant preference for higher trading ratios. Across $10,000$ draws, 0% (WTA) and 0.5% (WTP) agree with the null hypothesis. I fail to reject the null as regards the low engagement class. I report evidence in favour of Hypothesis III for the high-engagement class. In both scenarios, high-engagement respondents display a significant preference for lower transaction costs.

Table 4.6: *DCE II: Willingness to pay*

Attribute	Taste Parameters		Reference level
	Class I	Class II	
ASC$_{SchemeA}$	5.31	1.04	ASC$_{Optout}$
	(0.43)***	(0.54)	
ASC$_{SchemeB}$	4.87	1.01	ASC$_{Optout}$
	(0.38)***	(0.54)	
Trading Ratio = 10	0.73	0.09	Trading Ratio = 5
	(0.19)***	(0.30)	
Trading Ratio = 20	1.26	0.44	Trading Ratio = 5
	(0.29)***	(0.34)	
Transaction Fee (%)	-0.12	-0.04	
	(0.02)***	(0.03)*	
Payment	-7.21	-4.71	
	(0.69)***	(1.28)**	
Summary of class allocation for model: Class 1 (76%) and Class 2 (24%)			
Adj. R^2 vs observed shares: 0.19, BIC: 1945			

4.6.3 Monetary cost estimates

Taste parameters for attributes in preference space can be expressed in monetary terms by dividing them by the parameter for the payment or cost attribute. Such transformations invite us to think of the taste parameters in terms of the change in payment required to choose the attribute level over the reference level (Hess & Palma, 2019). In the case of accepting government payment, preference for a particular attribute of the NFM scheme would manifest as a *negative* value because the respondent is willing to accept *lower* compensation if the scheme features the preferred attribute.

Conversely, a positive value means that higher compensation is required to incentivise respondents to choose that option, indicating that it is less attractive. In the choice set testing respondents' willingness to pay to absolve themselves of their NFM obligations, a positive value indicate preference because the respondent is willing to pay more for that option. A negative value on the other hand means that respondents are willing to pay less.

Figure 4.10 shows the taste parameters from the first choice experiment 4.4 expressed in monetary values. The latent class model is also augmented with a mixed logit model allowing for individual-specific preferences. This allows me to visualise the distribution of monetary values across the sample. I also distinguish members in the high engagement class from the low engagement class. In addition to illustrating the smaller size of the low engagement class, it also reaffirms that values in this class are typically clustered closer to zero.

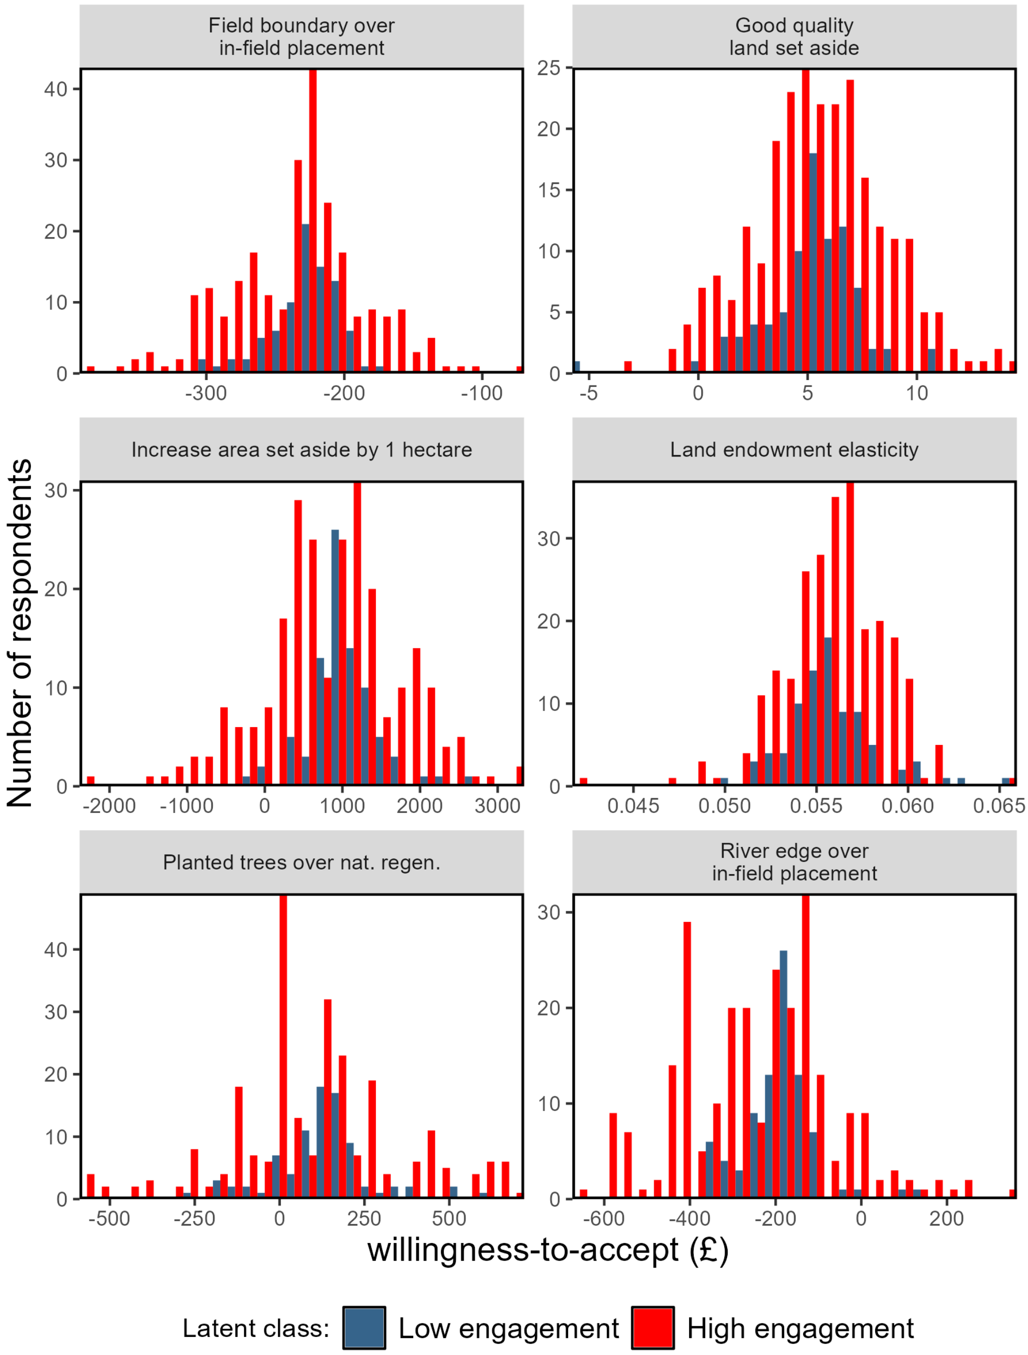

Figure 4.10: *Choice experiment 1: Monetary values for NFM scheme attributes estimated using a mixed logit model*

Respondents in the high engagement class are on average willing to accept ca £200 less per year in compensation if the NFM features can be placed along field- or river edges, instead of on the field. Most respondents are also willing to give up ca £100 per year to create natural regeneration features instead of planted trees. Respondents in both classes typically demand in the region of £500-£1500 more per year to retire an additional hectare of land. The estimated land endowment elasticity is distributed around 0.055 and not significantly different from zero, which indicates that the input factor productivity of agricultural land is close to 1.

Figure 4.11 shows the taste parameters from the willingness-to-accept scenario expressed in monetary values. The payment attribute is valued in £/hectare, and taste parameters therefore represent the change in payment per hectare required to choose the current level over the reference level. Respondents are willing-to-accept on average around £65-£75 less per hectare per year to be offered a trading ratio of 10 over a ratio of 5. They are willing to accept on average £300-£400 less for a trading ratio of 20 than a ratio of 5.

In terms of land value, a trading ratio of 5 means respondents would need to create an additional $^1/_5$ hectares of NFM to get the full payment. A ratio of 10 means an additional $^1/_{10}$ hectares, and a ratio of 20 an additional $^1/_{20}$. This means that the choosing the ratio of 20 over the ratio of 10 rewards a 475% reduction in the amount of land set aside. Conversely, the reduction in willingness-to-accept is more than 600%. The effect on willingness-to-pay from increases in the trading ratio is more than proportional. This suggests that there are few barriers to a functional trading market from farmers on high runoff potential land who would be facing high trading ratios. One remaining barrier may be transaction costs, depending on how the program is designed. An increase in the transaction fee by one percentage point of the per hectare payment increases the required payment by ca £20.

Figure 4.12 shows the effects of trading scheme features on respondents' willingness-to-pay to opt out of their NFM obligations. Preferences for a higher trading ratio are lower than in the willingness-to-accept scenario. Only ca 50 respondents are willing to pay more than £100 more per year to have the ratio of 20 instead of the ratio of 4. The the willingness-to-pay is ca £15 lower per percentage point increase in the transaction cost. The magnitude of this effect is approximately £5 per percentage point lower than for the WTA scenario. At the reference level for the trading ratio, the perceived value of opting into the scheme is greater in the WTP scenario than in the WTA scenario, although the value is positive in each case. This means that a meaningful transaction cost may further magnify the divide between the two sides of the hypothetical market.

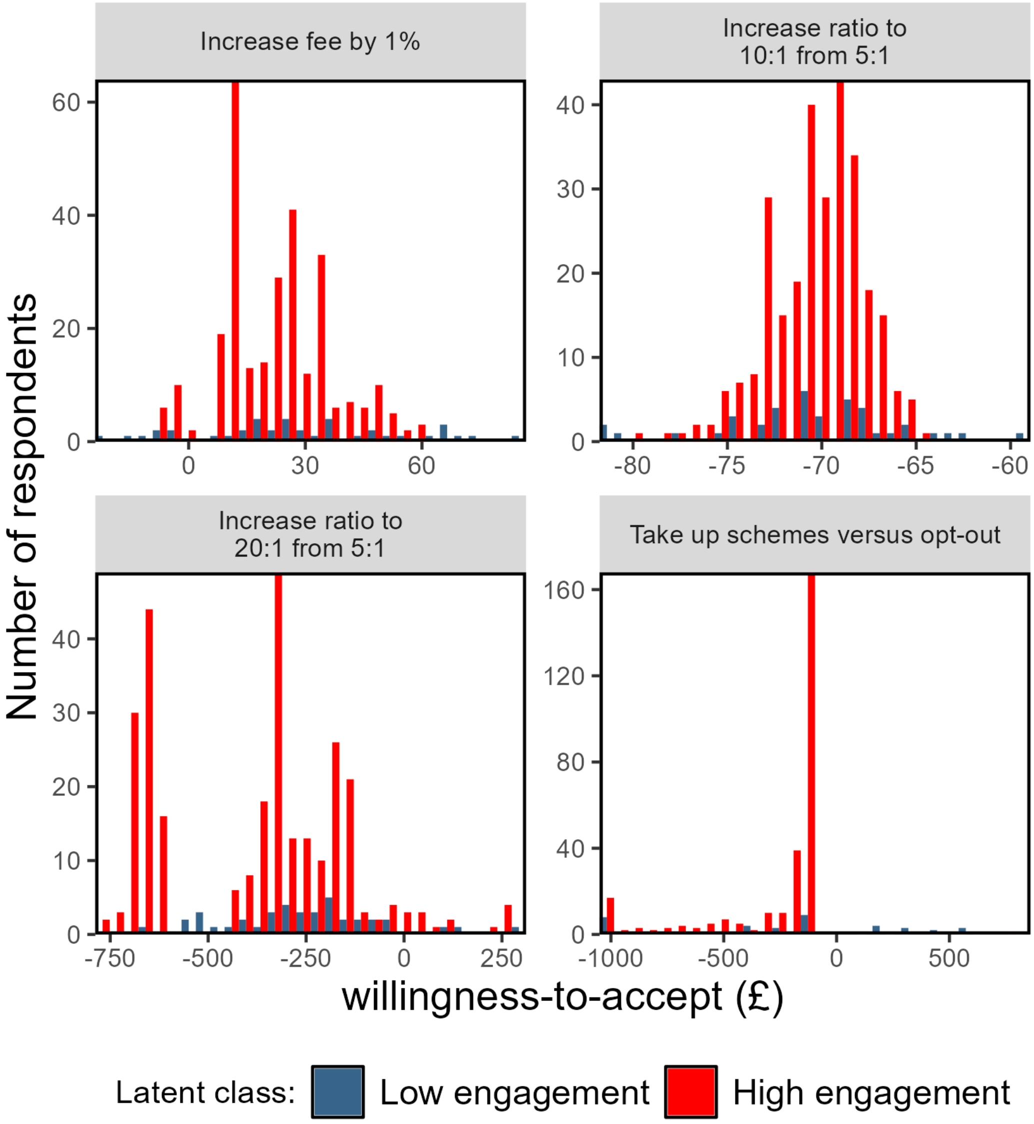

Figure 4.11: *Choice experiment 2a: Individual monetary values for NFM trading program (willingness-to-accept)*

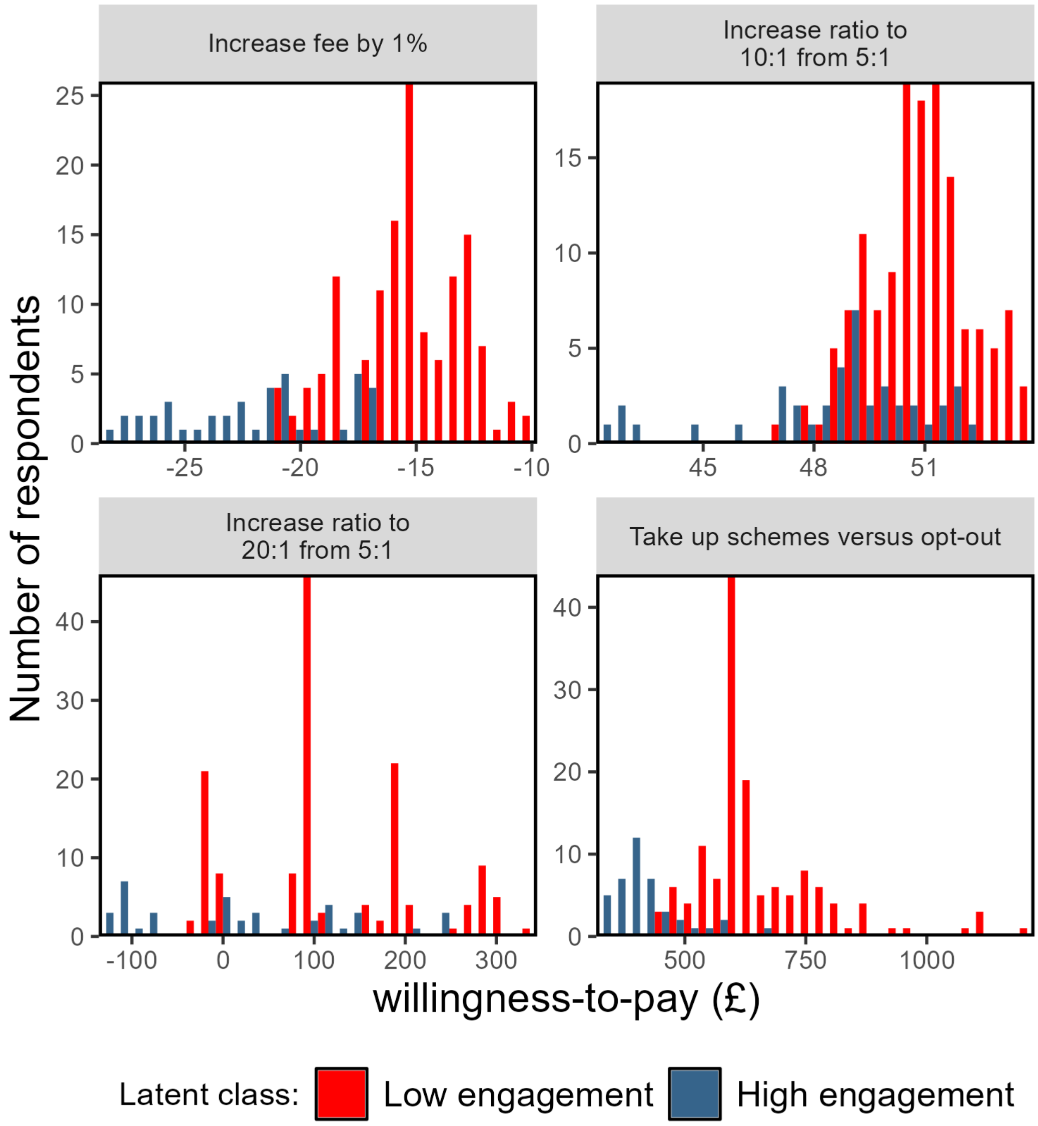

Figure 4.12: *Choice experiment 2b: Individual monetary values for NFM trading program (willingness-to-pay)*

4.6.4 Cost-effectiveness analysis of payments for NFM with a spatially targeted trading program

I use `SCIMAP-Flood` (Reaney, 2022) to identify two 10-by-10 kilometre samples of predominantly agricultural land from the Eden catchment. The two samples and their respective distribution of runoff generation risk scores are shown in figure 4.13. The average area-wide risk score for the high-risk sample is 0.11 and the average area-wide risk score for the low-risk sample is 0.03. This gives two hypothetical farm located in the high- and low-risk areas respectively a trading ratio between them of only 3.63. In practice, most farms in the study are significantly smaller then these sample areas, which allows for higher trading ratios. For example, the top 10% of the high-risk area (1000 hectares) has an average risk score of 0.2 while the bottom 10% of the low-risk area has an average score of 0.01. The trading ratios between these segments would be 20, which is the upper limit in the choice experiments. Runoff risk hotspots are typically clustered together as shown in figure 4.13 which makes ratios in the 5-20 range realistic between actual pairs of farms.

On each of these samples of geography, I simulate the two types of NFM features in four spatial configurations introduced in section 4.5. These include planted broadleaf trees and natural regeneration, arranged in a contiguous patch covering both active- and inactive farmland, in-field corridors, field-edge corridors, and in-field islands. Corridors and islands also come in widths of 10 and 20 meters. Benefits from the schemes are defined as the overall reduction in runoff generation risk scores per square meter of NFM features created.

The risk reduction can be linearly scaled up for larger features (Reaney, 2022). The effect in the low-risk and in the high-risk area are displayed in figure 4.14. I find

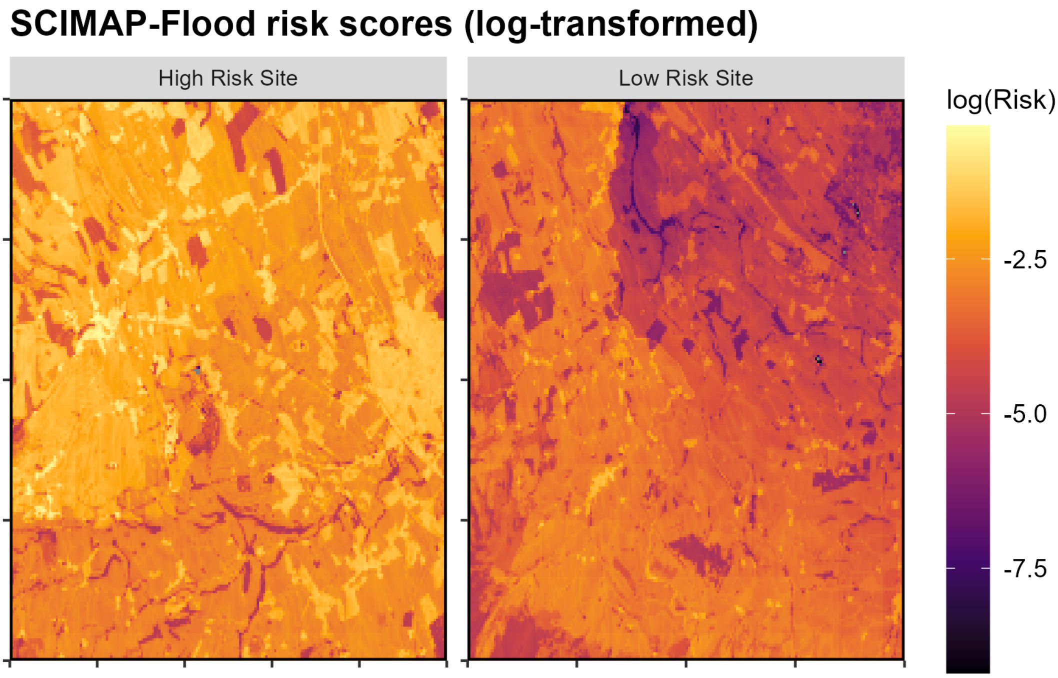

Figure 4.13: *Geographic distribution of runoff generation risk produced by* `SCIMAP-Flood` *(log-transformed) for two 10x10 kilometer sites in the Eden catchment, North West England.*

that the in-field islands consistently deliver the greatest benefit per area of NFM produced, followed by the singular, contiguous patch. This is within expectations, as compared to corridors, islands require far less land to be set aside for NFM at any given NFM intensity, expressed as the gap between the corridors/isles. In-field corridors are marginally the second most efficient option. This is also directionally within my expectations, as surface roughness and soil penetration are typically poorer on the field compared to field boundaries where more diverse vegetation may already contribute to runoff reduction.

Planted trees are approximately 40% more efficient than natural regeneration for in-field islands and 20% more efficient for other spatial configurations. These effects are directionally expected as trees contribute to prevent soil erosion and to absorption capacity (Weninger et al., 2021).

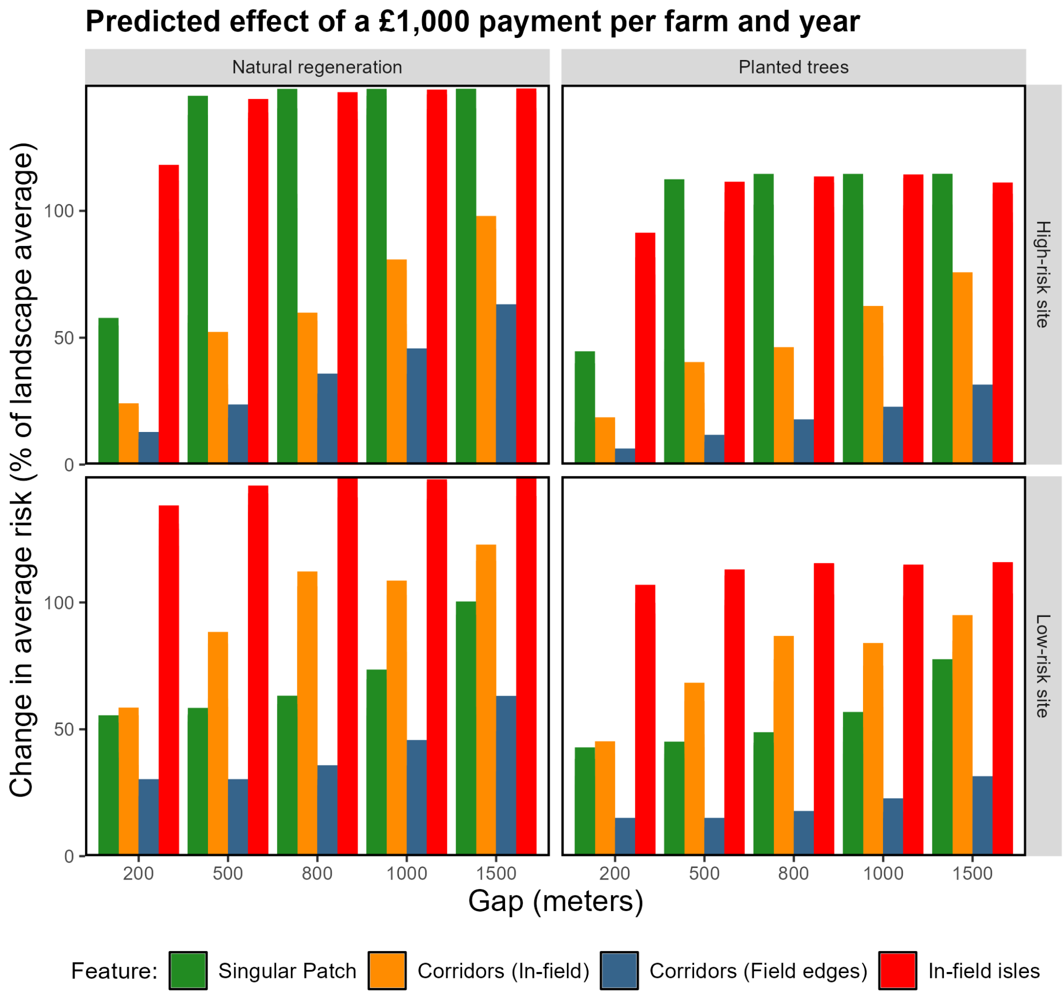

Figure 4.14: *High-risk (upper) and low-risk (lower) area-wide mean reduction in runoff risk per m^2 of NFM created, by feature type, spatial configuration, and NFM intensity*

I now move on to simulate trading between the two areas. Figure 4.15 shows for each NFM scheme how the resulting reduction in runoff risk changes as the NFM obligation (0.5% of productive farmland) is transferred from the low-risk area to the high-risk area. The white circles represent the post-NFM reduction in runoff risk for each NFM scheme before trading. The coloured circles represent the post-NFM reduction after trading. In this case, the low-risk farm does not maintain any NFM features, and the high-risk farm maintains their original obligation (0.5%) plus $^1/_{3.6}$ of the low-risk farm's obligation. The risk reduction per m^2 improves because the total amount of NFM created is only ca 65% of the amount without trading. By the same logic, in-field islands benefit the most, because a smaller total amount of land devoted to this configuration can be spread out across a larger area.

I now incorporate monetary values from the choice experiments to appropriately do a costing of each of these hypothetical schemes. I know from DCE I that farmers value a one hectare increase in the amount of NFM created at approximately £1000 per year. I also know that farmers perceive planted trees as ca £100 more expensive per $^1/_{20}$ hectare per year, compared to natural regeneration. Finally, I recall that the stated cost of in-field placement of features is approximately £200 per $^1/_{20}$ hectare in excess of the cost for field-edge features. I add these differences in cost to a baseline payment of £500 per $^1/_{20}$ hectare per year.

Figure 4.16 illustrates the effect on the total government spending required annually to maintain a NFM obligation of 0.5% of productive land for each farm. The white circles represent the required spending before a market for obligations in introduced and the coloured circles represent the required spending when trading is allowed. I also incorporate the additional costs associated with facilitating the trade, in the form a a transaction fee. For natural regeneration, trading represents a cost saving of ca £30,000 per year for a contiguous patch, in-field corridors and in-field islands.

Trading represents a £10-15,000 cost saving for field-edge corridors. The cost savings are consistently larger for a planted tree scheme than for natural regeneration. This is because a more expensive scheme (planted trees) benefits more as a result of the greater land use efficiency from trading. Across the board, a higher transaction cost reduce the cost savings achievable from trading, although even at 10% of the payment trading reduces the overall cost by no less than £10,000.

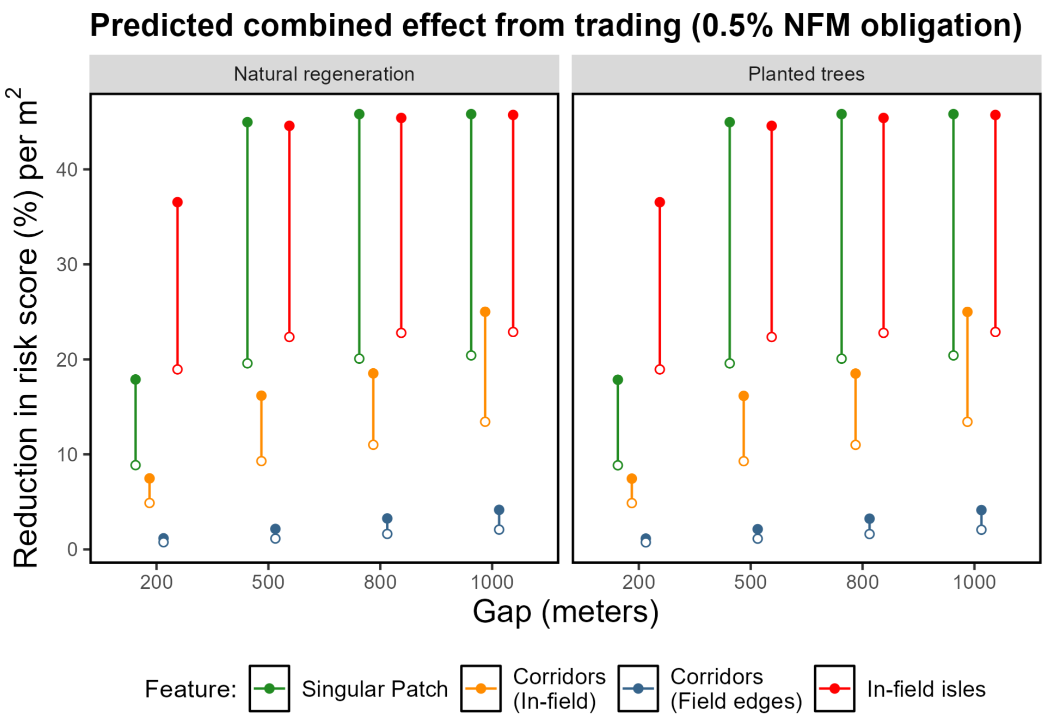

Figure 4.15: *Reduction in runoff risk per m^2 of NFM created, by feature type, spatial configuration, and NFM intensity. White circles represent the reduction without trading. Coloured circles represent the reduction with trading.*

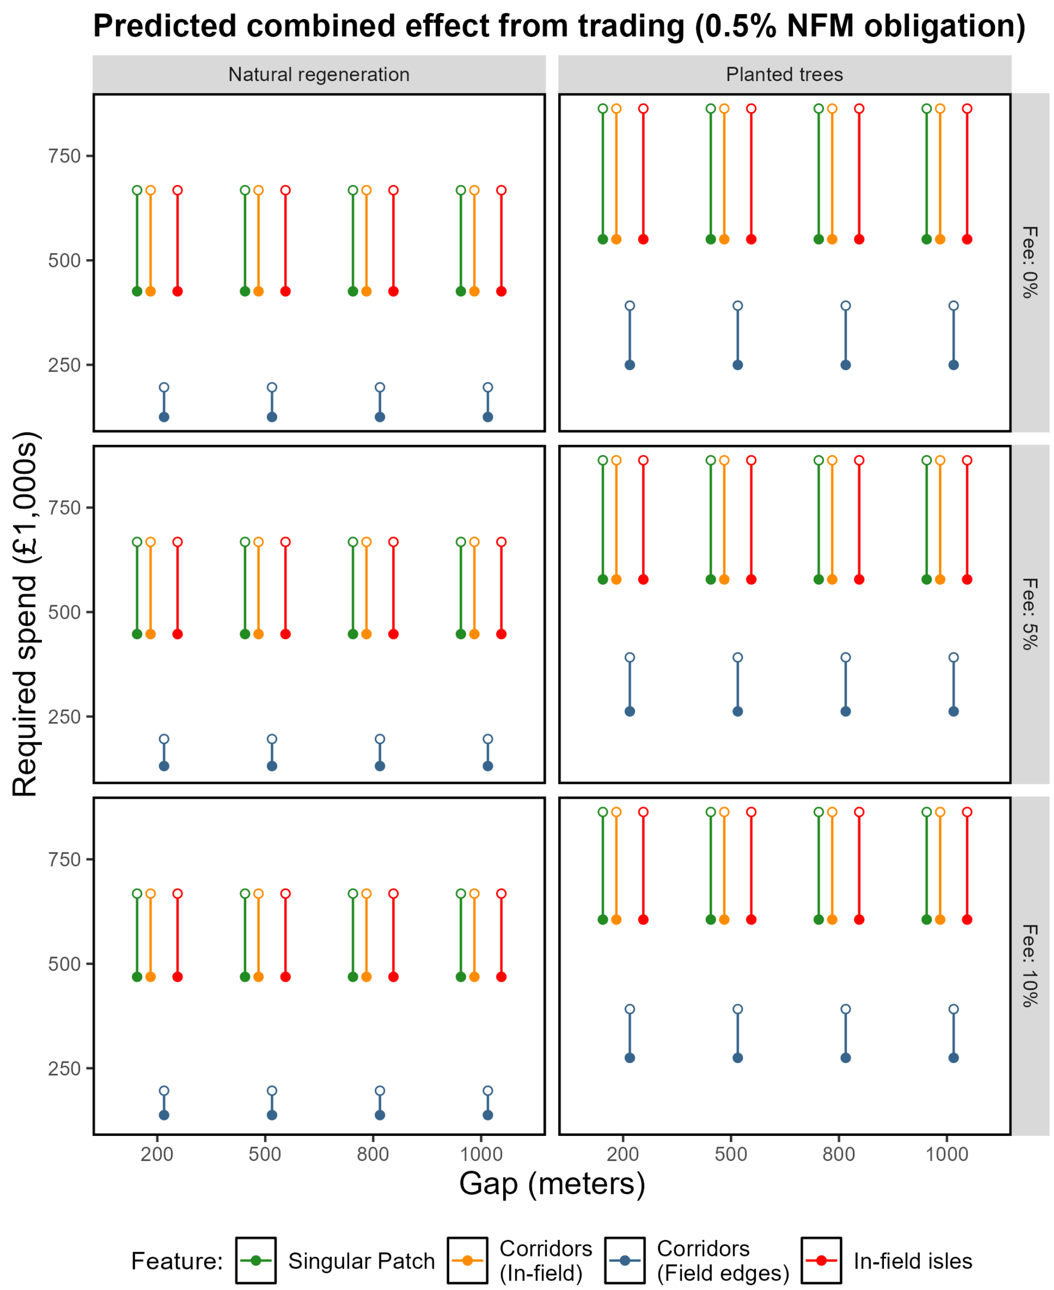

Figure 4.16: *Required government spending to ensure 0.5% NFM obligations, by feature type, spatial configuration, and NFM intensity. White circles represent the cost without trading. Coloured circles represent the cost with trading.*

4.7 Limitations

A key assumption underlying discrete choice modelling is that individuals have well-formed preferences that are stable over time. In practice, even repeated choices in the very near term such as the repeated choice tasks in the online survey, frequently display some degree of preference instability. This is a particular risk when the choice tasks are cognitively taxing for respondents, for example due to complexity or unfamiliarity (Hess et al., 2012). Figures 4.17 and 4.18 show preference stability for higher trading ratios across 12 choice tasks, six in the willingness-to-accept scenario and six in the willingness-to-pay scenario. With reference to the first choice task in each scenario, the figures display the proportion of respondents who maintain their preference in each subsequent choice task. The proportions are not cumulative, i.e. a respondent may choose inconsistently in the second choice task and return to their initial preference in the third. On the horizontal axes are shown the choice task as well as the payment trade-off for a higher trading ratio. A negative number means that the high trading ratio option is less attractive financially, and a positive number means that it is more attractive. Note for example task 2 (+£40) in the willingness-to-accept scenario and task 3 (+£20) in the willingness-to-pay scenario. In these cases, the option with the higher trading ratio is strictly dominant as it is also more advantageous financially. Nevertheless, only ca 70-75% of respondents maintain their preference for higher trading ratios from task 1.

Furthermore, 95 respondents chose the strictly dominated alternative in these two tasks, featuring a worse trading ratio, a higher transaction fee (paid by the respondent), and a higher payment. I hypothesise that this was a result of ambiguous wording on the choice card. In particular, the description of the payment attribute specified the £/ha amount from the trading counterparty's point of view. In other words, a higher trading ratio chosen by the respondent (e.g. 5 over 10) allows the counterparty to set aside less land for the same payment, attributing a higher £/ha value to the land. Indeed, their land *is* more valuable in terms of flood risk reduction. I suggest that respondents may have chosen irrationally if they perceived this higher £/ha value as a cost to them. These respondents display negative tastes for more favourable trading ratios. Any future attempt at replicating my results should express land values in terms of hectares per pound sterling: With a lower trading ratio, respondents in the willingness-to-pay experiment encourage their counterparty to set aside more land for NFM with the same payment.

More encouraging is that the rational respondents can be consistently distinguished from the irrational ones throughout the choice tasks. Looking at the comparison between rational and irrational respondents in figure 4.18, I observe that the rational group responds more clearly to incentives. In task 2 of the willingness-to-accept scenario with a comparatively large financial payoff from choosing the high-ratio alternative, the rational group (blue) is more likely to choose that option. Conversely in task 4, which imposes a steep payment penalty from choosing the high-ratio option, the rational group is less likely to choose it than is the irrational group. The same dynamic is observed in the willingness-to-pay scenario.

Two perspectives on the source of preference instability are commonly presented: Discovered versus constructed preferences (Matthews et al., 2017). The former hypothesises that when people have to make decisions about an unfamiliar issue or in an unfamiliar environment, their initial responses may be impulsive. As they learn about the decision environment (institutional learning) and their own attitudes (value learning), their decisions begin to exhibit less randomness and greater rationality. The latter posits that when faced with an unfamiliar or ambiguous choice, respondents may try to construct their preference on the spot, which may lead to instability. Study of the stability patterns in figures 4.17 and 4.18 reveals that the correlation between preference for high-ratio alternatives and positive financial payoffs from choosing the high-ratio alternative does not improve appreciatively over a series of choice tasks. This rejects the idea of significant learning over time. Instead, the trading scenario and the interpretation of the trading ratios would be rather abstract to respondents. This points towards a preference constructed with some confusion. In particular as I acknowledge that the choice cards could be worded more accurately in the willingness-to-pay case.

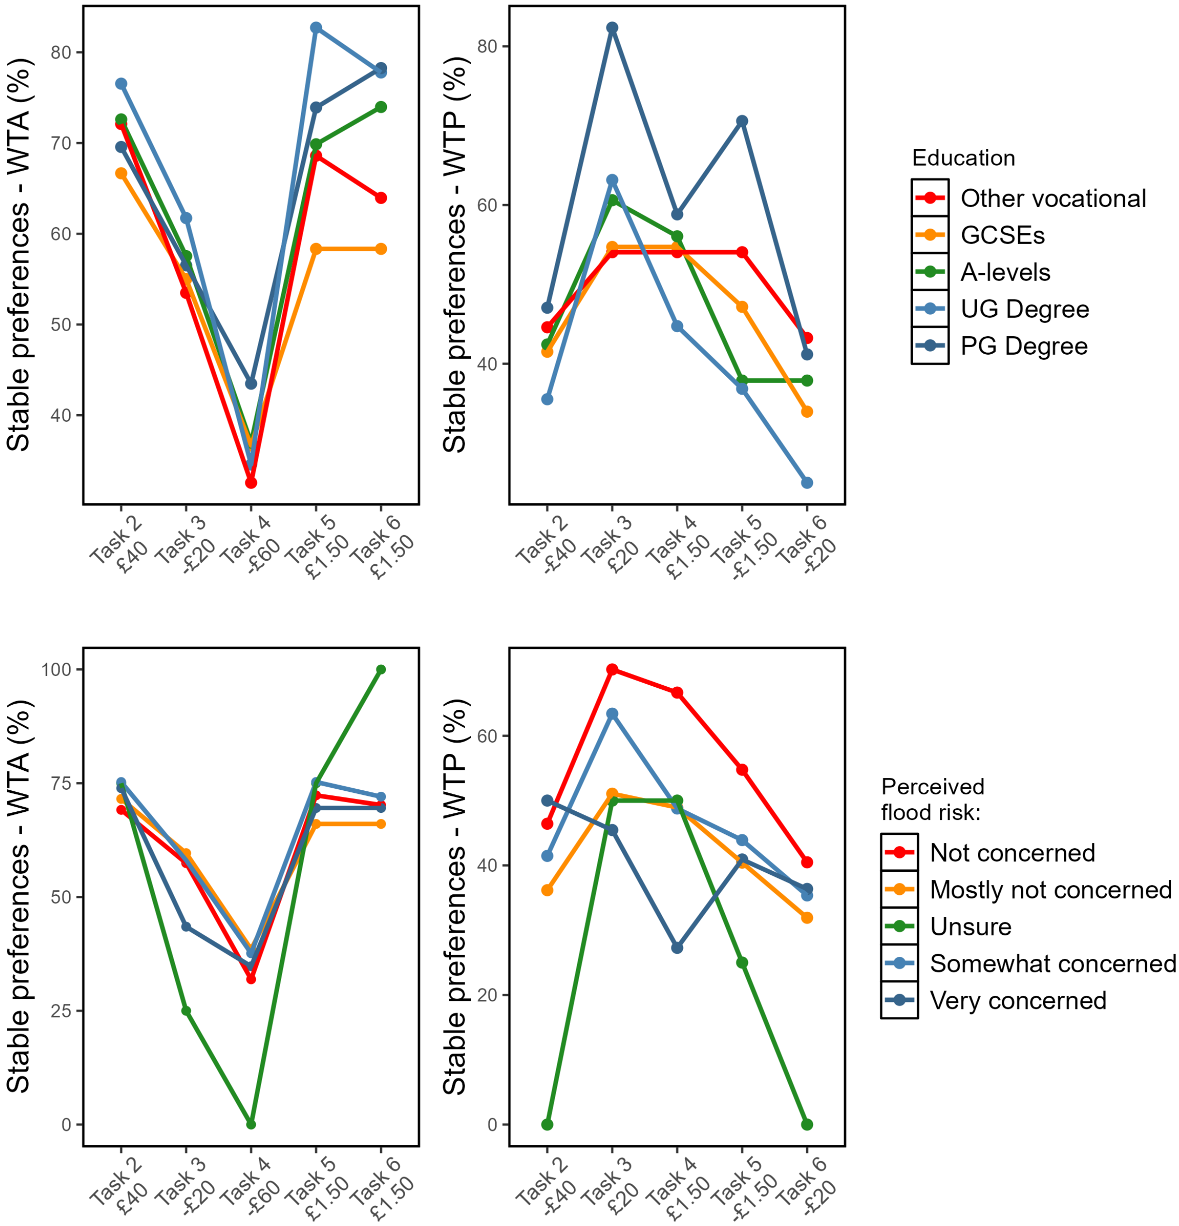

Figure 4.17: *Preference stability for higher trading ratios over six sequential choice tasks across two choice experiments, broken down by educational attainment and stated concern about flood risk.*

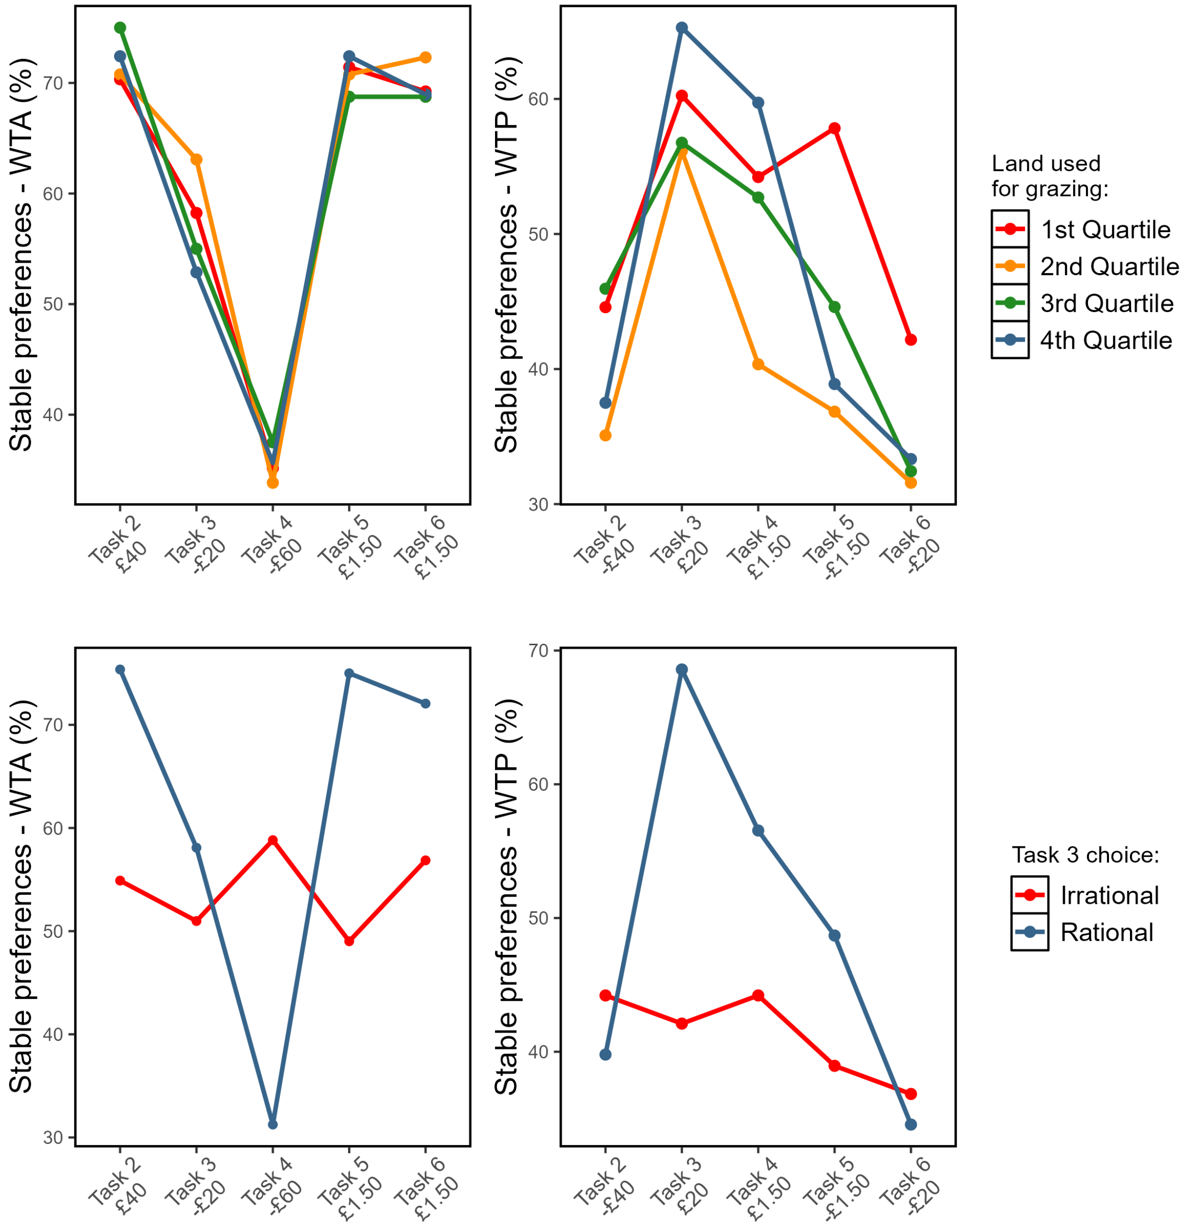

Figure 4.18: *Preference stability for higher trading ratios over six sequential choice tasks across two choice experiments, broken down by land used for grazing and irrational choice of dominated alternative.*

4.8 Discussion

Previous research has shown in different settings (including air- and water pollution) that producers of environmental externalities that are geographically removed from the source respond less to abatement incentives. It has been proposed that such geographically high risk polluters face less regulatory pressure as environmental damage may occur outside the jurisdiction in which they operate, and that such firms should face greater incentives to abate. Theoretical work has explored the application of trading ratios to traditional cap-and-trade programs as a way to better target incentives towards firms that are deemed geographically high risk (Fowlie & Muller, 2019; Stranlund & Chavez, 2000). Runoff generation from agricultural land use is a prime example of an environmental externality where flood damages may occur away from the farms. This research shows that, in general, there is a strong stated willingness among English farmers to enrol in NFM schemes providing payment for creating natural flood management. Participants in the DCE also express a willingness to take over the NFM obligation of another farms at a favourable trading ratio.

However, the research reveals notable differences between two groups that I call high- and low engagement respondents. The low engagement group is less likely to want to enrol in the schemes and is less likely to find the spatially targeted trading program attractive. This group is characterised by a much greater propensity to opt out of the schemes altogether, and is moderately less likely to currently be enrolled in a real-life environmental land management scheme. These quantitative results add to earlier qualitative work by e.g. Holstead et al. (2017) and Kenyon (2007).

Empirical research into the economic efficiency of markets in tradable pollution permits (Schmalensee & Stavins, 2013, 2017) has downplayed the concern about transaction costs highlighted in theoretical models (Xepapadeas et al., 1997). However, it has been an open question whether transaction costs presents a barrier to trade in ELM contracts between farmers Nguyen et al., 2022. Evidence from behavioural lab experiments suggests that they do (Banerjee et al., 2017). This work contributes novel evidence on this question by estimating active farmers' willingness to trade NFM obligations when transaction costs exist. Results show that the required base payment needs to increase by ca £15, per hectare per year, for each percentage point increase in the transaction cost. For context, given the range of per-hectare payments in the DCEs, the shift in WTA is approximately equal to the shift in transaction costs. This means that the regulator has limited scope to pass on the transaction costs to farmers. As a result, efforts to remove trading barriers should feature in an attempt to open a spatially targeted market in NFM contracts. For example, output from models of flood risk (such as `SCIMAP-Flood` presented in this work) can be used in an optimisation algorithm to match pairs of farms where the gradient between risk scores is maximised. The relevant regulator (such as Defra) may host an online platform with the matching algorithm running on the back-end. In this way, farmers who sign up can be matched with the most profitable trades.

This research finds significant environmental benefits from trading, both in terms of flood risk reduction and in terms of government spending. In particular, small disconnected islands of retired land with reduced runoff potential on managed fields is a comparatively cost-effective intervention. These results only further highlight the usefulness in incentivising a spatially targeted system. Transaction costs in the form of trading fees do reduce the cost savings from trading, at most by approximately 10%. This is a result of farmers demanding a higher base payment rate to offset the impact of transaction costs. Limited information about prices and potential buyers and sellers are typical sources of transaction costs in other permit markets (Tietenberg, 1990). Defra could increase transparency by providing a digital platform in a similar vein as how it communicates land parcels' eligibility for other ELM projects.

Finally, this work provides new lessons for DCE practitioners on the issue of preference instability. I find that certain respondents make irrational choices, possibly because they have misunderstood the choice task, which may have been poorly presented. These respondents can be clearly identified by tracking stability of their preferences over repeated choice tasks. Results from a hypothetical DCE, featuring relatively abstract schemes for trading NFM contracts, display clear differences in preference stability between two groups of farmers. Those who chose the cost-minimising option, in a choice task with a fully dominant scheme, respond rationally to changes in payoffs through six repeated choices. Farmers who did not choose the dominant option also do not adapt as expected to changes in payoffs. In contrast, differences in preference stability by educational attainment and awareness about flood risk were much less pronounced. Research on complicated schemes using hypothetical DCEs may therefore consider introducing a choice task with a dominant option to identify this group and to discuss their choices separately. This procedure offers an ex-post way to identify respondents who, despite efforts by the survey designer, have misunderstood one or several choice attributes. This is particularly relevant when hypothetical DCEs are used to estimate preferences for policies or products that do not yet exist.

Chapter 5

Voluntary spatial targeting in the presence of coordination costs

5.1 Introduction

As well as managing negative externalities (pollution, flooding) environmental land management can produce positive ones. Protecting local ecosystems by planting trees, hedgerows and flower strips contributes to what Costanza et al. (1997) call ecosystem goods and services. Economists would soon discuss the value of ecosystem services like climate regulation, nutrient recycling and pollination. As recognised by Heal (2000), although biodiversity and associated services may seem intuitively valuable and important, their market value is more ambiguous. Key uncertainties relate to the indirect use value of an ecosystem (Nijkamp et al., 2008) where it supports marketed natural resources, such as agricultural yields. Due to the complexity of ecological systems, such values are not obvious but scenarios in Kubiszewski et al. (2020) attribute changes in land management alone to a difference of $81 trillion by 2050.

Pollination is one of the most intensely studied ecosystem services due to its link with global food production (Hanley & Perrings, 2019), with Porto et al. (2020) estimating that US$155 million of research funding had been contributed by 2018. In a literature review, Klein et al. (2007) show that pollinators impact food supply globally, as pollinator-dependent crops contribute to 35% of overall crop production by volume. It is estimated that 87 of the 115 major crops grown worldwide depend on biotic pollination to set fruits and seeds to at least some degree. Globally, the economic value of pollination is estimated at US$195-387 billion (Porto et al., 2020). Pollination is essential for farming apples, cacao and vanilla and of great importance for buckwheat, pears, and berries Klein et al. (2007). The use of animal pollinated biofuel crops is growing, with the cultivation area of oilseed rape, sunflowers and soybeans increasing by 32% across Europe between 2005 and 2010 (Breeze et al., 2015), with ca $320,000$ hectares used for these crops in the UK

(Thompson, 2022). In total, pollination in the UK is valued between £189 million [1] (Breeze et al., 2021) and £379 million (Breeze et al., 2015) per year.

Powney et al. (2021) and Potts et al. (2016) show a reduction of wild pollinator populations at the regional level, especially within Europe and North America. Recent research suggests that the occupancy[2] of bee and hoverfly species has declined by an average of 25% across the UK since 1980 (Powney et al., 2019, 2021). A comparative study of European honeybee colonies showed that while there were honeybee deficits (insufficient stocks to supply 90% of national demands) in 22 countries in 2010, only the UK and Moldova had a pollinator stock capacity below 25% (Breeze et al., 2014).

The causes of pollinator decline include the indiscriminate use of pesticides, biological invasions, genetically modified (GM) crops, intensification and expansion of agricultural practices (Dicks et al., 2016; Potts et al., 2016), as well as habitat loss and fragmentation associated with farming and urbanisation (Donkersley et al., 2014; Potts et al., 2010; Xiao et al., 2016). Properly targeted environmental land management (ELM) schemes provide measurable improvement in fragmented landscapes (Donald & Evans, 2006). Understanding how land management affects pollinator abundance and diversity in combination with other drivers is necessary to design more targeted, adaptive management strategies at national scales (Halinski et al., 2020; Lucas et al., 2017).

While a developing literature is studying the targeting of ELM projects to achieve optimal pollination benefits (Halinski et al., 2020; Häussler et al., 2017; Image et

[1] This estimate is based on the market value of crops lost under a 30% reduction in insect pollination

[2] Occupancy rates are the proportion of occupied 1km grid squares each year based on presence-absence models

al., 2023), recognition that habitat connectivity is a driver of pollination (Jauker et al., 2013) calls for collaboration between farmers (Krämer & Wätzold, 2018). Meanwhile, work on agglomeration bonus payments do not typically treat pollinator dependence as a differentiator between land managers (Banerjee et al., 2014; Kuhfuss et al., 2016). A recent literature review of 55 studies finds only six empirical valuations of coordination bonuses in ELM schemes (Nguyen et al., 2022) and the few examples that incorporate pollination in the production function do not study cooperative equilibria (Kleftodimos et al., 2021). I fill this research gap by modelling a mixed agricultural catchment where creation of natural features may enhance productivity among pollinator-dependent farms. The rest of the chapter is structured as follows: First, I establish the current state of knowledge around landscape fragmentation and its impact on the economic value of pollination. Second, I apply for the first time a spatially explicit model of pollinator visitation to validate an agricultural production function incorporating pollinator dependence, and identify benefits from connectivity improvements. I explore whether pollination services can be enhanced via coordination between farmers to achieve optimal connectivity improvements. Third, I test the model's prediction that variation in coordination costs predict connectivity improvements using a discrete choice experiment with English farmers. Finally, I discuss the results in context of ongoing revisions to UK ELM schemes and their implications for policy making.

5.2 Background literature

Insect pollination is a well-studied ecosystem service that supports production in 75% of globally important crops (Klein et al., 2007). Insects visit flowering crops to forage for nectar and pollen, that is used for food. When moving from flower to flower, they fertilize the plant by depositing pollen stuck to their bodies (Lucas et al., 2017). Insects known to benefit crops grown for human consumption are: honeybees; sting-less bees; bumble bees; solitary bees; wasps; hover flies and other flies, and beetles (Klein et al., 2007). Of these, honeybees are the most important to agriculture. To date, the most comprehensive review of pollinator dependence for different crops was done by Klein et al. (2007) who designated insects essential to 13 out of 75 crops, with another 30 classed ”highly dependent”. Figure 5.1 shows dependence ratios for important crops in the UK agriculture industry defined as the proportion of yield lost in the absence of pollination (Breeze et al., 2021).

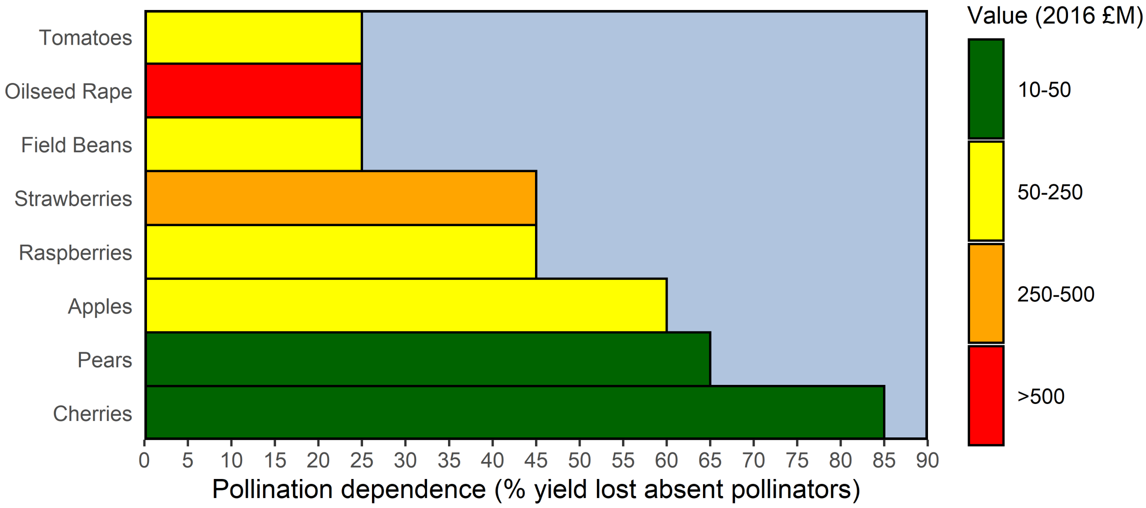

Figure 5.1: *Estimates of pollination dependence and crop values in the UK from a 2014-2016 survey by Breeze et al. (2021).*

As shown in figure 4.1, pollinator dependence as well as economic values vary across crops. Only 25% of oilseed yield is at risk from pollinator decline, but ecological studies at experimental fields suggest that restriction of insect visits results in yields $27 - 30\%$ lower (Stanley et al., 2013). Additionally, total UK production was valued at £662M annually in 2016 (Breeze et al., 2021), pricing pollination benefits over £150M per year. In general, the use of pollinated biofuel crops has grown, with the cultivation area of oilseed rape, sunflowers and soybeans increasing by 4.2 million hectares (32%) across Europe between 2005 and 2010 (Breeze et al., 2014).

5.2.1 Economic value of pollination

The magnitude of the threat from pollinator decline to rural economies and food security has been actively debated (Ghazoul, 2005a, 2005b; Steffan-Dewenter et al., 2005). Although in large commercial systems, stable service through the flowering period is often ensured by beekeepers, pollination services are primarily provided by wild insect communities (Breeze et al., 2021). Increased dependence on communities managed by beekeepers would also add costs for farmers. While Kleijn et al. (2015) find that most crop pollination is provided by a limited set of non-endangered species, recent accounting in Powney et al. (2021) shows that overall UK pollinator communities remain at 60% of their 1980 baseline, with no sign of recovery. Further, honeybee declines in the UK are largely a post-2007 phenomenon, qualifying certain sources in Kleijn et al. (2015) as either internationally oriented (Winfree et al., 2008) or premature (Kleijn et al., 2006). Additionally, while industrial farming contributes the majority of agricultural value added, the pollination benefit to yields is typically more pronounced on small farms more sensitive to downside risk (Garibaldi et al., 2016). Nonetheless, the complicated and obscure nature of pollinator benefits to crop yield has meant limited interest in conservation among farmers (Ghazoul, 2005a). Acceptable schemes need to be nonintrusive in day-to-day farm management. There is a need to achieve the greatest provision of ecosystem services (including other environmental goods such as flood management, proposed in e.g. Forbes et al. (2015)) at the lowest disruption to agricultural land.

Crucially, an emerging literature seeks to clarify the relationship between habitat connectivity and pollination services. Habitat connectivity - distinct from habitat *area* - refers to the degree to which individual pollinators can easily traverse a landscape to find mates and food (Xiao et al., 2016). For winged insects like bees, butterflies and hover flies, this primarily means access to land parcels of sufficient feeding quality within a certain foraging distance of a suitable nesting site (Lepais et al., 2010). Foraging distances vary quite significantly between species and range from 200 to 600 meters among important European pollinators (Häussler et al., 2017).

Highly connected habitats suffer from few obstructions like large, continuous areas of non-flowering crops (De Palma et al., 2015), intensive grazing (Le Féon et al., 2013), or water that may result in lower pollinator abundance. In an experimental study set in grazed grasslands and intensely farmed landscapes, Steffan-Dewenter and Tscharntke (1999) found that increasing isolation of small islands of habitats resulted in decreased pollinator abundance of bees. Overall, the academic consensus is that connectivity is causally linked to pollination by wild insects (Senapathi et al., 2017).

Habitat fragmentation is specifically recognised as a threat to ecosystem service provision in agricultural landscapes (Montoya et al., 2021). However, the impact on crop yields from fragmentation is complex. For example, strategically placed but disconnected patches of trees can support pollination by increasing provision of flower-rich grove edges (Halinski et al., 2020; Ren et al., 2023). The spatial coexistence of crops and natural land can also create spillover effects for provision of ecosystem services broadly defined. Despite a rapidly growing literature, concludes Montoya et al. (2021), our understanding of these interactions remains incomplete.

The consensus in the recent ecological literature is that ELM schemes aimed at supporting provision of pollination services should be spatially targeted and focused on interventions empirically proven to be effective, including hedgerows (Timberlake et al., 2019), planted trees (Halinski et al., 2020), and seminatural grassland management (Berg et al., 2019). Spatially explicit models of pollinator visitation, with high-resolution land cover data and parametrized to fit a representative landscape, are used (Image et al., 2023) to deal with the complexity of pollination provision highlighted in Montoya et al. (2021).

I collect land use and crop cover data around 495 surveyed farms in the north of England. I simulate counterfactual landscapes by altering these data with ELM features. I run the simulated landscapes through one such model to disentangle the effect of connectivity from that of habitat size. The model, `poll4pop`, is a probabilistic model of abundance and visitation rates previously applied to English farmland in Image et al. (2023). I focus on connectivity as it presents a possible channel to improve pollination without larger sacrifices of productive land (Image et al., 2023). These simulations help to specify the functional form of the connectivity-visitation and feature size-visitation relationships. This allows me to propose policy designs that optimise the cost-effectiveness of programs aimed at improving pollination services.

5.2.2 Spatial models of pollination

Open-source model `poll4pop` (Gardner et al., 2020; Häussler et al., 2017) is a spatially explicit model predicting pollinator visitation rates. Häussler et al. (2017) use the model to estimate the effect of establishing grassy field margins offering nesting resources and a low quantity of flower resources, and/or late-flowering flower strips offering no nesting resources but abundant flowers, on visitation rates to flowers in landscapes that differ in amounts of linear seminatural habitats and early mass-flowering crops. `poll4pop` adds to earlier models (Lonsdorf et al., 2009; Zulian et al., 2013) by (1) integrating preferential use of more rewarding floral and nesting resources; (2) considering population growth; (3) allowing for different movement distances for foraging and queen dispersal (Lepais et al., 2010); and (4) producing spatially explicit flower visitation rates.

The model is parametrised based on a survey of the literature by Häussler et al. (2017) on pollinator dispersion. As a result, certain parameters are "best guesses" about a pollinator species' nesting requirements, mean foraging distance, and survival rates. These estimates are shown in table 5.1. In addition, the model uses land cover rasters to represent the agricultural landscape.

Table 5.1: *poll4pop parameters*

Parameter	Description	Value	Source
n_{max}	# nests of max nesting quality	19.6 nests/ha	Osborne et al. (2008)
w_{max}	Max # workers per nest	600	Häussler et al. (2017)
p_w	% foraging workers	50%	Brian (1952)

Across a landscape of 10-by-10 meter parcels, land use classes are scored according to the flower resources they provide. Each land class is scored according to the amount of floral cover it provides, the attractiveness of those floral resources to each pollinator guild (floral attractiveness) and the attractiveness of the nesting opportunities the land class provides to each pollinator guild (Gardner et al., 2020). From Häussler et al. (2017), table 5.2 shows floral cover and attractiveness for a sample of land use classes. Floral cover was defined as the proportion area covered by flowers, between 0 and 1, and varies by land cover class. Floral attractiveness is defined as a score ranging from zero (not at all attractive, never used) to 20 (very attractive, preferred over other flowers).

Floral cover is multiplied by floral attractiveness to obtain the species-specific floral value scores. Nesting quality was defined as a score ranging from zero (totally unsuitable) to one (very suitable). As with floral resources, the expected number of nests per cell is defined as the empirical maximum times the nesting quality of the cell. For each landscape, the raster of floral resources for period F is the product of cell-specific floral coverage, expressed as the proportion of area covered by flowering plants, and a score of the species-specific attractiveness of the typical flowers in a land-use category. From here, visitation of bees from cell j to any other cell i is expressed in the following way:

$$VR_{j \to i} = X_j \frac{F_i e^{-d_{i,j}/\beta}}{\sum F_q e^{-d_{q,j}/\beta}} \rho_F^{d_{i,j}} \tag{5.1}$$

where the parameters including the mean dispersal rate for foraging β and the survival rate per meter during foraging $\rho_F^{d_{i,j}}$ are taken from published literature (Häussler et al., 2017). d_{ij} is the distance between i and j. Initial nests and flower resources are allocated from conditional Poisson density distributions based on the floral cover- and attractiveness scores of a given cell's land use class. Unlike presence-absence models, nests are distributed across the landscape in a probabilistic way. The model is parametrised for a social guild (ground- and tree-nesting bumblebees) and a solitary guild (solitary bees). Each is present active within the study area (Image et al., 2022). For the social guild, the model has two periods where queens forage during the first floral period and a subset of workers during the second period. The number of workers in the second period is determined by the resources gathered by the queen in the first period. For the solitary guild, only the queens forage (Häussler et al., 2017).

The total visitation rate at a cell is given by its proximity-weighted floral resources score, relative all other cells within the species' foraging distance. X_j is the number of foragers originating from cell j and is logically computed by multiplying the attractiveness-weighted number of nests with the number of foraging workers per nest: $p_w \times w_{max} n_{max}$ when nesting attractiveness is 1.

As discussed in the previous section, pollinator visitation has tangible economic importance (Garibaldi et al., 2016). Many agricultural crops are fertilised by pollen exchanged by foraging insects and insufficient visitation can result in lower genetic diversity and flower quality output. Pollination is an example of what Ellis et al. (2021) call a 'weak-link' problem where agricultural losses are attributed to the land parcels receiving the fewest visits. One variable of interest when evaluating the effect of an ELM scheme is therefore the minimum visitation across a crop field or orchard. Overall visitation is also a predictor of the landscape-level attractiveness to pollinators and community growth (Häussler et al., 2017). Improving the landscape-scale average can therefore also be target from a conservation perspective.

Table 5.2: *Land use categories*

	Ground-nesting bees		Tree-nesting bumblebees	
Land use class	FA	NA	FA	NA
Coniferous Woodland	4.31	0.28	1.33	0.53
Broadleaved Woodland	12.37	0.54	16.00	0.84
Improved Permanent Grassland	4.52	0.33	4.57	0.29
Unimproved Permanent Grassland	16.13	0.57	20.0	0.08
Growing cereals	1.83	0.40	1.00	0.00
Oilseed rape	16.21	0.41	20.00	0.00
Orchards	17.57	0.72	20.00	0.60
Strawberries	13.3	0.42	18.67	0.00

Notes: FA = floral attractiveness, NA = nesting attractiveness, as per Häussler et al. (2017). Unimproved permanent grassland later referred to as ”natural regeneration”.

5.2.3 Habitat connectivity

By connectivity I refer to the accessibility between land parcels suitable for pollinator nesting and foraging. Several measures of connectivity have been proposed. A seminal specification by Hanski (1994) and evaluated in Saura and Pascual-Hortal (2007) defines a probability of connectivity (PC) index across a landscape L with area A_L as follows:

$$PC = \frac{\sum_{i=1}\sum_{j=1} a_i a_j p_{ij}}{A_L^2} \tag{5.2}$$

where a is the area of a given disconnected habitat patch and p_{ij} represents the probability of dispersal between two patches i and j. The probability $p_{ij} = e^{-\alpha d_{ij}}$ depends on the distance d between i and j, as well as a constant α set such that $p = 0.5$ for the average dispersal distance of the species. Saura and Pascual-Hortal (2007) highlight a number of advantages of the PC in that it; a) indicates lower connectivity when the distance between patches increases; b) detects as more im-

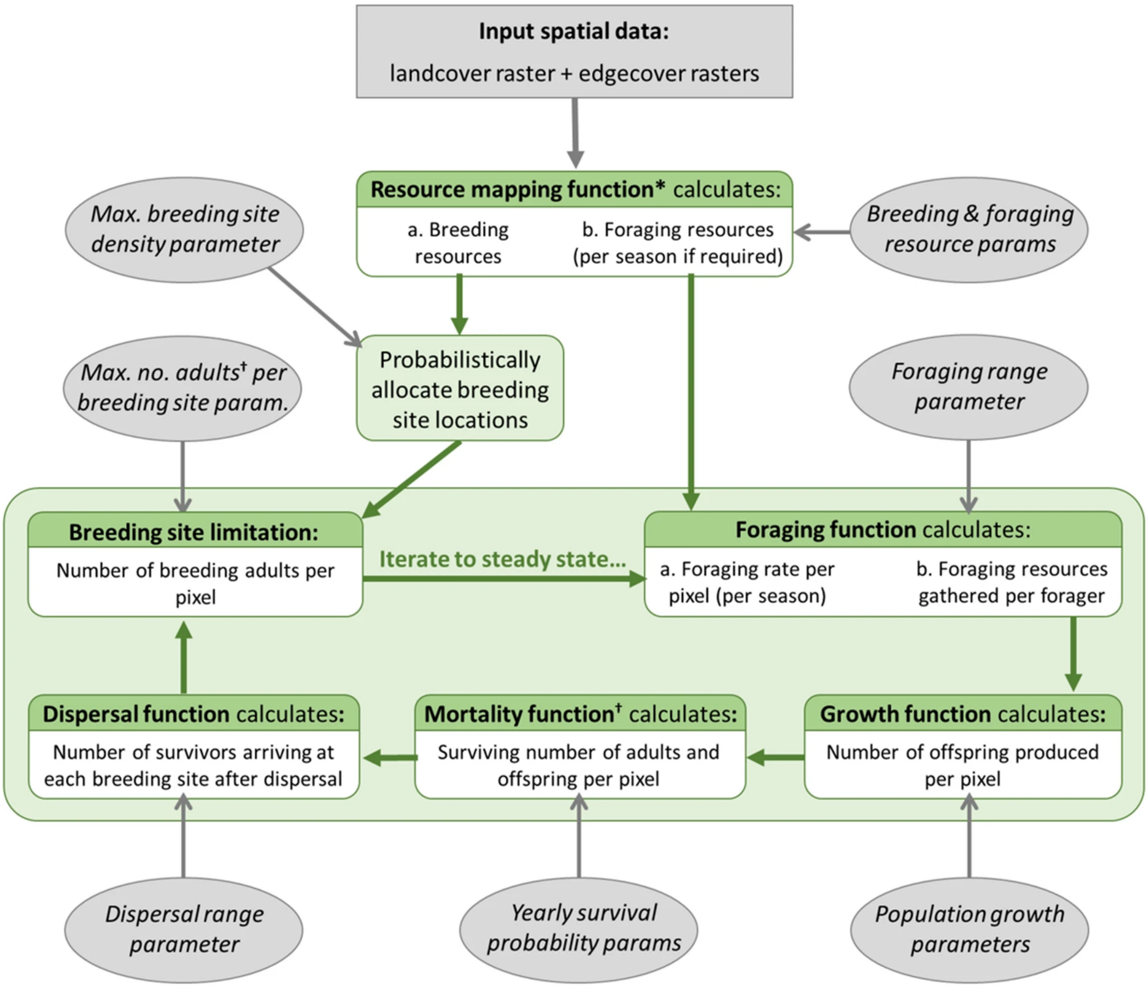

Figure 5.2: *Model process flowchart for* `poll4pop` *from Gardner et al. (2024).*

portant the loss of bigger patches; and c) detects as less important loss of those connecting patches that leave most of the habitat connected. Studying the relationship between connectivity and bee species richness (separate from visitation) Jauker et al. (2013) simplify the computation of the Hanski index by allowing a connectivity value for every patch. The specification widely applied to pollinators (Jauker et al., 2013; Marini et al., 2014) computes the connectivity index of patch i as $CI_i = \sum_{j \neq i} e^{-\alpha d_{ij}} A_j$. The landscape connectivity can then be defined as the average across patches.

Accounting for the distances between habitat patches ensures that different configurations of the same aggregated habitat size can result in different connectivity levels. Environmental policy scholars have taken note of this distinction because it provides an avenue to pursue conservation, potentially without taking much agricultural land out of production (Krämer & Wätzold, 2018). The following chapter will present a model to predict farmers' willingness to enrol in such a ELM scheme.

Algorithm 1: Count disconnected patches

To compute the connectivity index, I must define and count all disconnected habitat patches in the landscape. Here I apply an algorithm colloquially known as *Count Islands* (Cormen et al., 2022) and modified for this research to deal with large raster files. The algorithm counts fragmented patches by recursively checking if any direct neighbours of a suitable habitat cell are also classed as suitable. Once no more suitable neighbours can be found, the recursive search ends and the visited cells are designated an 'island'. The recursive method can easily overflow the stack of search functions to call when the land use raster is large. I solve this problem by recognising that any habitat cell connected on all sides to other habitat cells can not contribute or detract from the connectivity count. I declare them as already visited upon first calling the algorithm:

$\mathbf{r} \leftarrow M \times N$ `land use matrix`
$\mathbf{v} \leftarrow M \times N$ `matrix`
`count` $\leftarrow 0$
for $m \leftarrow 1$ to M **do**
    for $n \leftarrow 1$ to N **do**
        if $\mathbf{r}_{n,m} \in \texttt{nf}$ **and** $\mathbf{v}_{n,m} = 0$ **then**
            $\mathbf{v}_{n,m} \leftarrow 1$
            $\mathbf{v} \leftarrow \texttt{search}(\mathbf{r}, n, m, \mathbf{v}, \texttt{nf})$
            `count` $++$
        end if
    end for
end for

where the function `search` visits the neighbours of $\mathbf{r}_{m,n}$ and - if they are classed as natural features - recursively visits their neighbours. The recursive process continues until no `nf` neighbour can be found, when the updated matrix of visited cells is returned and the count of islands is increased by one.

5.3 Model

Consider a farm producing some agricultural output Y. Agricultural output is given by a Cobb-Douglas production function displayed in equation (5.3) taking as inputs land L_{ag}, non-land inputs X, and (for some products) pollination V (Dawson & Lingard, 1982). The Cobb-Douglas exponents $\alpha+\beta+\gamma=1$ means constant returns to scale. The Cobb-Douglas function has useful properties with respect to pollinator visitation rates, because when $\gamma=0$ - interpreted as no pollination dependence - we get $V^0=1$ for all levels of V and the production is given by $X^\alpha L_{ag}^\beta$. There is a saturation effect from increasing pollination inputs, as the probability of fertilising a flower is cumulative with insect visits. I assume that there is no impact of non-land inputs X on pollinator services supplied (e.g. no impacts of higher pesticide use). This simplifying assumption is made for two reasons. First, efforts to reduce pesticides is bundled up with other ELM schemes that were not explored for the survey (Defra, 2022). Second, organic farming typically means greater demands on L_{AG} to maintain yields (Finger & Möhring, 2024). In this case, incentives run counter to the provision of multifunctional benefits (e.g. flood management) explored elsewhere in this thesis.

$$Y = X^\alpha L_{ag}^\beta V^\gamma \tag{5.3}$$

I have previously established pollinator visitation as a so-called weak-link problem where insufficient visitation can result in lower yield quality and/or quantity, but increasing visitation from a healthy baseline is uncertain to increase yields in e.g. oilseed rape (Garratt et al., 2018). On this basis I assume that $0 \leq \gamma < 1$. Second, I assume a decreasing marginal product from land $0 < \beta < 1$. This is based on a generalisation that the availability of land suitable for specific crops is limited in the UK, that the market is characterised by a plurality of small farms, and that government support programs have traditionally focused on pluriactivity, enhancing the sum of agricultural and non-agricultural incomes (Marsden & Sonnino, 2008). The assumption that $0 < \alpha < 1$ goes as follows: Dedicating more labour and capital to a limited amount of land results in diminishing yield returns (Desiere & Jolliffe, 2018). In addition, a shortage of farm workers following Brexit has worsened the prospect of offsetting production losses with labour. I assume a competitive output market where the individual farmer can not influence the price or collude with competitors to do so. In the absence of government programs, the farm's objective is to minimise costs subject to meeting its residual demand. It is also constrained by its land endowment, which we assume to be fixed in the short-run, as is typical in production economics.

Consider a hypothetical ELM scheme designed to be conceptually similar to the Sustainable Farming Initiative (SFI) piloted by the UK Department for Environment, Forestry, and Agriculture (Defra). It provides funding for long-term, large-scale projects that "restore priority habitats, improve habitat quality, and increase species abundance" in England by, e.g. building or linking nature reserves, creating woodlands, or improving habitat connectivity (Defra, 2022). Specifically aimed to improve connectivity, participating farmers receive an annual payment per meter ℓ of a natural corridor created across their fields. These corridors should have high pollinator attractiveness scores such as flower strips. Additionally, suppose that on top of the annual payment π, the scheme features a bonus payment B for coordinating with n neighbouring farmers to connect habitats with strips of set-aside land that improve connectivity (Correa Ayram et al., 2016). I state the Lagrangian from the farmer's objective function:

$$\texttt{min}\ \mathcal{L} = p_X X + r L_{ag} + C(n) - \pi\ell - Bn - \mu_1(Y - X^{\alpha} L_{ag}^{\beta} V^{\gamma}) - \mu_2(\bar{L} - L_{ag} - w\ell) \tag{5.4}$$

The length of the corridors enters into the land endowment constraint because the scheme will plausibly specify a minimum width, w, such that the area $w\ell$ is added to the retired area. This implies that the farmer should view w as a scaling-up factor for the amount of natural features they need to create, which leads us to the first hypothesis used in validating the model:

Hypothesis I: Farmers require a larger government payment to increase the width of any natural corridors created on their land.

Pollinator visitation V is a function of connectivity enhanced both by the length of corridors cutting through the agricultural landscape and the geographical extent of those features increased by n. I assume that the corridors are placed in such a way that both ℓ and n increase connectivity. Coordination with neighbours may involve costs that we call coordination costs $C(n)$ where $C'(n) > 0$, as they need to communicate and agree on corridor placements that may be suboptimal for the individual. The farmer chooses their amounts of X, L_{ag}, ℓ, and the number of neighbours to collaborate with.

$$\frac{\partial \mathcal{L}}{\partial X} = p_X + \mu_1 \alpha X^{\alpha-1} L_{ag}^{\beta} V^{\gamma} = 0 \tag{5.5}$$

$$\frac{\partial \mathcal{L}}{\partial L_{ag}} = r + \mu_1 \beta X^{\alpha} L_{ag}^{\beta-1} V^{\gamma} + \mu_2 = 0 \tag{5.6}$$

$$\frac{\partial \mathcal{L}}{\partial \ell} = -\pi + \mu_1 \gamma X^{\alpha} L_{ag}^{\beta} V^{\gamma-1} V'(\ell) + \mu_2 w = 0 \tag{5.7}$$

$$\frac{\partial \mathcal{L}}{\partial n} = -B + C'(n) + \mu_1 \gamma X^{\alpha} L_{ag}^{\beta} V^{\gamma-1} V'(n) = 0 \tag{5.8}$$

From the first-order conditions and the constraints we can derive the demand functions for the cost-minimising allocations of ℓ. The demand for corridor length is:

$$\ell^* = \frac{1}{w}\left[\bar{L} - \left(\frac{Y}{\left[\frac{\alpha}{\beta}\frac{\pi + rw - (B - C'(n))\phi}{p}\right]^{\alpha} V^{\gamma}}\right)^{\frac{1}{\alpha+\beta}}\right] \tag{5.9}$$

where $\phi = V'(\ell)/V'(n) > 0$, i.e. the ratio of marginal visitation rate from corridor length to the marginal rate from coordination. When $\phi < 1$, the marginal effect from coordination with neighbours is larger, for example because connectivity on the farm level is already sufficient or because the farm itself is small compared to surrounding ones. We refer to ϕ as the *connectivity insensitivity ratio*, recalling the Hanski connectivity index. The following reasoning provides the name: When $\phi > 1$, the pollinator visitation rate increases more from the marginal increase in the amount of habitat created in a given fixed-size plot of land ($V'(\ell)$, than from the marginal increase in habitat connection with neighbouring fixed-size plots ($V'(n)$). Connecting corridors across neighbouring plots increases the connectivity index but does not increase the amount of habitat in each pixel. $\phi > 1$ implies a relative insensitivity to marginal connectivity improvements.

Figure 5.3 displays 3-D space as 2-D contours from the demand function for corridors ℓ for variation in ϕ and γ, the pollinator dependence of the farmer's crops. Under Cobb-Douglas production, the marginal demand for an additional meter of corridor is positive and diminishing in γ. Conditional on the assumption that farmers consciously internalise pollination benefits, those who grow crops more reliant on pollinators are expected to create more ecological corridors, given certain payment and coordination bonus. When marginal coordination costs are increasing with the number of coordinating neighbours ($C'(n) = n^2$), the optimal corridor length will increase in ϕ when a) there is zero coordination, be independent of ϕ when b) there is one coordinating neighbour, and decrease in ϕ when c) there are two coordinating neighbours. In case a) marginal coordination costs are zero and demand for NFM on the farmer's own land will decline at a steeper rate as $V'(n) > V'(\ell)$, i.e. $\phi < 1$.

In this case, substituting own NFM for more coordination is not only lower cost but also improves visitation at the margin. In case b) the marginal coordination cost of 1 is equal to the coordination bonus (assuming unit prices) and so ϕ disappears from the demand function. To see why this result is necessary, imagine that $V'(n) \to 0$ and therefore that $\phi \to \infty$. For example, one could imagine hypothetical pollinators with a maximum foraging distance of only a meter, at which point coordination with habitats on neighbouring farms would be close to useless. But when the marginal coordination cost equals the coordination bonus $C'(n) = B$, the cost of additional coordination is zero. It follows that there is no value for $V'(n) > 0$ small enough to dissuade the farmer from coordinating with one additional neighbour. Then, the amount of NFM created by the farmer will depend only on their pollinator dependence. Similar reasoning applies for linear- and diminishing marginal coordination costs.

Häussler et al. (2017) also suggest that the introduction of natural features such as flower strips can induce competition for pollinators among pollinated species, including flowering crops. Under such competition the creation of flower-rich corridors in a field where pollinator-dependent crops are growing may result in a marginal decline in visits to these economic crops. In this edge case $V'(\ell) < 0$ and $\phi < 0$. $V'(n)$ is assumed to be strictly positive.

With increasing marginal costs, farmers favour the payment for ecological corridors over the coordination bonus. From the first order conditions (5.5) and (5.8), I show that the farmer is expected to choose their level of coordination n such that $B - C'(n) = -(\gamma/\alpha)pXV'(n)/V$. The farmer may engage in coordination even if the marginal coordination cost exceeds the bonus if the loss is offset by savings in capital inputs from improving pollination. Under the assumption that the

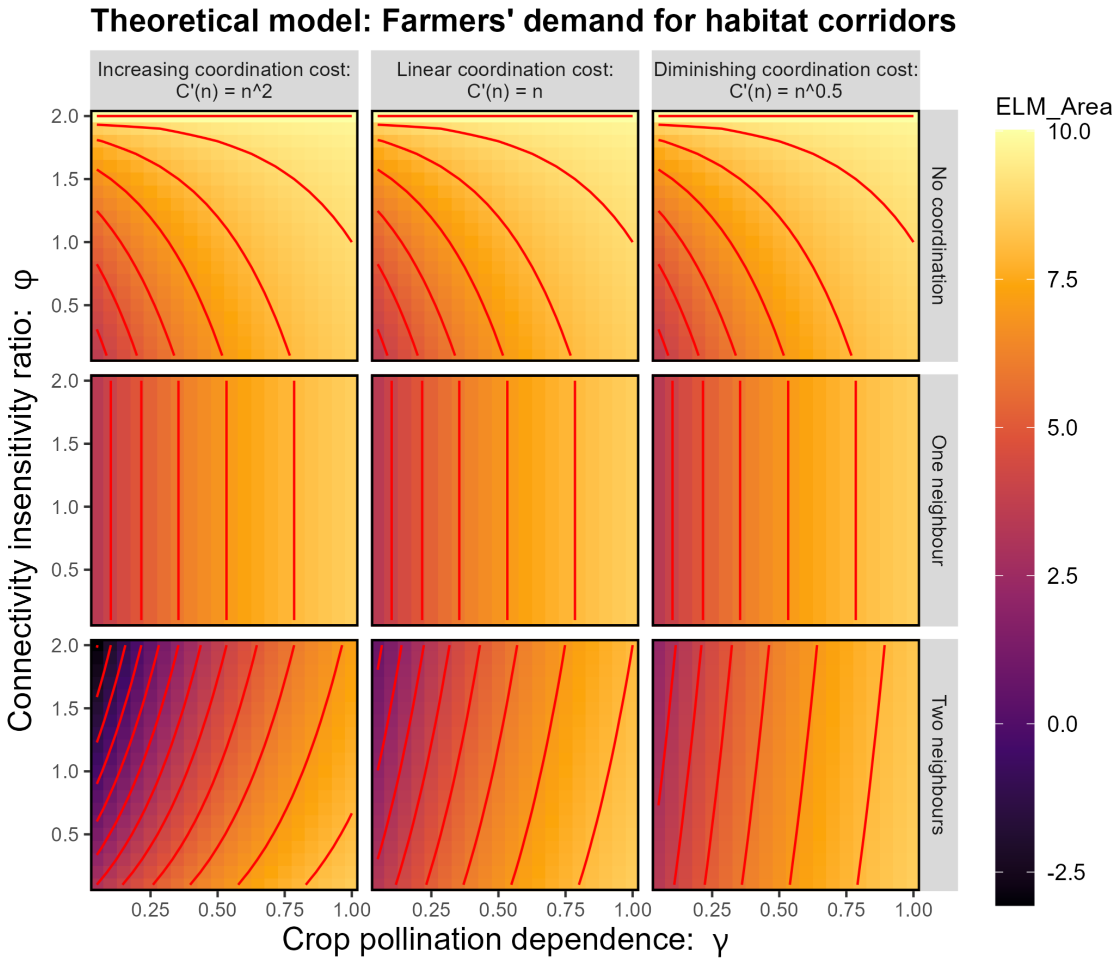

Figure 5.3: *Simulated demand for ℓ plotted against pollinator dependency (γ) and the connectivity insensitivity ratio ϕ. From left to right, demand contours are shown for increasing, stable, and diminishing coordination costs, respectively.*

marginal coordination cost is positive, $C'(n) > 0$, we can formulate the following hypotheses:

Hypothesis II: Farmers will express a preference for a lower level of coordination over a higher level of coordination when the bonus is held constant.

Increased information sharing has been found to increase efficient coordination in an experimental setting (Nguyen et al., 2022). However Banerjee et al., 2014 caution that this efficient equilibrium may deteriorate over time. For example, a farmer may falter in trust that her immediate neighbour will cooperate if she learns that more distant peers in the local network have not. This risk is mitigated when there is a strong social connection between a farmer and their immediate neighbours (Banerjee et al., 2017). Familiarity facilitates greater trust between neighbouring farmers that coordination will persist. As a corollary to hypothesis II, the effect of social connection on the preference against coordination is tested in hypothesis III:

Hypothesis III: Farmers expressing a strong social connection with their neighbours display a weaker preference for a lower level of coordination when the bonus is held constant.

In this section we have presented a model of agricultural production which acknowledges the contribution of pollination services by incorporating crop visitation and pollinator-dependence. The model predicts that visits to crops by pollinating insects can be enhanced by increasing the habitat connectivity of the agricultural landscape, and possibly by replacing productive farmland with more attractive habitats such as broadleaf trees or corridors set aside for rewilding. The exact prediction from the model follows from my assumptions about ϕ, verbalised as Hypothesis IV in table 5.3.

If $\phi < 0$, the natural features compete with pollinated crops for visits from pollinators. In this case, hypothesis III predicts that increasing either the amount of natural features, L_{NF}, or the degree of connectivity-enhancing coordination between neighbouring farms, n, will have a negative effect specifically on pollination of economic crops. If $\phi > 1$, increasing L_{NF} also increases crop visitation. Specifically, visitation depends primarily on the amount of habitat created, more than connectivity improvements. This could be true if, for example, more natural features support more nests but long foraging distances mean that the placement of features can not inhibit visitation. Conversely, if $0 < \phi < 1$, increasing total area of created natural features increases crop visitation, but primarily through the increasing habitat connectivity channel. In this case, scaling up the amount of land used for natural features will improve visitation but mostly insofar as it improves connectivity.

Crucially, this also determines whether visitation can be improved by arranging a set amount of features in a way which improves connectivity. By testing this hypothesis, we provide policy-relevant guidance on the viability of programs that seek to maximize connectivity while retiring a limited amount of productive land.

Table 5.3: *Predictions from the magnitude of connectivity insensitivity ratio ϕ*

Hypothesis IV	Prediction
$\phi < 0$	$\frac{\partial V}{\partial L_{NF}} < 0$
$0 < \phi < 1$	$0 < \frac{\partial V}{\partial L_{NF}} < \frac{\partial V}{\partial n}$
$\phi > 1$	$0 < \frac{\partial V}{\partial n} < \frac{\partial V}{\partial L_{NF}}$

5.4 Econometric modelling

Hypotheses I-III are tested by estimating taste parameters for individual attributes. This is done by estimating a latent class model, using data recorded from the questionnaire. Following the procedure from chapter 3, taste parameters from the hypothetical DCE are estimated using a latent class logit model (Greene & Hensher, 2003). This approach helps to identify and understand the different consumer segments that may exhibit diverse decision-making patterns, which can be helpful for designing targeted policy interventions (Tyllianakis et al., 2023). It was communicated to respondents that the natural features were the same as in DCE I and that he bonus scales linearly with the number of neighbours. If the respondent does not coordinate with anyone, the bonus payment is always zero. If they coordinate with at least one neighbour, the payment to each coordinating farmer is multiplied by their total number (including the respondent). In addition to the five attributes listed in table 5.4, I interact the coordination attribute with an indicator variable which takes the value of 1 if the respondent states that they regularly share farm equipment with a neighbour, following Sheremet et al. (2018). Approximately 45% of survey respondents indicated that they share equipment with neighbours. The frequency is greater than in Sheremet et al. (2018) and can be explained by noting that northern England is more densely populated than Finland. I treat this variable as an indicator of generalised collaboration costs. The hypothesis is that current regular collaboration with neighbours, e.g. sharing of farm equipment, is indicative of lower coordination costs due to habit formation and pro-social attitudes (Banerjee et al., 2014).

Table 5.4: *Discrete choice attributes and levels*

Attribute	Levels
Type: *The corridor feature*	Natural Regeneration, Planted Broadleaf Trees
Width (w): *The required width of corridors*	10 meters, 20 meters
Coordination (n): *The number of connected farms*	None, One, Two
Bonus (B): *One-time bonus payment per connected farm*	£100, £200, £300, £400
Payment (π): *Annual payment per 100m of corridor*	£200, £300, £400, £500

Continuing to follow Boxall and Adamowicz (2002), the number of classes is decided based on minimising the Bayesian Information Criterion (BIC). Models with $2-4$ classes were estimated, but with no more than two classes did the model converge. The BIC for the two-class model was 4612 compared to 4796 for the base MNL model. Accordingly, the model with two latent classes is estimated, with the utility from option (ELM scheme) i specified as follows:

$$\begin{aligned} U_{s,i} = ASC_{i,s} + ASC_{i,s} \times GRAZING + \\ \beta_{TREES,s} \times TREES + \beta_{WIDTH_{20m},s} \times WIDTH_{20m} + \\ \beta_{COORDINATION_{n=1},s} \times COORDINATION_{n=1} + \\ \beta_{COORDINATION_{n=2},s} \times COORDINATION_{n=2} + \\ \beta_{BONUS,s} \times BONUS + \beta_{PAYMENT,s} \times PAYMENT + \\ \beta_{(n=1)\times SHARING} \times (COORDINATION_{n=1} \times SHARING) + \\ \beta_{(n=2)\times SHARING} \times (COORDINATION_{n=2} \times SHARING) + \delta_s \end{aligned} \tag{5.10}$$

Equation (4.10) models the utility that farmers in class s derive from choosing option i. The attributes are described in table 4.6. The alternative-specific constant, $ASC_{i,s}$, is interacted with a variable indicating the proportion of land the respondent uses for grazing. Hypothesis I is stated as the following null and alternative hypotheses. It is evaluated using a one-sided t-test.

H0: $\beta_{WIDTH_{20m}} = 0$

H1: $\beta_{WIDTH_{20m}} < 0$

Hypothesis II is stated as the following null and alternative hypotheses. H1 is a joint inequality and is evaluated via $10,000$ draws from the bivariate distribution of $\beta_{COORDINATION_{n=1},s}$ and $\beta_{COORDINATION_{n=2},s}$, following the procedure in section 4.4 of chapter 4:

$$\text{H0: } \beta_{COORDINATION_{n=2}} = \beta_{COORDINATION_{n=1}} = 0$$
$$\text{H1: } \beta_{COORDINATION_{n=2}} < \beta_{COORDINATION_{n=1}} < 0$$

To test hypothesis III, the coordination attribute (deciding whether the farmer has to connect ELM features with zero, one, or two neighbours) is interacted with a binary variable indicating whether the they regularly share farming equipment with neighbours. This is used as a proxy for coordination costs, assuming that farmers who collaborate with neighbours professionally find it easier to coordinate. The null hypothesis is rejected if farmers facing low coordination costs are significantly more likely to coordinate:

$$\text{H0: } \beta_{(n=2)\times SHARING} = \beta_{(n=1)\times SHARING} = 0$$
$$\text{H1: } \beta_{(n=2)\times SHARING} > \beta_{(n=1)\times SHARING} > 0$$

5.5 Simulation of pollination services

The `poll4pop` model calculates visitation rates for each cell in a raster based on land cover data over the same extent and resolution and estimates of ecological parameters from published literature (Häussler et al., 2017). I use the most recent 10m^2 resolution land cover data provided by the UK Centre for Ecology and Hydrology (Rowland et al., 2020) and crop cover data provided by the Rural Payments Agency of the UK. I select a 4km^2 area around the location of each farm in my sample as a baseline in an effort to capture the possible effect from connectivity and coordination between neighbouring farms (the area can fit four average-sized farms of 100 ha) while ensuring estimates that are relevant to the individual farm.

The model was first applied to Swedish data (Häussler et al., 2017) but have since been used to evaluate the effectiveness of environmental land management interventions in the UK (Image et al., 2023). This recent work has shown that hedgerow or woodland edge management had the largest positive effect on pollination service change, due to high resource quality. Fallow areas were also strong drivers, despite lower resource quality, implying effective placement. Interventions had stronger effects where there was less pre-existing semi-natural habitat. The visitation model has been validated for application to English agricultural landscapes (Gardner et al., 2020; Image et al., 2023) but specifying its relationship with connectivity is outside the scope of these studies. In this article, following the suggestions in Image et al. (2023), I study two hypothetical interventions; (1) planted broadleaf trees and (2) natural regeneration where land is taken out of production and flowers protected from grazing. I calculate expected visitation rates before and after each intervention.

As in chapter 4 I study four different spatial configurations of these natural features: i) Corridors along field edges, ii) in-field corridors, iii) evenly distributed in-field islands, and iv) a contiguous patch of land at the edge of field, but nonetheless taking a portion of farmland out of production. In each case i) to iii), I let the width of the features be either 10 meters or 20 meters across, mirroring the attributes in the choice experiment. The size of the contiguous patch was determined so as to match the combined area set aside for field-edge, and in-field corridors. The in-field islands are small 10×10 or 20×10 meter patches that are distributed evenly across the field. I let the gaps between corridors and islands vary between 200, 300, 500, and 800 meters. Larger gaps between natural features mean less farmland taken out of production and less need for coordination between farmers, at the expense of fewer habitats and less connectivity.

For each combination of feature type, spatial configuration, feature width, and feature gap, I compute average crop visitation rates and total pollinator abundance using `poll4pop`. I compare these two metrics for the treated and untreated landscape, without added natural features. I repeat this procedure for a 4 km^2 area around each farm in the survey sample.

5.5.1 Visitation model inputs

The `poll4pop` model takes as inputs two sets of data. Species-specific parameters and land use data. The species-specific parameters include the nesting- and foraging attractiveness of each land use class and the foraging distance for each pollinator species. The species-specific parameters are provided in Häussler et al. (2017) and summarised here in tables 5.1 and 5.2. As land use inputs I use the crop map of England (CROME).

Crop Map of England: CROME (Rural Payments Agency, 2021) is a polygon vector dataset mainly containing the crop types of England. The dataset contains approximately 32 million hexagonal cells classifying England into over 15 main crop types, grassland, and non-agricultural land covers, such as Woodland, Water Bodies, Fallow Land and other non-agricultural land covers. The classification was created automatically using supervised classification (Random Forest Classification) from the combination of Sentinel-1 Radar and Sentinel-2 Optical Satellite images during the period late October 2021 – September 2022. The results were checked against survey data collected by field inspectors and visually validated. CROME has been repeatedly used for research in ecology and agricultural science, including in Image et al. (2022), Image et al. (2023), and Upcott et al. (2023). Examples of CROME maps and simulated natural features are displayed in figure 3.2.

Distributions of crops across a sample of 306 farms are displayed in figure 5.4. Grassland is by some margin the most common land use type in the agricultural landscapes, with a median land cover share of 81%. Broadleaf woodland serves as naturally occurring habitat for tree-nesting bumblebees, but rarely makes up more than 10% of the land surrounding sampled farms, and most commonly less than 5%. Pollinator-dependent economic crops occur in the form of oilseed rape and field beans but make up only a minority of the agricultural land use.

Following Häussler et al. (2017), I focus on three groups of pollinators, ground-nesting solitary bees, ground-nesting bumblebees and tree-nesting bumblebees. Ground-nesting bumblebees (*Andrena*) make up 75% of foraging bees species. Notable examples native to the UK include the red mason bee and the tawny mining bee. Many species in this group are small in body size (1-2cm) which is associated with comparatively short foraging distances of 100-300 meters (Antoine & Forrest, 2021). Less mobile species are of particular interest when estimating the value of connectivity improvements, as these may be vulnerable to habitat fragmentation even at small scales.

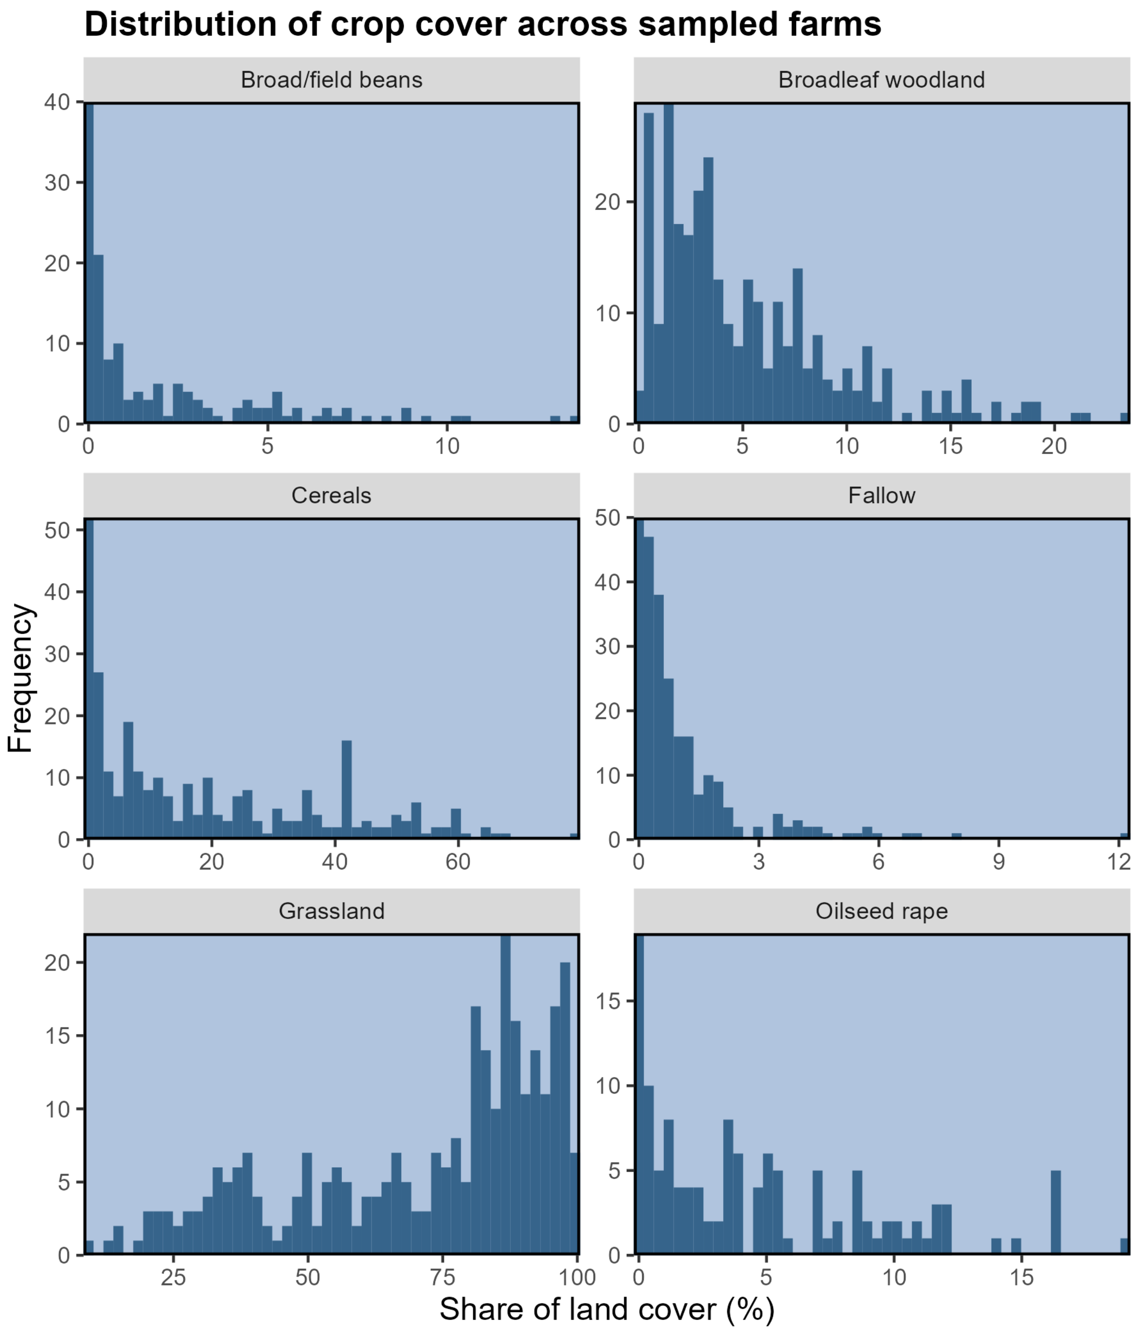

Figure 5.4: *Land use distributions across sampled farms*

Bumblebees fill an important niche as effective pollinators in temperate and wet climates (Liczner & Colla, 2019). A growing body of evidence highlight the importance of trees as nesting sites (Mola et al., 2021). Running `poll4pop` on the simulated landscapes, I estimate lower quartile average, and upper quartile visitation rates to economic crops at both the landscape and field scale. I focus in particular on oilseed rape and broad and field beans, that display moderate pollinator-dependence (Breeze et al., 2021). Specifically, the visitation rate VR_{ij} is the rate at which flowering crops within cell i in the field is visited by foraging bees from cell j.

I judge the effectiveness of each scheme on the basis of the resulting change in average visitation rates across cells used for pollinated economic crops. The insect pollinated crops featuring in the crop cover data are oilseed rape, field beans, each with ca 25% of yields at risk from pollinator declines. While no attempt is made in this thesis to translate the change in visitation to a change in yields, the pollinator dependence chart in figure 5.1 provides guidance. I compute the aggregate visitation from all three pollinator species in the model. The change is calculated as the post-treatment change in visitation as a percentage of the pre-treatment visitation. Finally, I divide the change by the area of economic farmland set aside for natural features. This yields the effect of the scheme on visitation per m^2 of natural features created. In this way, inefficient land use is penalised, and allow for cost estimates of the schemes based on choice experimental results.

5.5.2 Quantifying the connectivity insensitivity ratio ϕ

As shown in my theoretical model, benefits to pollination services from natural features and coordination between farmers depend on the ratio between the marginal rate of visitation per m^2 of natural features and marginal visitation per connectivity improvements facilitated by greater coordination. In particular, coordination in this case means maintaining the same amount of natural features per farm but arranged in a way which improves connectivity between neighbouring farms. The magnitude of this ratio $\phi = V'(L_{NF})/V'(n)$ governs how the model predicts that visitation rates will change as farmers substitute more natural features for more coordination, and vice versa.

It is therefore important to establish at least a directional understanding of ϕ. First, it allows us to validate the predictions arising from the model. Second, policy recommendations for future AES schemes of this type depend on understanding whether or not increases in connectivity via coordination can compensate for reductions in the total amount of productive farmland set aside for natural features.

Consider a 2-D surface with connectivity, driven by increasing coordination n with a set amount of L_{NF}, along the y-axis and with the amount of natural features along the x-axis. At every point of the surface representing a n-L_{NF} pairing is an associated change in the visitation rate. Consider first the case where ϕ positive and large which means that $V'(L_{NF})$ is much greater than $V'(n)$. In this case we would expect a horizontal gradient in V' as L_{NF} increases but not much change vertically in n. Conversely, in the case where $V'(n)$ is much larger and ϕ approaches zero, we expect a vertical gradient in n to dominate.

Figure 5.5 shows the empirical 2-D visitation change gradient along dimensions of connectivity and NF area within the sample. Each tile represents an aggregation of point observations across farms into brackets of connectivity and L_{NF}. A diagonal pattern is observed in the underlying scatter plot which illustrates the correlation between the two. However, I do not observe a clear gradient in the visitation

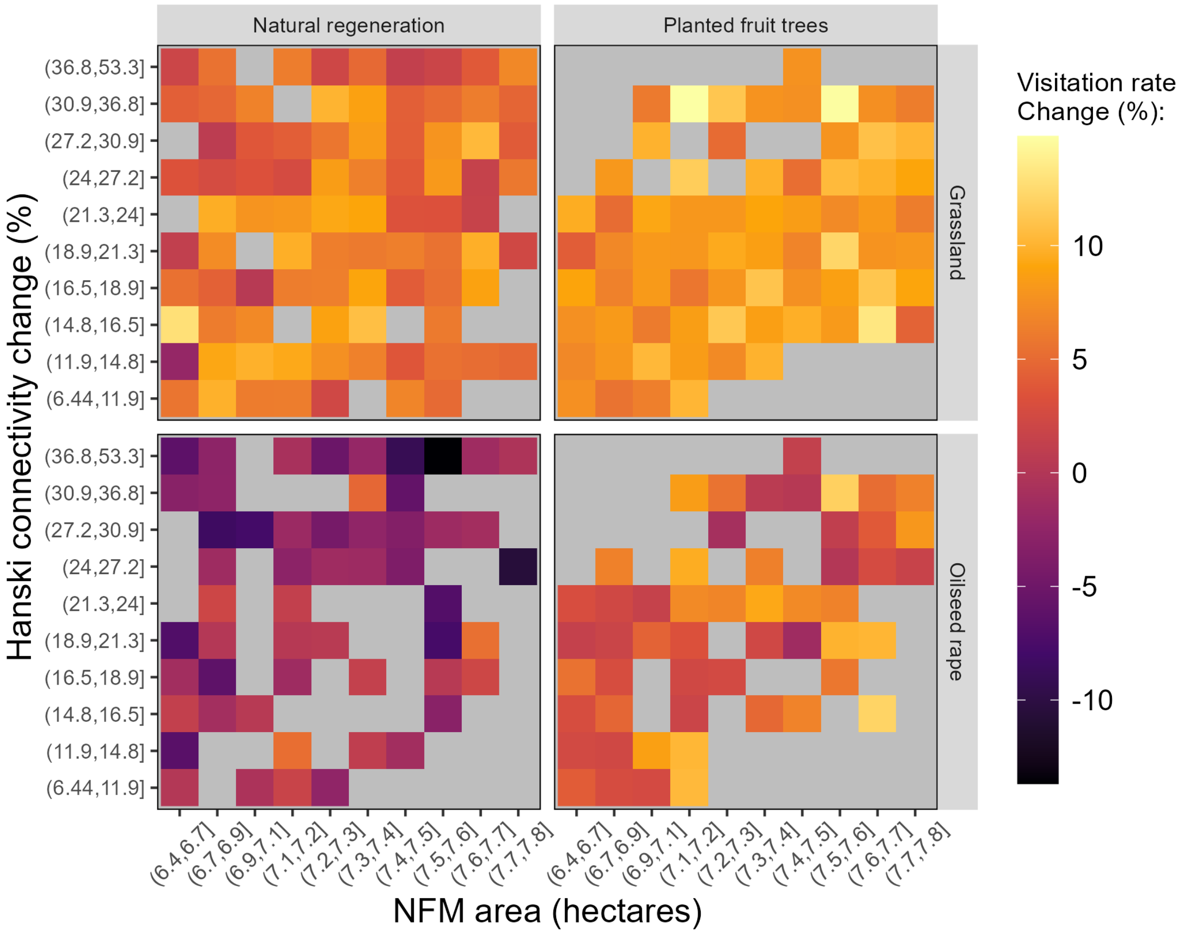

Figure 5.5: *2-D gradient in visitation rates change*

change.

Figure 5.6 illustrates the challenge of separating $V'(\ell)$ from $V'(n)$. The box plots show summary statistics on increases in the Hanski connectivity index (Hanski, 1994) resulting from each of the schemes. The magnitudes of connectivity improvements decline as the gap between natural features widen. This result agrees with the theoretical framework, as increases in the gap between features is a form of habitat fragmentation. However, increasing the gap also reduces the amount of natural features in the landscape. The significant correlation (0.75) between the amount of land used for natural features L_{NF} and the Hanski connectivity introduces severe autocorrelation issues in attempts at using regression analysis to identify the individual effects of L_{NF} and connectivity. I address this issue having designed the schemes in such a way that the contiguous patch, in-field corridors, and field-edge corridors set aside equal amounts of land for any given farm. By grouping farms and spatial configurations into subsamples where L_{NF} is broadly identical within each group, I can isolate the effect from connectivity differences. I fit individual linear models within each group, with regression coefficients and standard errors shown in figure 5.7. The vertical axis shows the average percent change in oilseed rape visitation per percentage point change in the connectivity index. I show the group average pre-treatment connectivity index on the horizontal axis.

I show that $V'(n)$ is positive when the initial pre-treatment connectivity is low, and $V'(n)$ is negative when the pre-treatment connectivity is high. This means that in most cases, $V'(n)$ is expected to be low, resulting in a large magnitude for ϕ. It is nonetheless difficult to determine ϕ with accuracy as it depends on the pre-treated land use configuration as well as the dominant species of pollinators in the landscape. Directionally, I have shown that ϕ can be negative in some cases. This is more likely when pre-treatment connectivity is already high, and might result from economic crops facing competition for pollinators by flower resources on natural features. This is supported by findings in Häussler et al. (2017) who report reduced post-treatment visitation in diverse landscapes due to competition between established and pre-existing flower resources. This allows me to formulate a test of hypothesis IV:

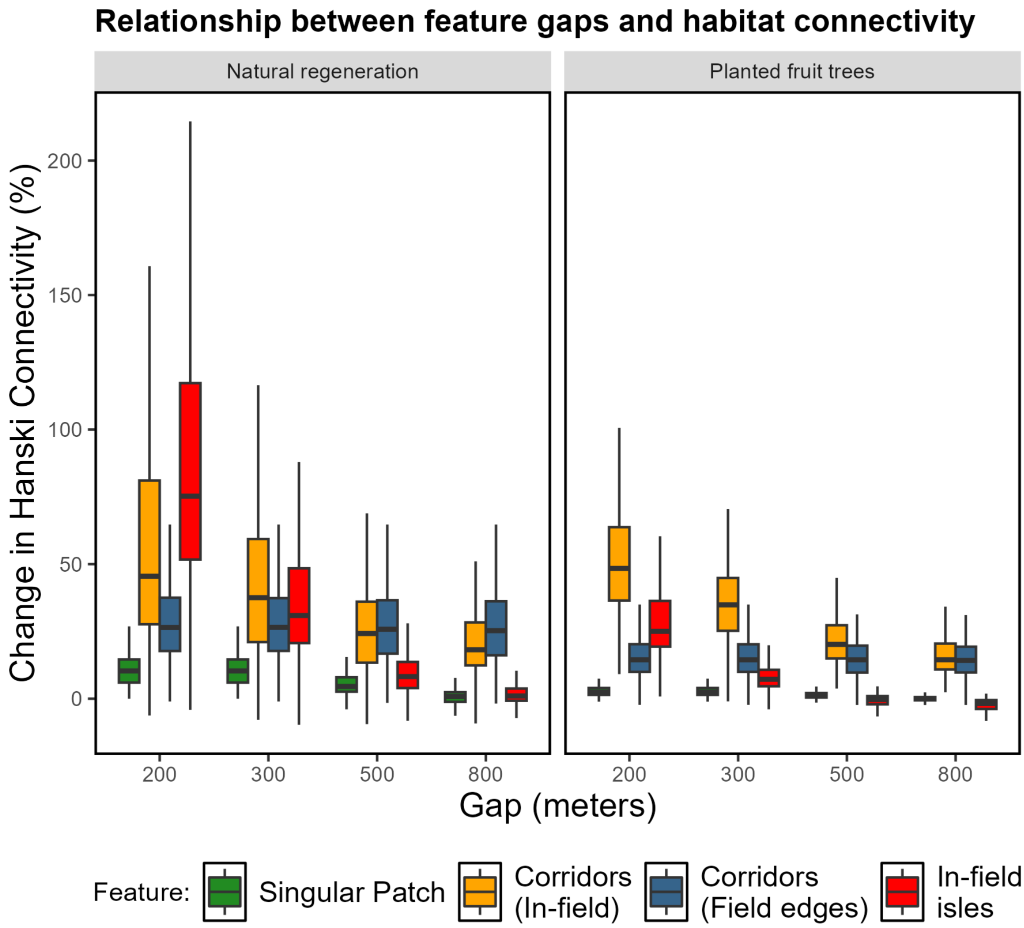

Figure 5.6: *The relationship between habitat connectivity (Hanski, 1994) and gaps between natural features in agricultural landscapes. Edges of boxes represent the lower- and upper quartiles of connectivity improvements across farms in the sample. Middle bands on boxes represent the median.*

Hypothesis IV: On average, increases in the amount of natural features L_{NF} will have a positive effect on economic crop visitation, independent of any increases in connectivity. Resource rich habitats may see a decline in economic pollination from added natural features.

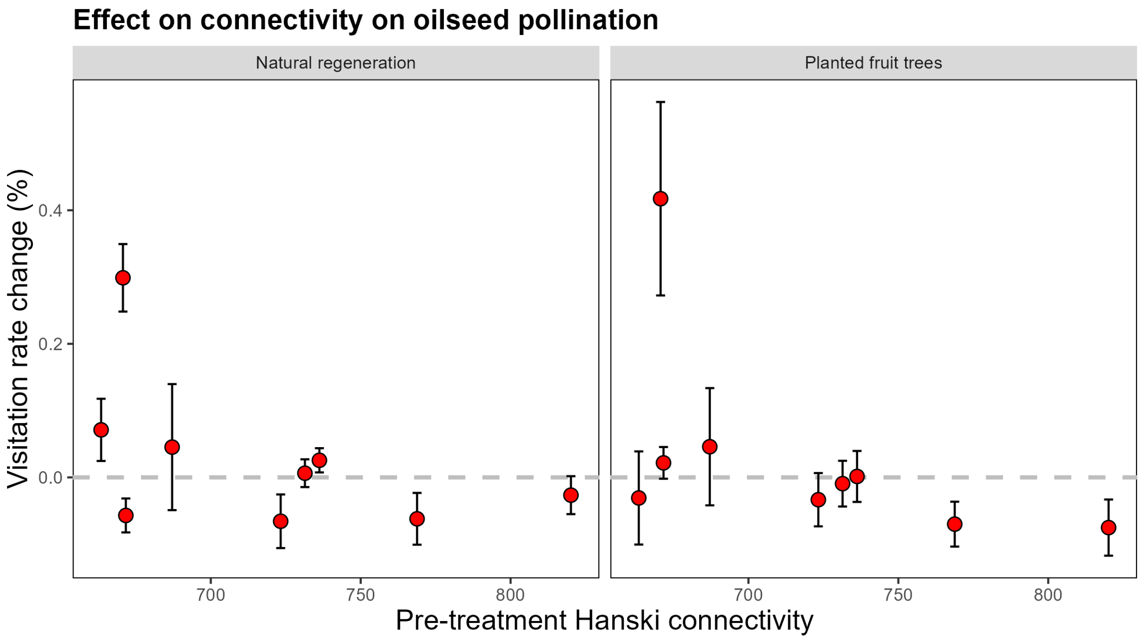

Figure 5.7: *Coefficients β_2 and standard errors for the model $\Delta V_i = \beta_1 + \beta_2 \Delta CI_i + u_i$, where i is unique farm-scheme combinations within groups where increases in L_{NF} are the same. The vertical axis shows the average percent change in oilseed rape visitation per percentage point change in the connectivity index. I show the group average pre-treatment connectivity index on the horizontal axis*

5.6 Results

In this section I report on the results from the discrete choice experiment and the spatially explicit crop visitation model. Combining results from these two methodologies then allows me to perform a cost-effectiveness analysis for each of the hypothetical habitat connectivity schemes. I rely on results from chapter 4 for the relative reduction in payment required to place corridors along a field- or river edge instead of in the field. This cost-effectiveness analysis reports for each scheme the estimated effect on economic crop visitation achievable from a given payment per farmer and year.

5.6.1 Barriers to coordination

I begin by reporting results from a latent class model shown in table 5.5, where I test for the existence of distinct classes of respondents in terms of preferences for the coordination schemes. Each respondent has a posterior conditional probability of belonging to each class. I assign respondents to the class where their conditional probability is at least 80%. The magnitude of the difference between classes I and II is indicated by the parameter δ_{II}. I see that δ_{II} is negative and significant. This indicates that a smaller portion of respondents belong to class II, which means that the classes are of different sizes and that class II is the smaller one. This mirrors the results from the choice experiments in chapter 4.

The alternative-specific constants for each scheme relative to the opt-out, status quo alternative are positive for members of class I and negative for class II. The interpretation is that respondents in class I have a native preference for participating in the scheme, while members of class II prefer to opt out before any changes to the schemes' attributes are considered. This mirrors results from chapter 4. The preference for enrolling in the scheme is moderately higher among farmers who manage more grazing land, indicated by the interaction between the alternative-specific constant and the proportion of land used for grazing.

The taste parameter for planted trees β_{Trees} is negative and significant for both classes. This means that respondents strongly prefer to create corridors of natural regeneration over rows of planted broadleaf trees. Similarly, there is a strong preference for narrower corridors of 10 meters in favour of a width of 20 meters. The taste parameter for wider corridors $\beta_{w=20}$ is also negative and significant for both classes.

Compared to the reference level of no coordination, respondents in the larger class II display a distaste for coordination with one neighbouring farmers. However, the taste parameter for class I is insignificantly different from zero, which means that respondents in this class are indifferent between no coordination and coordination with one neighbour. The taste parameter for coordinating with two neighbours is negative and significant for both classes. In other words, coordinating with two neighbours is less attractive than coordination with one neighbour, and less attractive still than no coordination. It is important to recall that the taste parameters are effects while holding other attributes to be constant. There is an overall preference for less coordination before considering any coordination bonus. These results are in expectation with my model, which predicts that farmers suffer a marginal coordination cost for each additional neighbour they coordinate with.

The taste parameters for the increases in the coordination bonus and the base payment are each positive and statistically significant. This indicates that farmers in both classes would behave in a cost-minimising fashion, preferring more compensation for costly activities. Finally, I interact the coordination attribute with a binary dummy variable which takes a value of 1 if the respondent states that they regularly share farm equipment with neighbours, and a value of 0 otherwise. The taste parameters for the interactions are positive and significant within class II but insignificant within class I. For class II, a positive taste parameter for the interaction means that the distaste for more coordination is weaker if the farmer regularly shares farm equipment with neighbours. This in line with the theory that the marginal coordination cost is lower among farmers who regularly collaborate.

Figure 5.8 compares respondents in latent classes I and II in terms of socio-economic

Table 5.5: *Latent class model: Preferences for coordination*

Attribute	Taste Parameters		Reference level
	Class I	Class II	
$ASC_{SchemeA}$	2.84	-0.70	ASC_{Optout}
	(0.51)***	(0.32)**	
$ASC_{SchemeB}$	2.87	-0.62	ASC_{Optout}
	(0.51)***	(0.30)***	
Trees	-0.39	-1.24	Natural Regeneration
	(0.05)***	(0.16)***	
20 meter width	-0.47	-1.07	10 meter width
	(0.06)***	(0.16)***	
Coordination (n=1)	0.04	-1.04	No coordination
	(0.12)	(0.26)***	
Coordination (n=2)	-0.33	-1.32	No coordination
	(0.17)**	(0.36)***	
Coordination bonus	0.42	0.95	
	(0.26)**	(0.61)*	
Payment	2.99	3.65	
	(0.24)***	(0.64)***	
$ASC_{Scheme} \times$ % Grazing	0.01	0.003	
	(0.005)**	(0.002)	
$\beta_{n=1} \times Sharing$	0.004	0.50	
	(0.13)	(0.25)**	
$\beta_{n=2} \times Sharing$	0.12	0.67	
	(0.15)	(0.31)**	
δ_{II}	-0.94		δ_I
	(0.14)***		
Summary of class allocation for model: Class I (72%) and Class II (28%)			
Adj. R^2 vs observed shares: 0.21, BIC: 4612, MNL BIC: 4796			

and behavioural differences between them. The principal difference is that members of class II are much more likely to select the opt-out alternative than are members of class I. Farmers in class II are also less likely to currently be enrolled in an ELM scheme, less likely to collaborate with neighbours, and less likely to grow pollinator-dependent crops. In this sample of farms, this refers to oilseed rape and broad- or field beans. Members of class II are also moderately more likely to have opted for a vocational- or non-traditional qualification opposed to academic exams or degrees.

I do not find evidence of class allocation based on age, gender, or land endowment. Similarly, the effect of educational attainment is ambiguous. Instead, class allocation is based on behavioural differences: Members of class II are significantly more likely to choose the opt-out alternative, less likely to collaborate, and engage with current ELM schemes. Therefore, I go on referring to class I as the high engagement class and to class II as the low engagement class. This follows the same pattern observed in chapter 4. This dynamic helps explain the comparatively stronger distaste for increased collaboration within the low engagement class. The low engagement class is characterised by less ELM participation and less collaboration, which is indicating a higher coordination cost.

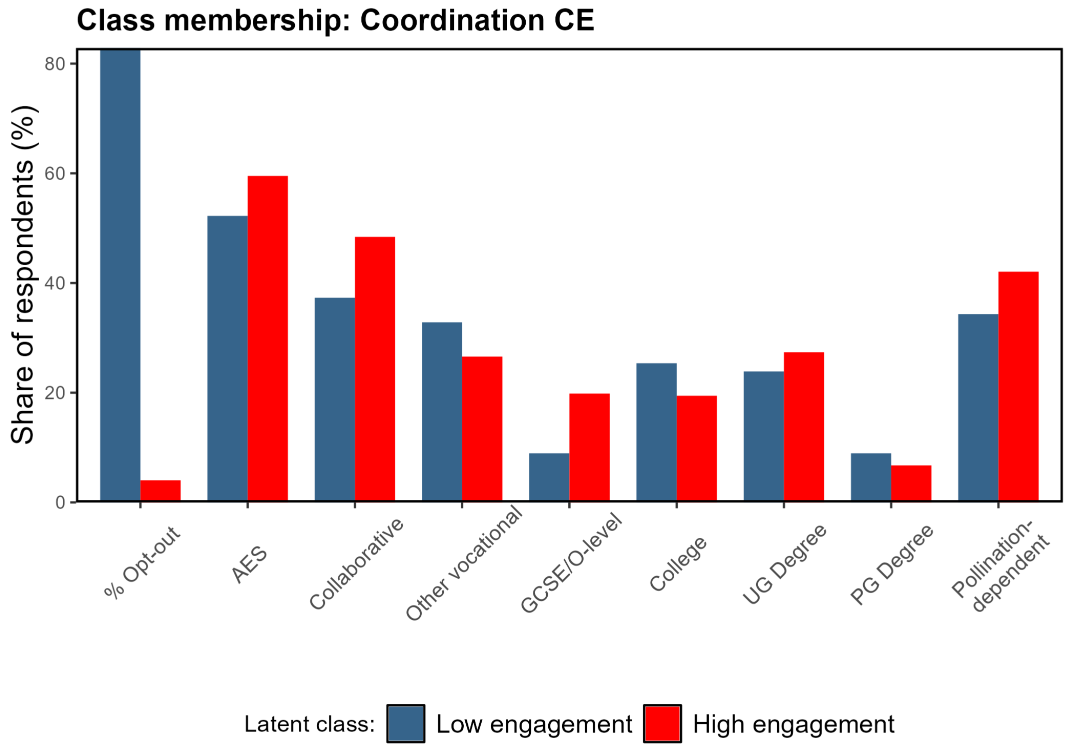

Figure 5.8: *Socio-demographic and behavioural predictors of latent class membership in choice experiment estimating willingness to coordinate with farm neighbours*

5.6.2 Monetary cost estimates

Figure 5.9 shows the taste parameters expressed in monetary terms. This monetary expression is obtained by dividing the taste parameter for the attribute by the parameter for the base payment. The distributions of demanded compensation to create corridors of broadleaved trees instead of natural regeneration are large for both classes. However, the average compensation is ca £100 for the high-engagement class. This means that the payment per 100 meters of corridors needs to be on average £100 higher to incentivise engaged farmers to maintain rows of trees instead of natural regeneration. The distribution of values for the low engagement class is skewed higher, which means that these farmers demand comparatively higher compensation. This is what is expected given the characteristics of the class, as their lower propensity to engage in either real or hypothetical schemes indicate higher perceived costs.

Farmers in both classes demand on average £200 more per 100 meters to make the corridors 20 meters wide instead of 10 meters wide. Invoking the result from section 5.3 that in the corridor creation scheme, $L_{NF} = \ell \times w$, I can compare these results to the results from chapter 4. The increase from a width of 10 meters to 20 meters represents a 1,000 m^2, or $1/10$ hectare, increase per 100 meters of corridors. I therefore estimate an increase in demanded compensation of approximately £2,000 per hectare. I can compare this against the result from chapter 4 where I estimate the willingness to create contiguous patches of natural features instead of corridors. There, the value was closer to £1,000 per hectare. I attribute this difference to the fact that creating corridors is a more complicated activity, with less freedom when it comes to feature shape and placement.

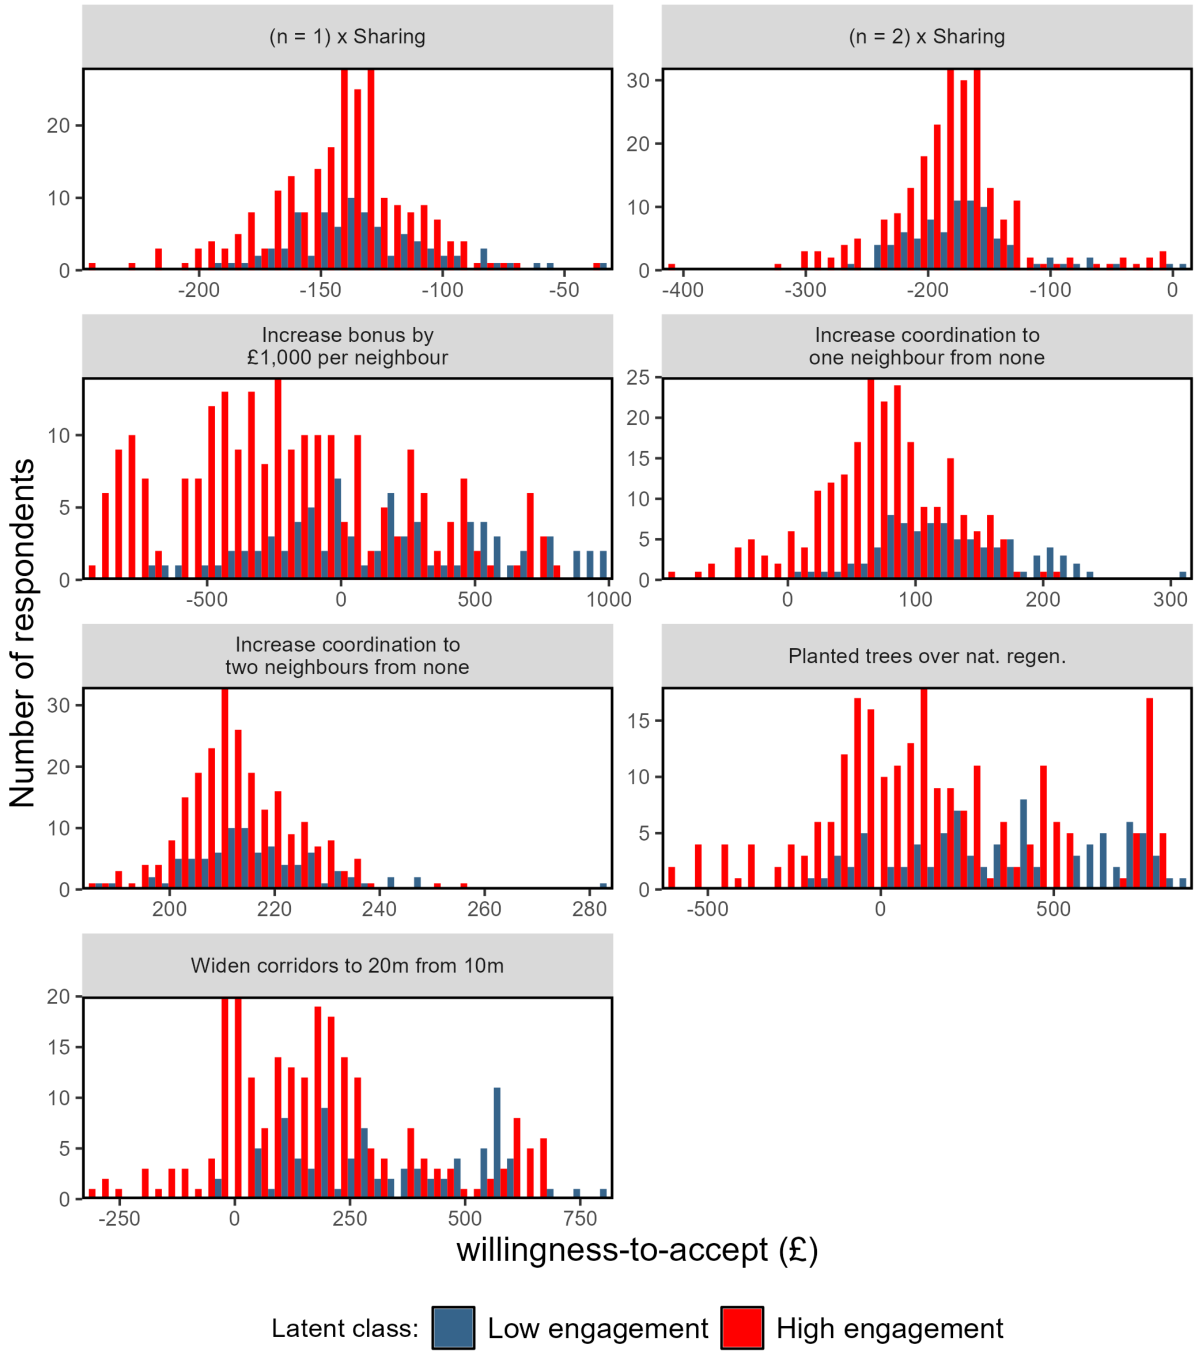

Figure 5.9: *Farm-specific monetary values for corridor scheme attributes estimated using a mixed logit model*

Testing Hypothesis I: I reject the null for hypothesis I, which states that respondents correctly perceive that the combined corridor length times width equals area set aside, if $\beta_{w=20} < 0$. The taste parameter is negative and significant for both classes. Directionally, the results lend support for hypothesis I. However, the magnitude of monetary values per hectare of corridors does not match exactly the value of contiguous features. This may result from perceived cost differences between corridors and contiguous features.

Compared with no coordination between neighbours, respondents demand on average £75-£100 per 100 meters to coordinate with one neighbour, and £210 extra per 100 meters to coordinate with two neighbours. The demanded compensation is skewed narrowly higher in the low engagement group. This suggests an approximately constant marginal cost of coordination for the average farmer. The interaction between an increase in the number of collaborators in coordination n and stated regular sharing of farm equipment is negative for $n = 1$ and $n = 2$. Farmers who state that they currently collaborate with neighbours demand on average £100-£150 less to go from no coordination to coordination with one neighbour, than do non-collaborative farmers. The "premium" placed on collaborating with two neighbours is approximately £150-£170 lower among farmers who regularly collaborate.

Testing Hypothesis II: For the low engagement class, 12% of draws match the joint equality of the null hypothesis. For the high engagement class, the proportion is 0.1%. For the high engagement class, I reject the null for hypothesis II, which states that there is a positive marginal coordination cost $C'(n)$, if $\beta_{n=2} < \beta_{n=1} < 0$. I am able to reject the null, and confirm that farmers account for a positive marginal coordination cost involved in coordinating connectivity improvements with neighbouring farmers.

Testing Hypothesis III: I partially reject the null for hypothesis III, which states that collaborative farmers are those that face lower marginal coordination costs and therefore are more willing to coordinate. The interaction parameters are greater than 0 (9% of draws for class I, 2% of draws for class II) which means that farmers who regularly share equipment are less sensitive to greater coordination requirements. However, I fail to reject the null for the joint inequality $\beta_{(n=2)\times sharing} > \beta_{(n=1)\times sharing} > 0$ (66% of draws). This means that I cannot reject that $\beta_{(n=2)\times SHARING} > \beta_{(n=1)\times SHARING}$.

5.6.3 Cost-effectiveness analysis of habitat connectivity

Figure 5.10 shows the change in average visitation rates attributed to the implementation of each type of habitat creation scheme. Changes in visitation rates are displayed for planted trees and natural regeneration, spatially arranged as field-edge corridors, in-field corridors, in-field islands, and singular contiguous patch. I calculate changes in visitation rates per m^2 of natural features created by farmers. The x-axis represents the gap between corridors in meters. For each gap size and farm, the contiguous patch, field-edge corridors, and in-field corridors have been placed such that the combined amount of land retired for natural features is identical. Larger gaps imply a smaller amount of natural features, L_{NF}, as well as a lower habitat connectivity.

The amount of land use change has been determined such that it can be achieved with a £1000 payment per farmer per year. This means that a scheme with planted trees will have less land converted into natural features than does a scheme with natural regeneration. Similarly, a £1000 payment affords fewer features if they are placed in-field compared to placement along field edges. This is because farmers in the choice experiments demand more compensation for these features, which are perceived to be more expensive or disruptive to create and maintain. Simulating uptake based on an equal payment allows me to compare the cost-effectiveness of the different schemes.

For pollinator visits to broad- and field beans, corridors along field-edges are by far the most cost-effective solution, increasing farm-wide visitation rates by on average 4% with natural regeneration and 1% with planted fruit trees that also provide flower resources. This is expected as placing corridors along field-edges is significantly cheaper than in-field, with only a modest penalty on connectivity at narrow gaps. After controlling for cost, effects on visitation rates are largely independent of the gap between natural features. The exception is natural regeneration features arranged as evenly distributed 100m^2 islands, where increasing the gap to 800 meters improves the cost-effectiveness of the scheme from by a factor of five to six. The cost-effectiveness comparisons are comparable for economic grassland, while the magnitude of visitation improvement is lower in the range of 0.2% to 1%. The economic value of flower pollination on grassland is ultimately negligible, as it is used for grazing.

For visits to oilseed rape, field-edge corridors remain the most cost-effective scheme overall. Unlike beans, the effect of in-field features on oilseed pollination depends meaningfully on whether features are natural regeneration or fruit trees. £1000 spent on in-field corridors or islands of natural regeneration results in reduced visitation rates. In their validation of the `poll4pop` model, Häussler et al. (2017) find that land use heterogeneity has a negative impact on oilseed pollination. Intersecting the oilseed fields with corridors and islands of natural regeneration increases the heterogeneity of the field, which may explain the observed negative effect.

Figure 5.11 reproduces 5.10 using the upper quartile of farms in terms of improved visitation rates. In this subsample, I observe a clear cost-effectiveness advantage of in-field islands over corridors as the gap between islands increases. The combined area of natural features scales significantly more with the feature gaps when configured in the form of islands. This is because increases in the gap between islands reduce the available area for features in two directions while corridors are only constrained in one direction. Therefore, the required cost of islands declines significantly compared to corridors as the gap widens. Figure 5.13 illustrates how the upper quartile of farms differ from the average in terms of land use. Farms where the schemes produce the greatest improvements to crop visitation have less grassland cover and more cereal fields. I attribute the higher cost-effectiveness of in-field islands within this subgroup to the removal of cereal fields, which score very poorly both for floral resources and nesting resources within the `poll4pop` model.

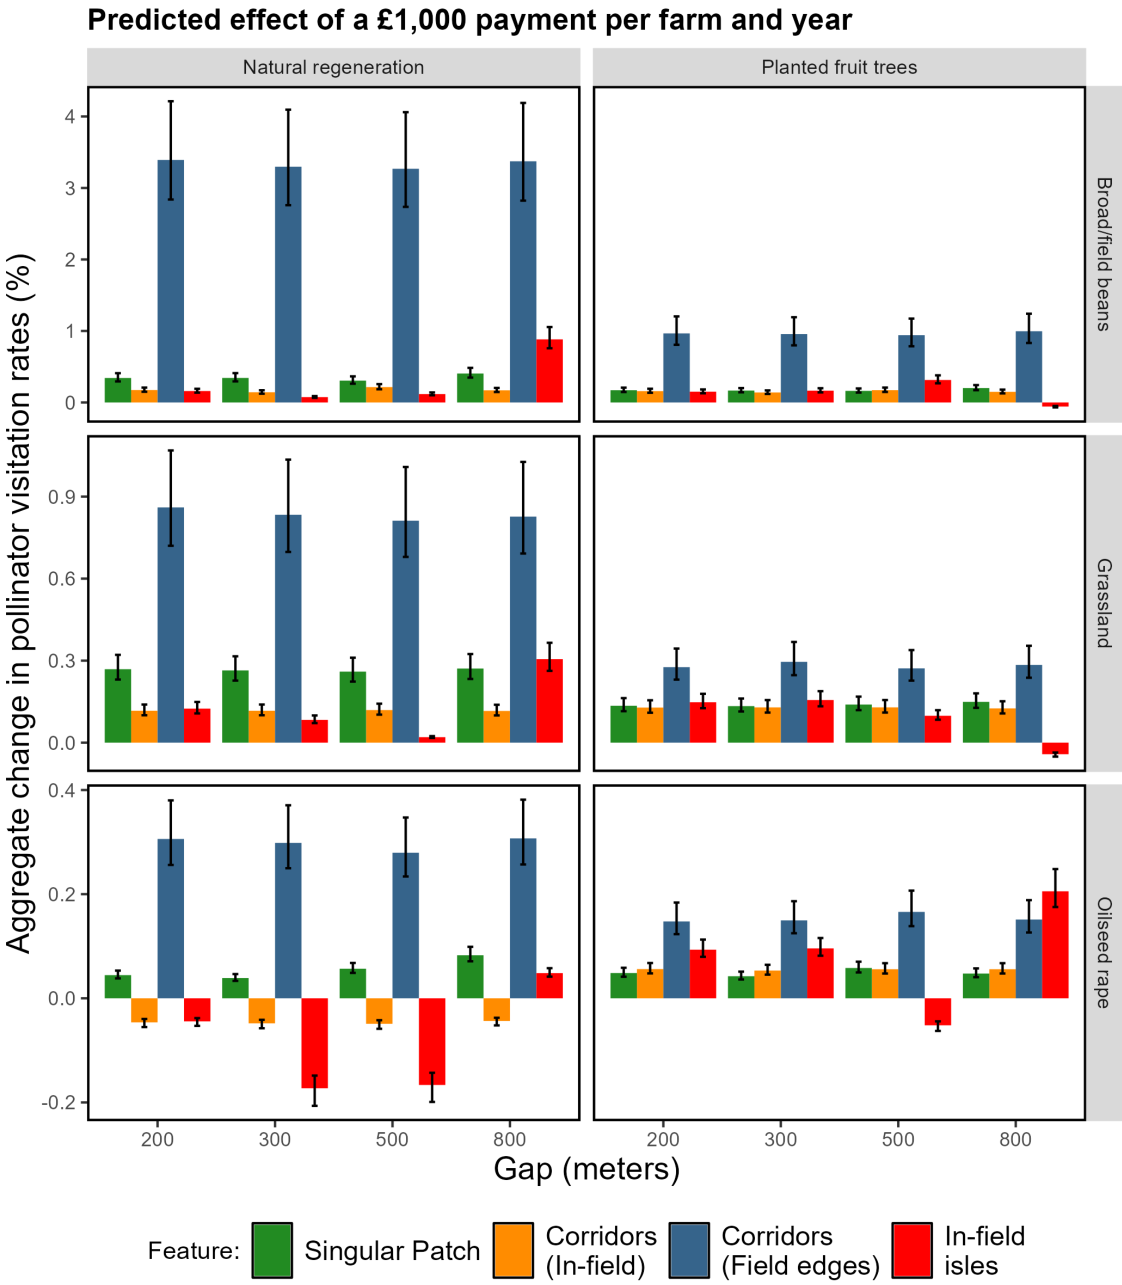

Figure 5.10: *Average aggregate change in pollinator visitation rates for three economic cover crops per £1000 payment per farm and year. All farms. Changes are reported by natural feature type and spatial configuration. The percentage change in visitation is reported per m^2 of natural features created. The x-axis denotes the gap between corridors.*

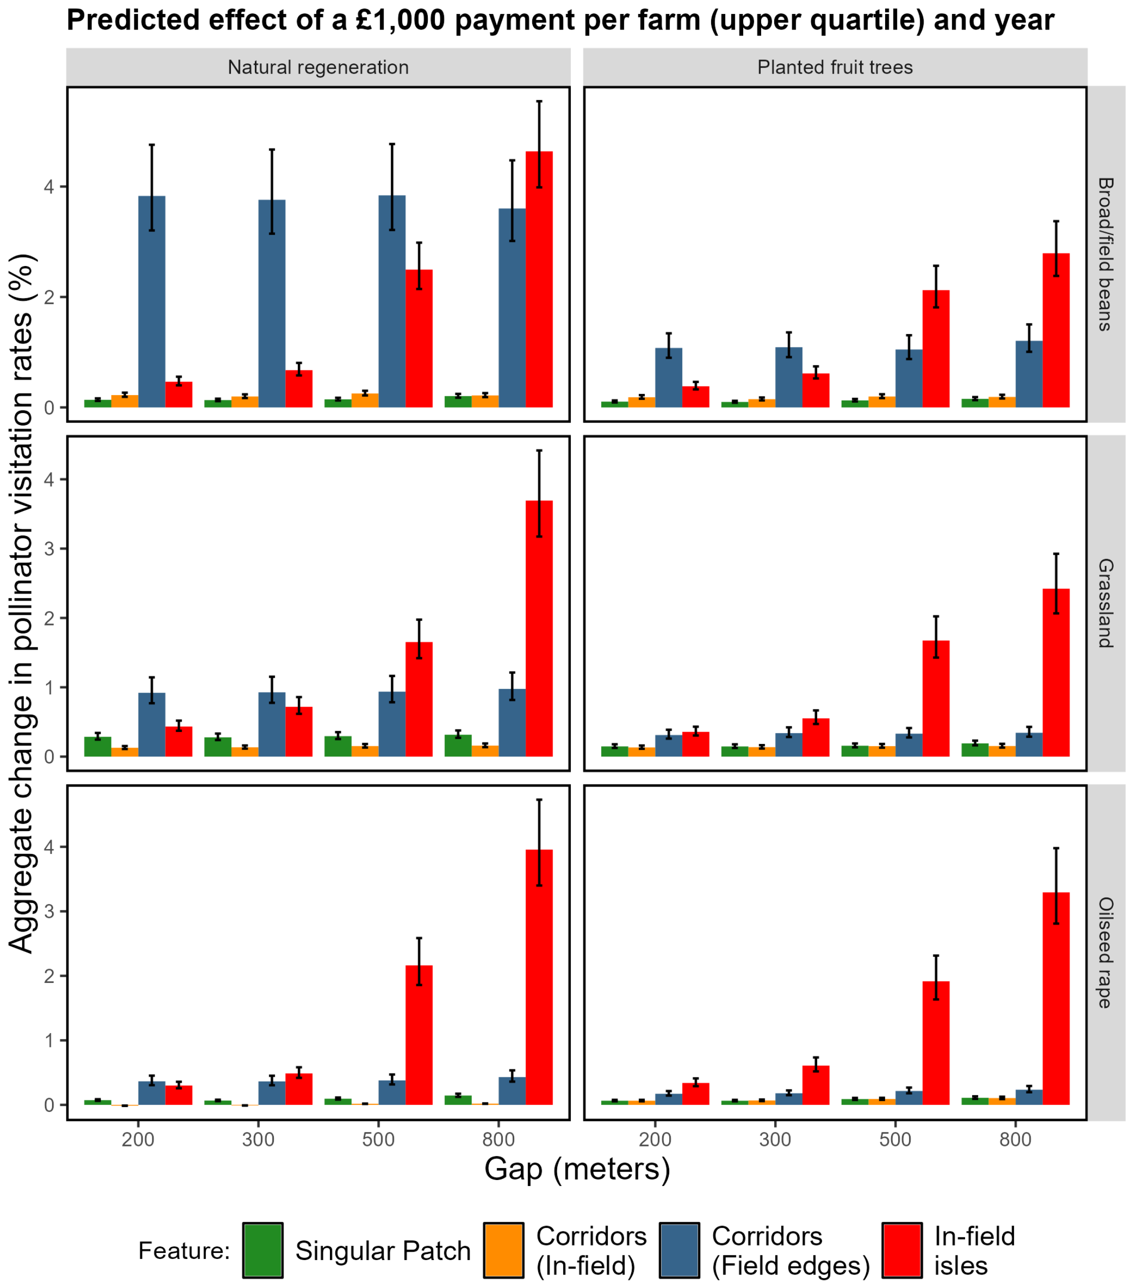

Figure 5.11: *Average aggregate change in pollinator visitation rates for three economic cover crops. Upper quartile of farms. Changes are reported by natural feature type and spatial configuration. The percentage change in visitation is reported per m^2 of natural features created. The x-axis denotes the gap between corridors.*

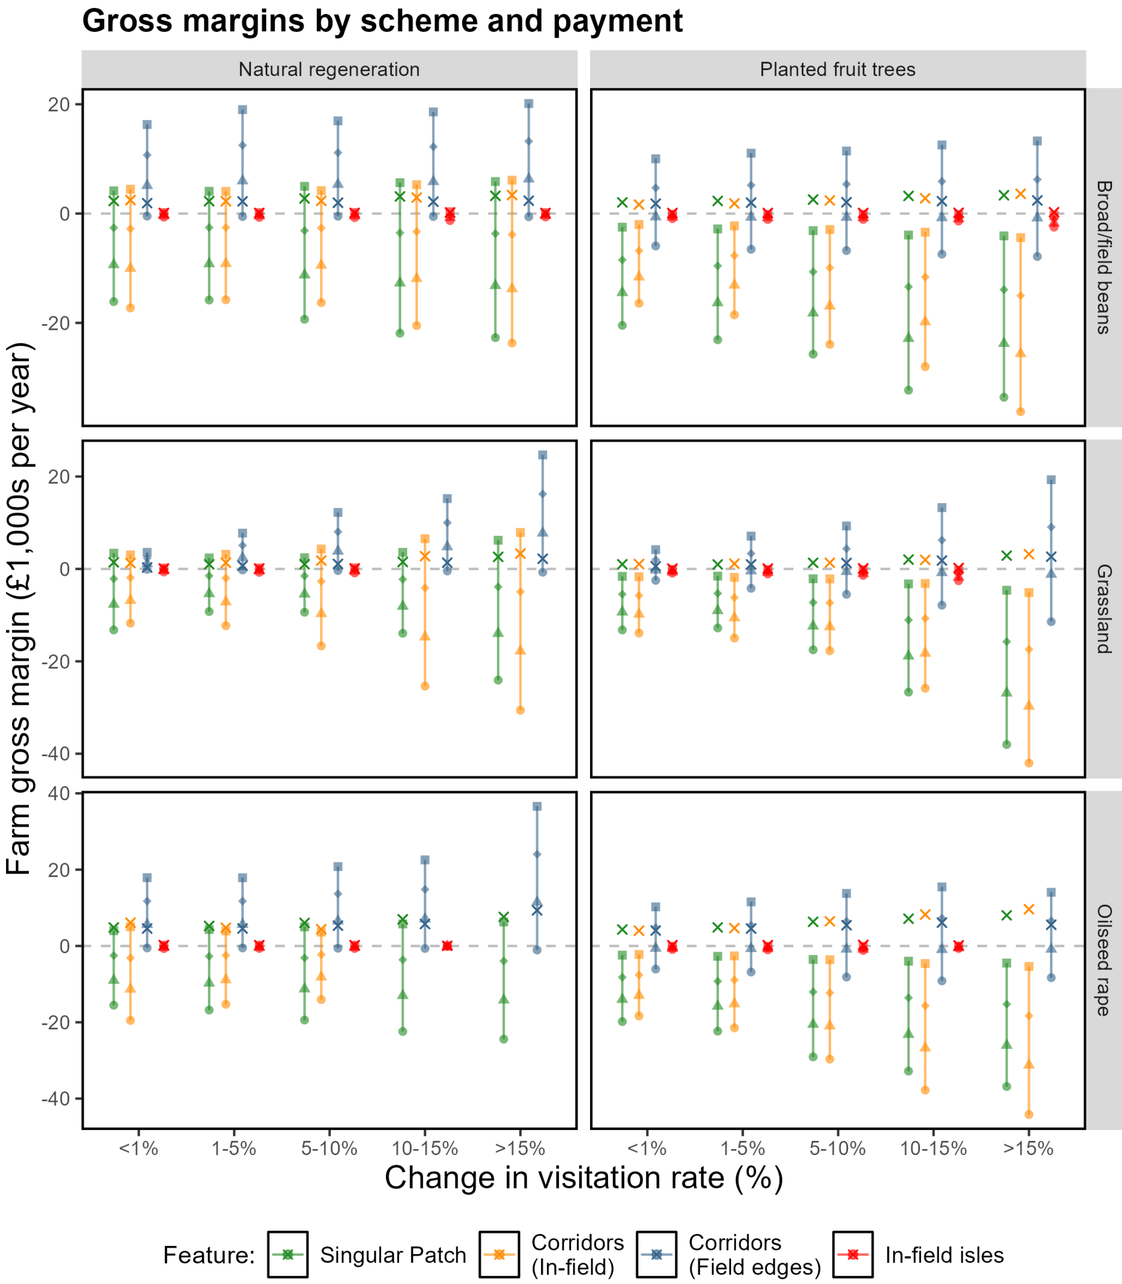

Figure 5.12: *Average farm gross margins (y-axis) resulting from sufficient amounts of natural features to produce visitation increases from 1% to 15% (x-axis). Margins were based on WTA from DCE I and an assumed annual payment of £2000/ha (○), £4000/ha (△), £6000/ha (◇), or £8000/ha (□). The × symbol denotes the 2022/23 gross margin for each land use class in the Farm Accounts for England (Department for Environment, Food and Rural Affairs, n.d.)*

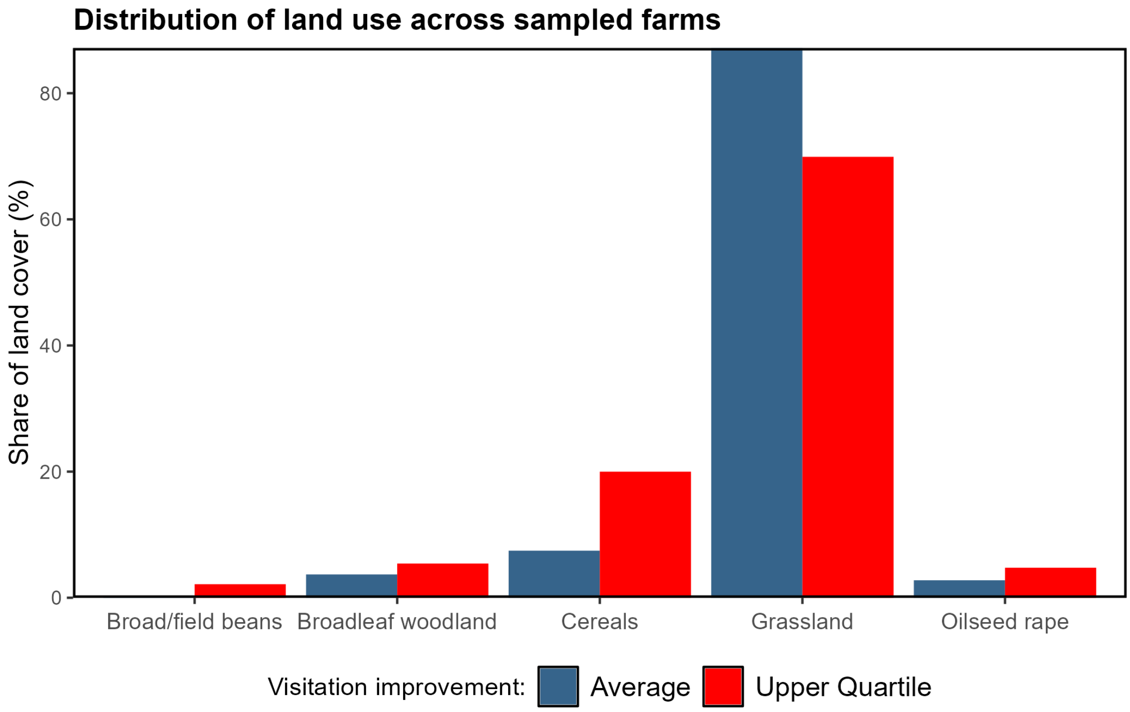

Figure 5.13: *Differences in land cover between the upper quartile of farms in terms of economic crop visitation improvements, and the average farm*

Testing hypothesis IV: The directional analysis of the connectivity insensitivity ratio ϕ in section 5.5.2 showed that ϕ can be negative, in which case more area devoted to natural features has a negative impact on pollination of economic crops such as oilseed rape and beans. I confirm that this is the case when the pre-treatment land use mix features a moderate level of flower resources. Natural features are unambiguously positive when the pre-treatment fields are low in resources.

5.7 Limitations

The economic model developed in section 5.3 makes an additional prediction of interest: When $\phi > 0$, farmers who grow pollinator dependent crops can benefit economically from pollination services. Because this research has established that, depending on the species, pollinators can forage up to 500 meters from their nests, pollinator-dependent farms benefit from natural features created on neighbouring farms. The model predicts that such farmers will be more inclined to engage in coordination, as the coordination bonus will incentivise their neighbour(s) to participate. This hypothesis can be tested by interacting the taste parameter for coordination with a variable indicating whether or not a respondent relies on pollinator-dependent crops.

I do not attempt in this paper to provide a conclusive answer to this question. The reason is that the survey was designed before this particular piece of the model had been derived, and therefore the survey did not give respondents sufficient information about pollination services and pollinator dependent crops. A direction for future research is to collect more detailed crop cover data on the farm level, and repeat the discrete choice experiment with an information treatment educating respondents about pollination services.

This research studies the value of connectivity improvements for pollination services specifically. The reported limited benefit from connecting natural features is attributed to the biology of insect pollinators in particular. Foraging distances of up to 500 meters reduce the harms from resource fragmentation in agricultural landscapes. As Correa Ayram et al., 2016's review of the habitat connectivity literature illustrates, a diverse range of species are affected. For example, 41% of studies surveyed focus on mammals compared to 8% that focus on insects. Affected species may also be of interest to policymakers and the public, for example out of conservation concerns or different ecosystem services (Hanley & Perrings, 2019). There is scope to expand the research in this chapter by incorporating species distribution models for other species of interest and to aggregate potential benefits.

5.8 Discussion

This research adds to a growing literature which studies the effect of ELM schemes on pollination services (Berg et al., 2019; Image et al., 2022, 2023; Kleftodimos et al., 2021). The current literature is primarily situated within ecology and agricultural sciences, although Kleftodimos et al. (2021) is a pioneering attempt at bringing the ecology and economics together. My work adds to this strand of the literature by incorporating a spatially explicit pollinator visitation model with the economic model. The model allows me to simulate the visitation effects from hypothetical ELM schemes at a very high (10m) resolution. As a result, this work contributes an early and comprehensive cost-effectiveness analysis of connectivity improvements via ELM schemes.

This chapter has shown that perceived coordination costs are a significant barrier to uptake of ELM schemes that require farmers to coordinate with neighbours to create connected natural features. It identifies a class of farmers whose coordination costs are lower, evidenced not only by a stated enthusiasm for collaboration, when compensated for it, but a revealed willingness to engage with neighbours and share farm equipment in the past. It is plausible that these differences are not only caused by differences in personality, such as agreeableness, but have grown over time due to growing familiarity and trust. Riley et al. (2018) state that whilst working relations between UK farmers are often collegiate, and in places collective, several watershed events over past decades have led to a shift from community-level to process-based (peer-to-peer) trust and a move toward land management being depicted as a squarely individual rather than collective issue. Against such a backdrop, environmental regulators may want to incentivise increased collaboration between neighbouring farms, anticipating that coordination costs will come down over time.

While I do not dispute this line of reasoning, results from this research call into question the value of farm-farm coordination, at least so far as it pertains to improving pollination services. When incorporating the ecology of three important pollinator species, the interdisciplinary cost-effectiveness analysis reveals that added habitat connectivity via coordination does not offer the most cost-effective outcomes. Given current marginal coordination costs, connecting features to reduce gaps between them in a 4 km^2 landscape does not offer more crop pollination for a given government payment. In fact, in landscapes where the share of crop cover that is unsuited to pollinators (such as cereals and grains) is high, the most cost-effective ELM project is disconnected, in-field islands as far as 800 meters apart.

Some caveats to these results should nonetheless be acknowledged. The first is that although the visitation modelling is based on a sample of over 300 representative landscapes at the surveyed farms, land use diversity was generally high. All but one land use class covered no more than half of the area, and significantly less for the majority of farms. As shown in figure 5.4, only grassland (in this context for grazing) dominated the landscape at some farms. Hence the effect of coordination on visitation rates may be greater on farms that are larger and more of a monoculture. In addition, as discussed in section 5.7, there are other species that may benefit significantly more from connectivity and that the public has an interest in protecting.

While these caveats are important for defining the boundaries of my findings, the results are likely to be generalisable to other UK upland farming systems with similarly diverse landscapes. The conclusion that untargeted investment in connectivity may be inefficient holds important policy implications for these specific, yet common, agricultural environments. The core finding is not that connectivity is "bad", but rather that its effectiveness is highly context-dependent, and my results highlight the conditions under which alternative interventions may provide greater value for money. This research argues that policymakers may want to advertise the benefits of disconnected islands, that may look disruptive for farmers but make very small demands in terms of retiring productive farmland. The analysis in chapter 4 also reveals that these in-field islands is among the most cost effective schemes for the purpose of reducing flood risk.

Chapter 6

Conclusion

This thesis has contributed to the literature on spatially targeted environmental policy in several ways. Chapter 2 presents the first incorporation of pollution dispersion modelling within a difference-in-differences analysis of a cap-and-trade scheme. The resulting ambient pollution maps allow me to estimate the causal effect of the scheme in terms of *cross-border pollution*. In turn, this allows me to address the first research question posed in the introduction in chapter 1: How do firms respond to spatially differentiated compliance costs and what is the resulting environmental impact? My research finds lower abatement among regulated power plants that export SO_2 amounting to at least 1% of ambient air quality standards (NAAQS) outside the state. The average reduction in abatement compared to non-exporters is approximately 20%. Crucially, this result extends beyond those plants that actually export pollution outside of the state where they are regulated. A lower effect of the tightened emissions cap is observed among plants that are merely close to the state border. This implies that, for at least some plants, energy firms' beliefs about compliance enforcement (via the 'good neighbour' provisions) may themselves contribute to lower abatement. This result adds to previous work, e.g. by Fowlie et al. (2012), Heo et al. (2023) and Cai et al. (2016), by explicitly evaluating the cross-border externality and putting it in context of the firms' response. In a spatially targeted market, with trading ratios reflecting the cross-border pollution risk, these plants could receive proportionally greater revenue from selling permits. Such a system may incentivise enough additional abatement to compensate for the observed treatment heterogeneity.

Chapter 4 contributes the first cost-effectiveness analysis of NFM schemes which integrates cost estimates from hypothetical DCEs and hydrological connectivity modelling. It indicates that the majority of farmers would be open to enrol in the schemes if compensation was in the region of £$200 - 500$ per annum for $^1/_{20} - ^1/_{10}$ hectares of NFM features. I leverage the hydrological model to simulate the benefits side of the cost-effectiveness comparison. I find that NFM features created in typical English agricultural landscapes could result in measurable reductions in water runoff. Spatially targeted trading in NFM contracts result in significantly better cost-effectiveness. In particular, contiguous patches or small in-field islands of planted trees or natural regeneration are advantageous. In simulations, these schemes reduce flood risk by $10 - 20\%$ without trading and by $20 - 40\%$ with trading. The second research question posed in chapter 1 asks whether spatially targeted trading in NFM contracts can facilitate more cost-effective mitigation of flood risk. This research answers affirmatively, and identifies that types of schemes that offers the best 'bang' for the taxpayers' buck. However, this research also identifies some important barriers to uptake. By including transaction costs in the DCEs, I address the third research question: How do transaction costs impact the feasibility of a hypothetical market in ELM obligations? The research finds that, contrary to previous evidence from trade in pollution permits, transaction costs are likely to be noticeable barriers in a hypothetical market for NFM contracts. However, with transaction costs in the region of $5 - 10\%$ of base payments, the required compensation is dwarfed by the aggregate cost savings from trading.

Chapter 5 contributes to the growing literature on the effect of ELM schemes on pollination services. While the current literature is primarily situated within ecology and agricultural sciences, this chapter adds to this strand of the literature by incorporating a spatially explicit pollinator visitation model with the economic model. The hypothetical DCE used to estimate farmers' willingness to coordinate is augmented by also surveying their professional connections with neighbours. The fourth research question asks: What role does social or professional networks play in farmers' perceived barriers to coordination? This research does report some evidence that farmers who have a professional relationship with their neighbours are more likely to opt to engage in coordination of connected habitats. However, this is among the weakest and most uncertain results in this thesis. In general, coordination costs are a barrier to this type of scheme. In this context, the simulation of pollination services for a broad range of spatial configurations of natural features in informative. I show that small, evenly distributed in-field islands of natural features deliver the most cost-effective improvement in pollinator visits to oilseed rape and field beans. This particular scheme can be implemented on a larger scale and demand very little coordination between farmers. This research informs policymakers that although coordination costs are substantial, pollination services could be improved while bypassing this barrier.

Finally, the fifth research question asks: How can spatially explicit simulation models contribute to cost-effectiveness analysis of spatially targeted schemes? In this thesis, I have integrated three different spatially explicit simulation models (Gaussian air pollution dispersion, hydrological connectivity, and pollinator visitation) to predict the benefits from several environmental policies. In particular, for chapters 4 and 5, I develop a common algorithm to simulate counterfactual landscapes under a common set of ELM schemes. This allows me to conduct the first *multifunctional* cost-effectiveness analysis in terms of both flood risk reduction and pollination service provision. This research enables policymakers in the UK and across Europe to compare potential schemes in terms of joint benefits. In summary, I demonstrate the value in integrating hypothetical DCEs with different spatial simulation models to conduct multifunctional cost-effectiveness analyses. Benchmarking alternative environmental schemes in terms of cost-effectiveness can account for multiple benefits. These can be weighted according to policy priorities.

Appendices

Derivation of demand for NFM in the targeted trading regime

The Lagrangian for farmer q's cost minimisation problem is shown in equation (6.1). The transaction cost T takes the value $(1+\tau)$ when q on net is buying out of their NFM obligation and $(1-\tau)$ when q accepts payment to take up additional NFM.

$$\begin{aligned}\mathcal{L} = p_X X + c_{NF} L_{NF} + T\pi\left(\tilde{L}_{NF} - r_q L_{NF}\right) - \\ \mu_1\left(\bar{Y} - X^{\alpha} L_{AG}^{\beta}\right) - \\ \mu_2\left(\bar{L} - L_{AG} - L_{NF}\right)\end{aligned} \tag{6.1}$$

When q chooses their levels of X, L_{AG} and L_{NF}, the first-order KKT conditions are shown in equations (6.2) through (6.4):

$$[X]: \quad p_X + \mu_1 \alpha X^{\alpha-1} {L_{AG}}^{\beta} = 0 \tag{6.2}$$

$$[L_{AG}]: \quad \mu_1 \beta X^{\alpha} {L_{AG}}^{\beta-1} + \mu_2 = 0 \tag{6.3}$$

$$[L_{NF}]: \quad c_{NF} - T\pi r_q + \mu_2 = 0 \tag{6.4}$$

By rearranging (6.2), recover the function for μ_1 in terms of the variables that drive agricultural production cost p_X, X, L_{AG}. The equality (6.5) is then substituted into (6.3) to solve for μ_2 in terms of the net costs of agricultural output and NFM features. The maximally simplified function for μ_2 substituted into (6.4) is displayed in (6.7).

$$\mu_1 = -\frac{p_X}{\alpha X^{\alpha-1} L^{\beta}} \tag{6.5}$$

$$\mu_2 = \frac{p_X}{\alpha X^{\alpha-1} L^{\beta}} \beta X^{\alpha} L_{AG}^{\beta-1} \tag{6.6}$$

$$\mu_2 = p_X \frac{\beta X}{\alpha L_{AG}} = c_{NF} - T\pi r_q \tag{6.7}$$

From equation (6.7) solve for X (6.8) and substitute into the production function to recover the cost-minimising agricultural land inputs in terms of residual output demand and costs (6.9). Equation (6.10) shows the properly simplified form of (6.9).

$$X = \frac{\alpha}{\beta} \frac{c_{NF} - T\pi r_q}{p_X} L_{AG} \tag{6.8}$$

$$\left(\frac{\alpha}{\beta} \frac{c_{NF} - T\pi r_q}{p_X} L_{AG} \right)^{\alpha} L_{AG}^{\beta} = \bar{Y} \tag{6.9}$$

$$L_{AG}^{*} = \left(\frac{\frac{\beta}{\alpha} p_X \bar{Y}^{1/\alpha}}{c_{NF} - T\pi r_q} \right)^{\frac{\alpha}{\alpha+\beta}} \tag{6.10}$$

The cost-minimising level of L_{AG} is substituted into the land endowment constraint to recover the demand for L_{NF}.

$$L^*_{NF} = L_{AG} - \left(\frac{\frac{\beta}{\alpha} p_X \bar{Y}^{1/\alpha}}{c_{NF} - T\pi r_q}\right)^{\frac{\alpha}{\alpha+\beta}} \tag{6.11}$$

Differentiating the demand for natural features with respect to the payment rate yields the marginal demand:

$$\frac{\partial L^*_{NF}}{\partial \pi} = -\left(\frac{\alpha}{\alpha+\beta}\right) r_q \frac{\left(\frac{\beta/\alpha p_x \bar{Y}^{1/\alpha}}{c_{NF}-T\pi r_q}\right)^{\frac{\alpha}{\alpha+\beta}}}{c_{NF} - T\pi r_q} \tag{6.12}$$

Survey data processing

This script processes the survey responses downloaded from Qualtrics into a format which is compliant with the R package Apollo (Hess & Palma, 2019). The script also geocodes the survey data based on postcodes volunteered by respondents.

```
### Clear memory
rm(list = ls())
#install.packages("dplyr")
#install.packages("reshape2")
#install.packages("writexl")
#install.packages("sp")
#install.packages("sf")
#install.packages("geosphere")
### Load libraries
library(dplyr)
library(reshape2)
library(writexl)
library(sp)
library(geosphere)
library(sf)
# -----------------------------------
```

```r
# User guidance: This script accepts survey data in the raw Qualtrics results format. However it must first be cleaned manually by:
# 1) replacing missing values with NA
# 2) correcting typos in survey respondent inputs
# 3) ensuring numeric data is coded as such
# I recommend that this is done in Excel prior to loading the data file into R, as the following script will assume correct data entry and types.
# In addition, this script incorporates additional data from the Ordnance Survey, Defra and Natural England, including a) areas (ha) of fruit orchards within farm perimeters,
# b) areas (ha) of oilseed rape fields within farm perimeters, c) NFM priority areas, d) post offices within respondent's post district, e) pubs within respondent's district.
# Users of this code should download the data and adjust location paths to the files accordingly.
# ------------------------------------
# functions used
replace <- function (x) {ifelse(x == "", "I want neither A nor B", x)}
# load survey answers
survey = read.csv("survey_220523.csv") %>%
# keep only respondents who have completed the survey and consented to have their answers recorded
filter(Consent == 'I CONSENT' & CE3Task8_1 %in% c('Option A','Option B','I want neither A nor B')) %>%
# generate respondent identifier variable and geographic identifiers
mutate(id = c(1:nrow(.)), postcode = toupper(gsub(" ", "", Form_4)), n_pubs = NA, n_post = NA, Q6 = as.numeric(Q6)) %>%
mutate(postcode = ifelse(postcode == "", NA, postcode))
```

```r
# load UK postcodes with coordinates from the Ordnance Survey Code-Point Open dataset
filenames <- list.files("postcodes", pattern="*.csv", full.names=TRUE)
# initialize a data frame to collect coordinates matching respondent postcodes
postcode_coords <- data.frame(postcode = character(), eastings = integer(), northings = integer())
for (i in 1:length(filenames)) {
# for each county, bind coordinates by postcode contained in the survey responses data set
postcodes <- read.csv(filenames[i])
names(postcodes) <- c("postcode","quality","eastings","northings","country_code","NHS_regional_HA_code","NHS_HA_code","admin_county_code","admin_district_code","admin_ward_code")
postcodes <- postcodes %>% mutate(postcode = toupper(gsub(" ", "", postcode))) %>% select(postcode, eastings, northings)
# bind to initialized data frame
postcode_coords <- rbind(postcode_coords, postcodes[postcodes$postcode %in% survey$postcode, ])
}
# merge coordinates into survey data by respondent postcode
survey <- survey %>% left_join(postcode_coords, by = "postcode")
# create a spatial points data frame geo-locating respondents
survey_sf <- survey %>%
filter(!is.na(eastings)) %>%
mutate(Q6 = ifelse(is.na(Q6), mean(survey[!is.na(survey$Q6),]$Q6), Q6)) %>%
st_as_sf(coords = c("eastings","northings"))
survey_sf_buff <- st_buffer(survey_sf, sqrt(survey_sf$Q6*1e+4)/2)
```

```r
# load Crop Map of England (CROME) data
crome_list = list('list of CROME shapefiles')
names(crome_list) <- c("Durham","North Yorkshire","South Yorkshire","West Yorkshire","East Riding of Yorkshire","Lancashire","Lincolnshire","Cumbria","Northhumberland","Tyne and Wear","Cheshire")
oilseed_dat <- data.frame(id = character(), oilseed = double(), lucode = character())
for (i in 1:length(crome_list)) {
# for each county
crome_sf <- st_read(crome_list[[i]])
# collect all cells classed oilseed rape
st_crs(survey_sf_buff) <- st_crs(crome_sf)
print(paste(names(crome_list)[i], "CROME data loaded successfully."))
# find CROME parcels of oilseed that fall within farm polygon by farm ID
intersect <- st_join(crome_sf, survey_sf_buff["id"], join = st_within) %>%
st_drop_geometry() %>%
filter(!is.na(id)) %>%
group_by(id) %>%
summarise(oilseed = sum(st_area_sh*1e-4))
# collect number of cells and ID
oilseed_dat <- rbind(oilseed_dat, intersect)
# print progress to console
print(paste(names(crome_list)[i], " county completed.", sep=""))
}
# load shapefile of orchards from Natural England
orchards_sf <- st_read("Traditional_Orchards_HAP_(England)___Natural_England.shp")
# select only pollinated fruit orchards
```

```
orchards_sf <- orchards_sf[!is.na(orchards_sf$Apple) | !is.na(orchards_sf$Pear) | !is.na(orchards_sf$Cherry) | !is.na(orchards_sf$Plum), ]
st_crs(survey_sf_buff) = st_crs(orchards_sf)
# find NE fruit orchard polygons that fall within farm polygon by farm ID
fruit_dat <- st_join(orchards_sf, survey_sf_buff["id"], join = st_within) %>%
st_drop_geometry() %>%
filter(!is.na(id)) %>%
group_by(id) %>%
summarise(fruits = sum(Area_Ha*1e-4))
# load shapefile of NFM priority areas (Defra 2020)
nfm_prio_sf <- st_read("Spatial_Prioritisation_of_Catchments_Suitable_for_Using_NFM.shp")
st_crs(survey_sf) = st_crs(nfm_prio_sf)
# join survey with NFM priority areas
survey_sf <- st_join(survey_sf, nfm_prio_sf, join = st_within)
# load shapefile of flood risk areas (Defra)
# level 3: >1% risk
floodmap3_sf <- st_read("Flood_Map_for[...]Sea_Flood_Zone_3.shp")
st_crs(floodmap3_sf) = st_crs(survey_sf)
floodmap3_sf = floodmap3_sf %>%
select(layer, geometry) %>%
rename(level3_risk = layer)
# level 2: 0.1-1% risk
floodmap2_sf <- st_read("Flood_Map_for_Planning_Rivers_and_Sea_Flood_Zone_2.shp")
st_crs(floodmap2_sf) = st_crs(survey_sf)
floodmap2_sf = floodmap2_sf %>%
```

```
select(layer, geometry) %>%
rename(level2_risk = layer)
# merge flood risk scores with survey data
survey_sf <- st_join(survey_sf, floodmap2_sf, join = st_within)
↪  %>%
st_join(floodmap3_sf, join = st_within)
# create dataframe
survey <- survey_sf %>%
st_drop_geometry() %>%
left_join(oilseed_dat, by = "id") %>%
left_join(fruit_dat, by = "id") %>%
left_join(postcode_coords, by = "postcode") %>%
mutate(fruits = ifelse(is.na(fruits), 0, fruits),
oilseed = ifelse(is.na(oilseed), 0, oilseed),
# generate a variable for the post code sector
sector = substr(postcode, 1, nchar(postcode)-2),
county = NA,
flood_risk = ifelse(!is.na(level3_risk), 2,
↪  ifelse(is.na(level3_risk) & !is.na(level2_risk), 1,
↪  0)))
# generate variable indicating number of pubs and post offices in
↪  respondent's postcode
for (i in 1:nrow(survey)) {
if (!is.na(survey$postcode[i])) {
        # resolution: postcode sector
        survey$n_pubs[i] =
        ↪  nrow(pubcoords[grepl(survey$sector[i],
        ↪  pubcoords$postcode), ])
        survey$n_post[i] =
        ↪  nrow(postcoords[grepl(survey$sector[i],
        ↪  postcoords$postcode), ])
        # get county
```

```r
        survey$county[i] = postcoords[grepl(survey$sector[i], postcoords$postcode), "county"][1]
}
}
survey = survey %>% rename(email = Form_5) %>% mutate(email = toupper(email)) %>% filter(email != "")
# load designs from Ngene
dgn1 = read.csv('design1.txt', sep = '\t')
dgn_wta = read.csv('design2_wta.txt', sep = '\t')
dgn_wtp = read.csv('design2_wtp.txt', sep = '\t')
dgn3 = read.csv('design3.txt', sep = '\t')
# collect column names for choice tasks
choices = list()
choices[[1]] = names(survey_full[ , grepl('CE1Task', names(survey_full))])
choices[[2]] = names(survey_full[ , grepl('WTATask', names(survey_full))])
choices[[3]] = names(survey_full[ , grepl('WTPTask', names(survey_full))])
choices[[4]] = names(survey_full[ , grepl('CE3Task', names(survey_full))])
# collect column names for control questions (same across CEs)
controls = names(survey_full[ , grepl('Q', names(survey_full)) | names(survey_full) %in% c("fruits","oilseed","ea_nfm_pri","flood_risk","eastings","northings","county")])
times = list()
times[[1]] = names(survey_full[, grepl('Page.Submit', names(survey_full))])[1:8]
times[[2]] = names(survey_full[, grepl('Page.Submit', names(survey_full))])[9:14]
times[[3]] = names(survey_full[, grepl('Page.Submit', names(survey_full))])[15:20]
```

```
times[[4]] = names(survey_full[, grepl('Page.Submit', names(survey_full))])[21:28]
# collect attribute names
attributes = list()
attributes[[1]] = names(dgn1)[2:12]
attributes[[2]] = names(dgn_wta)[2:8]
attributes[[3]] = names(dgn_wtp)[2:8]
attributes[[4]] = names(dgn3)[2:12]
# collect designs from CE 1, CE 2 (WTA), CE 2 (WTP) and CE 3
designs = list()
designs[[1]] = dgn1 %>% rename(task = Design) %>% mutate(task = choices[[1]])
designs[[2]] = dgn_wta %>% rename(task = Design) %>% mutate(task = choices[[2]])
designs[[3]] = dgn_wtp %>% rename(task = Design) %>% mutate(task = choices[[3]])
designs[[4]] = dgn3 %>% rename(task = Design) %>% mutate(task = choices[[4]])
# clean data and collect CEs in list
data = list()
# number of tasks per choice experiment
len = c(8,6,6,8)
attribute_names <- list(c("task","c1.type","c1.loc","c1.qual","c1.area","c1.pay","c2.type","c2.loc","c2.qual","c2.area","c2.pay"),
c("task","c1.ratio","c1.fee","c1.pay","c2.ratio","c2.fee","c2.pay"),
c("task","c1.ratio","c1.fee","c1.pay","c2.ratio","c2.fee","c2.pay"),
c("task","c1.type","c1.coord","c1.width","c1.bonus","c1.pay","c2.type","c2.coord","c2.width","c2.bonus","c2.pay"))
# for each choice experiment i:
for (i in 1:4) {
```

```
data[[i]] = survey_full %>%
# select responses to choices and control questions
select(c(id, choices[[i]], controls, times[[i]])) %>%
# replace missing responses with status quo
mutate_at(choices[[i]], funs(replace(.))) %>%
# only keep observations where choice tasks are answered
filter_at(vars(choices[[i]]), any_vars(. %in% c('Option A',
↪  'Option B', 'I want neither A nor B'))) %>%
# transform data into long format
melt(id = c('id', controls, times[[i]])) %>%
rename(task = variable, choice = value) %>%
arrange(id) %>%
# add attribute levels to each choice task
inner_join(designs[[i]], by = 'task') %>%
select(c(id, choice, attributes[[i]], controls, times[[i]])) %>%
rename_with(~ attribute_names[[i]], all_of(attributes[[i]])) %>%
# numeric data
mutate(Q2 = as.numeric(Q2),
Q5 = as.numeric(Q5),
Q6 = as.numeric(Q6),
Q7_1 = as.numeric(Q7_1),
Q11_1 = as.numeric(Q11_1),
Q11_2 = as.numeric(Q11_2),
Q11_3 = as.numeric(Q11_3),
Q11_4 = as.numeric(Q11_4),
Q11_5 = as.numeric(Q11_5),
Q11_6 = as.numeric(Q11_6),
Q11_7 = as.numeric(Q11_7),
Q12   = as.numeric(Q12),
Q15   = as.numeric(Q15),
Q17_2 = as.numeric(Q17_2)) %>%
mutate_at(times[[i]], as.numeric) %>%
# generate control variables
```

```r
# NOTE: as.numeric() transforms non-numeric formats to NA - ensure data type is correct to avoid data loss
mutate(female = ifelse(Q1 == 'Female', 1, 0), # is respondent female
age = 2023 - as.numeric(Q2), # respondent age
farm_age = 2023 - as.numeric(Q5), # age of farm
# educational attainment
edu = case_when(Q4 == 'Other vocational/technical training' ~ 0,
Q4 == 'GCSEs, O-levels or equivalent' ~ 1,
Q4 == 'College (A-levels or equivalent)' ~ 2,
Q4 == '3-year university degree' ~ 3,
Q4 == 'Postgraduate degree' ~ 4),
hectare = Q6, # farm size (ha)
owned = Q7_1, # % of land owned
primary = ifelse(Q8 == 'Yes', 1, 0), # is respondent's primary income from agriculture
cereals = Q11_1, # % used for each product
cropping = Q11_2,
grazing = Q11_3,
pigsbird = Q11_4,
horticult = Q11_5,
dairy = Q11_6,
other = Q11_7,
sum_total = ifelse(Q11_1 + Q11_2 + Q11_3 + Q11_4 + Q11_5 + Q11_6 + Q11_7 == 100, 1, 0),
n_tracts = Q12, # tracts of land farmed
aes = ifelse(Q13 == 'None', 0, 1), # does respondent currently participate in ELM scheme
# self rated community participation
social = case_when(Q14 == 'Much less than average' ~ 0,
Q14 == 'Less than average' ~ 1,
Q14 == 'About average' ~ 2,
Q14 == 'More than average' ~ 3,
```

```r
Q14 == 'Much more than average' ~ 4),
# number of neighbouring farms
boundary = as.numeric(Q15),
# does respondent share farm equipment with neighbours
sharing = ifelse(Q16 == "Yes", 1, 0),
ea_respect = Q17_2,
# fconcern = concern about flooding on the farm
fconcern = case_when(Q18 == 'Not concerned' ~ 0,
Q18 == 'Mostly not concerned' ~ 1,
Q18 == 'Unsure' ~ 2,
Q18 == 'Somewhat concerned' ~ 3,
Q18 == 'Very concerned' ~ 4),
# cconcern = concern about flooding in the catchment, surrounding communities
cconcern = case_when(Q19 == 'Not concerned' ~ 0,
Q19 == 'Mostly not concerned' ~ 1,
Q19 == 'Unsure' ~ 2,
Q19 == 'Somewhat concerned' ~ 3,
Q19 == 'Very concerned' ~ 4),
poll_dep = ifelse(fruits + oilseed > 0.5, 1, 0), # does respondent grow pollinator-dependent crops
# what is the flood risk management priority level of farm
nfm_prio = case_when(is.na(ea_nfm_pri) ~ NA,
ea_nfm_pri == 'Low' ~ 0,
ea_nfm_pri == 'Medium' ~ 1,
ea_nfm_pri == 'High' ~ 2),
choice = case_when(choice == 'Option A' ~ 1,
choice == 'Option B' ~ 2,
.default = 0)) %>%
group_by(id) %>%
mutate(serial_sq = ifelse(mean(choice) == 0, 1, 0),
resp_time = mean(c_across(times[[i]]))) %>%
ungroup()  %>%
```

```
select(-c(controls[1:29], times[[i]], ea_nfm_pri))
}
# save data
sheets = list("dce1" = data[[1]], "dce_wta" = data[[2]],
    "dce_wtp" = data[[3]], "dce3" = data[[4]])
write_xlsx(sheets, 'apollo_data_full.xlsx')                    # END
```

Run poll4pop

This script loops through a list of simulated landscapes to model a) crop pollinator visitation rates and b) the Hanski (habitat) connectivity index for each hypothetical landscape. Visitation rates are calculated using `poll4pop` (Häussler et al., 2017). Each simulated landscape represents a hypothetical spatial configuration / feature type of the ELM scheme.

```
rm(list = ls())

library(readxl)
library(plyr)
library(dplyr)
library(sf)
library(raster)
library("EBImage")
library(progress)
library(foreach)
library(doParallel)

# load external functions
source("kerncalc.R")
source("latfordisp.R")
source("./rawPoll4Pop/computeFloralNesting.R")
```

```r
source("./rawPoll4Pop/growth.func.R")
source("./rawPoll4Pop/runpoll_3seasons.R")
source("ci_index_fun.R")
source("corridor_fun.R")

# Load land use raster map
cc <- raster("lc.tif")
# load poll4pop parameters
load(file = ".\\data\\parameters.rda")
# declare the pollinator species
bees <- c(1,2,8)
names(bees) <- c("GNBumblebees","GNSolitaryBees","TNBumblebees")
# declare widths of habitat corridors / isles
widths <- c(10, 20)
# declare gaps between corridors / isles
gaps <- c(200, 300, 500, 800)
# declare type of ecological feature (natural regeneration and
↪   fruit trees)
feature_infield <- c(11, 34)
feature_edge <- c(24, 34)
# declare spatial structure for features
spatial <- c("isle","corrs","edges","mosaic","islands")
# load farm locations
dat = read_excel("apollo_data_full.xlsx", sheet = "dce3") %>%
dplyr::select(id, northings, eastings) %>%
dplyr::filter(!duplicated(id))
dat_sf <- st_as_sf(dat, coords = c("eastings","northings"))
# for each pollinator, define what land use classes are suitable
↪   for nests
nf <- list()
for (j in 1:3) {
nf[[j]] <- parameters[["florNestInfo"]][["attract"]] %>%
dplyr::filter(species == bees[[j]] & Nest_P1_b > 0.6) %>%
```

```
dplyr::select(lu)
}
# natural features are placed on grazing grassland
# and crop fields poorly suited for pollinators
product <- c(6,13,14,28:33,37:39)
names <- c("id","area","l_nf_by_lu","feature","width","gap", "spatial","landuse","share",
"ci_gn_pre","ci_gn_post","ci_tn_pre","ci_tn_post","ci_sb_pre", "ci_sb_post",
"vr_gn_pre","vr_gn_post","vr_tn_pre","vr_tn_post","vr_sb_pre", "vr_sb_post",
"q_gn_pre","q_gn_post","q_tn_pre","q_tn_post","q_sb_pre", "q_sb_post")

# set up parallel processing
cores <- detectCores()
cl <- makeCluster(cores - 1)
registerDoParallel(cl)
pb <- progress_bar$new(
format = "  downloading [:bar] :percent eta: :eta",
total = 430, clear = FALSE, width= 60)

# for each farm "i"
foreach (i = 1:nrow(dat_sf)) %dopar% {

library(plyr)
library(dplyr)
library(sf)
library(raster)
library("EBImage")

# move on if farm located outside of the accepted map
if (is.na(sum(unlist(extract(cc, dat_sf[i,]))))) {next}
```

```
count <- 1
# crop a 2000 by 2000 meter tile centered on farm "i"
tile <- crop(cc, extent(dat_sf[i, ]) + c(-1000, 1000, -1000, 1000))
tile[tile[]==22] = 14
# extract land uses as vector
lu_vec <- unique(values(tile))
if (length(lu_vec)<3) {next}
# initialise the number of loops 'n'
n <- length(lu_vec) * length(widths) * length(gaps) * length(spatial) * 2
out <- matrix(NA, nrow = n, ncol = length(names))
colnames(out) <- names
empty <- tile
values(empty) <- 0
ci_pre <- list()
# for each pollinator species "s"
for (j in 1:3) {
# calculate connectivity index for base scenario
ci_pre[[j]] = connect_index_fun(tile, nf[[j]]$lu, parameters$distance[parameters$distance$species==bees[[j]] & parameters$distance$activity=="foraging", "best_guess"])
}
# pre-intervention
# simulate nests
nf_pre <- computeFloralNesting(landuseMap=tile, edgesMap=stack(empty,empty), unitEdges = "m", widthEdges=10,
landuseCodes, bees=names(bees), num_floral=3,
florNestInfo=parameters$florNestInfo, codeEdges=c(11,21), cell.size = 10,
paramList=parameters)
# simulate visitation rates
```

```
vr_pre <- runpoll_3seasons(M_poll0 = numeric(0), firstyear=TRUE,
↪   firstyearfactor = c(1, 1, 1),
bees = names(bees), cell.size = 10, paramList=parameters,
↪   nest=nf_pre$nest,
floral=nf_pre$floral, cutoff = 0.99, loc_managed)
for (f in 1:2) {
for (w in widths) {
for (g in gaps) {
# determine the placement of natural features in the landscape
treatments <- corridorFun(tile, product, w, g,
↪   10)
for (s in spatial) {
tile_post <- tile
if (s %in% c("corrs","islands"))
↪   {
action <-
↪   feature_infield[f]
}
else {
action <-
↪   feature_edge[f]
}
# create structures s = {corrs, isles, edges} of features f =
↪   {natural regeneration, hedgerows, fruit trees}
tile_post[treatments[[s]][]==1] <- action
ci_post <- list()
#
# # for each pollinator species "s"
for (j in 1:3)
↪   {
# calculate connectivity index for corridors, and islands
```

```
ci_post[[j]] = connect_index_fun(tile_post, nf[[j]]$lu,
↪  parameters$distance[parameters$distance$species==bees[[j]] &
↪  parameters$distance$activity=="foraging",⌋
↪  "best_guess"])
}
# post-treatment
nf_post <- computeFloralNesting(landuseMap=tile_post,
↪  edgesMap=stack(empty,empty), unitEdges = "m", widthEdges=10,
landuseCodes, bees=names(bees), num_floral=3,
florNestInfo=parameters$florNestInfo, codeEdges=c(11,21),
↪  cell.size = 10, paramList=parameters)
vr_post <- runpoll_3seasons(M_poll0 = numeric(0),
↪  firstyear=TRUE, firstyearfactor = c(1, 1, 1),
bees = names(bees), cell.size = 10, paramList=parameters,
↪  nest=nf_post$nest,
floral=nf_post$floral, cutoff = 0.99, loc_managed)

for (lu in unique(lu_vec)) {

out[count, "id"]      <- dat_sf$id[i]
out[count, "feature"] <- action
out[count, "width"]   <- w
out[count, "spatial"] <- s
out[count, "gap"]     <- g
out[count, "area"]  <- length(
↪  treatments[[s]][treatments[[s]][]==1])*100
out[count, "l_nf_by_lu"] <- length(tile_post[tile[]==lu &
↪  tile_post[] == action])*100
out[count, "landuse"] <- lu
out[count, "share"] <- length(tile[tile[] == lu]) / length(tile)
↪  * 100
out[count, "ci_gn_pre"] <- ci_pre[[1]][["ci"]]
out[count, "ci_gn_post"] <- ci_post[[1]][["ci"]]
```

```
out[count, "ci_tn_pre"] <- ci_pre[[3]][["ci"]]
out[count, "ci_tn_post"] <- ci_post[[3]][["ci"]]
out[count, "ci_sb_pre"] <- ci_pre[[2]][["ci"]]
out[count, "ci_sb_post"] <- ci_post[[2]][["ci"]]
out[count, "vr_gn_pre"] <- mean(vr_pre[["flowvis"]][["GNBumblebees"]][[3]][tile[]==lu])
out[count, "vr_gn_post"] <- mean(vr_post[["flowvis"]][["GNBumblebees"]][[3]][tile[]==lu])
out[count, "vr_sb_pre"] <- mean(vr_pre[["flowvis"]][["GNSolitaryBees"]][[1]][tile[]==lu])
out[count, "vr_sb_post"] <- mean(vr_post[["flowvis"]][["GNSolitaryBees"]][[1]][tile[]==lu])
out[count, "vr_tn_pre"] <- mean(vr_pre[["flowvis"]][["TNBumblebees"]][[3]][tile[]==lu])
out[count, "vr_tn_post"] <- mean(vr_post[["flowvis"]][["TNBumblebees"]][[3]][tile[]==lu])
out[count, "q_gn_pre"] <- mean(sum(vr_pre[["M_poll"]][["GNBumblebees"]][[2]][tile[]==lu]), sum(vr_pre[["M_poll"]][["NBumblebees"]][[3]][tile[]==lu]))
out[count, "q_gn_post"] <- mean(sum(vr_post[["M_poll"]][["GNBumblebees"]][[2]][tile[]==lu]), sum(vr_post[["M_poll"]][["GNBumblebees"]][[3]][tile[]==lu]))
out[count, "q_sb_pre"] <- sum(vr_pre[["M_poll"]][["GNSolitaryBees"]][[1]][tile[]==lu])
out[count, "q_sb_post"] <- sum(vr_post[["M_poll"]][["GNSolitaryBees"]][[1]][tile[]==lu])
out[count, "q_tn_pre"] <- mean(sum(vr_pre[["M_poll"]][["TBumblebees"]][[2]][tile[]==lu]), sum(vr_pre[["M_poll"]][["TNBumblebees"]][[3]][tile[]==lu]))
out[count, "q_tn_post"] <- mean(sum(vr_post[["M_poll"]][["TNBumblebees"]][[2]][tile[]==lu]), sum(vr_post[["M_poll"]][["TNBumblebees"]][[3]][tile[]==lu]))
```

```
count <- count + 1
}
}
}
}
}
write.table(out, file=paste("./output/farm_",i,".txt",sep=""), row.names=FALSE, col.names=TRUE)
pb$tick()
}
print("Simulations complete!")
stopCluster(cl)
```

Landscape simulation

This script takes as input the current land use raster, the gap between features, the width of features, and a vector of land use classes, `product`, describing economic crops where ELM features should be applied. It simulates five spatial configurations of the ELM features: In-field corridors, in-field islands, field-edge corridors, field-edge mosaic, and contiguous patch.

```
corridorFun <- function (lu, product, width, gap, res) {
# 'w' determines the number of 10 m2 layers to add when increasing the width of natural features. E.g. if width is 20, add 20/10-1=1 extra layer
w <- width / 10 - 1
gap <- gap / res
corridors <- lu
islands   <- lu
edges <- lu
```

```
mosaic <- lu
isle <- lu
# initialise each spatial configuration
values(corridors) <- 0
values(islands)   <- 0
values(edges)     <- 0
values(mosaic)    <- 0
values(isle)      <- 0
# in-field corridors
verticals   <- seq(1, ncol(lu), gap)
horizontals <- seq(1, ncol(lu), gap)
# populate islands and in-field corridors
corridors[, verticals] <- 1
islands[horizontals, verticals] <- 1
for (i in 0:w) {
corridors[, verticals+i] <- 1
for (j in 0:w){
islands[horizontals+i, verticals+j] <- 1
}
}
# mosaic covering same area as corridors
gap_mos <- ncol(lu) /
↪  round(sqrt(length(corridors[corridors[]==1])), -1)
verticals_mos   <- seq(1, ncol(lu), gap_mos)
horizontals_mos <- seq(1, ncol(lu), gap_mos)
mosaic[horizontals_mos, verticals_mos] <- 1
for (i in 0:w) {
mosaic[horizontals_mos, verticals_mos+i] <- 1
}
for (p in product) {
edges.p <- lu
edges.p[edges.p[] != p] <- NA
edges.p <- boundaries(edges.p)
```

```
edges[edges.p[]==1] <- 1
}
# sample pixels from edges totalling the excess pixels in edges compared to corridors, revert these to original land use class
sample <- sample(which(edges[]==1),
max(length(edges[edges[]==1])-length(corridors[corridors[]==1]), 1),
replace=FALSE)
edges[sample] <- 0
# singular island of natural feature
l <- round(sqrt(length(corridors[corridors[]==1])), -1)
isle[1:l, 1:l] <- 1
corridors[!lu[] %in% product] <- 0
islands[!lu[] %in% product] <- 0
mosaic[!lu[] %in% product] <- 0
isle[!lu[] %in% product] <- 0
projects <- list(isle, corridors, edges, mosaic, islands)
names(projects) <- c("isle","corrs","edges","mosaic","islands")
return(projects)
}
```

Calculate Hanski connectivity

This script calculates the Hanski (Hanski, 1994) index of habitat connectivity.

```
connect_index_fun <- function(raster, scheme, for_dist) {
        # pixels that are not natural features set to 0
        raster[!raster[] %in% scheme] = 0
        # we reduce the resolution
        raster = aggregate(raster, 10, max)
        poly <- as(raster, "SpatialPolygonsDataFrame")
```

```
        lu_sf <- st_as_sf(poly, wkt = "geom")
        names(lu_sf) <- c("lu","geometry")
        # identify pixels that are natural features
        lu_sf <- lu_sf[lu_sf$lu %in% scheme, ]
        lu_buff <- st_buffer(lu_sf, dist = 1, nQuadSegs = 2)
        lu_buff<-st_union(lu_buff)
        # create polygons (parcels) of natural features
        lu_buff<-st_cast(lu_buff, "POLYGON")
        lu_buff <- st_as_sf(lu_buff, wkt = "geom")
        # calculate the size of each parcel in m2
        areas <-st_area(lu_buff[[2]])
        class(areas) = "numeric"
        avg_area = mean(areas)
        ci = vector(length = length(areas))
        coords = lu_buff[[2]]
        d = st_distance(coords, coords, by_element = FALSE)
        class(d) = "numeric"
        # compute the connectivity index following Hanski (1994)
        ci = rowSums(exp(-d/for_dist)*(areas)**0.5)
        # compute average index over landscape cells
        ci <- mean(ci)
        output = list(ci,avg_area)
        names(output) = c("ci","area")
        return(output)
}
```

Latent class modelling and hypothesis testing

The following scripts estimate latent class models and specify hypothesis tests via one-sided t-tests and joint inequality tests.

DCE I Estimation

Reproduces table 4.4 and tests Hypothesis I of chapter 4.

```
### Clear memory
rm(list = ls())
#install.packages("apollo")
#install.packages("readxl")
library(apollo)
library(readxl)
library(dplyr)
library(ggplot2)
# load DCE1 data
database = read_excel("apollo_data_full.xlsx", sheet = "dce1") %>%
# only keep answered choice tasks
filter(serial_sq == 0 &
between(hectare, 15, 800)) %>%
mutate(size = hectare,
mean_size = median(size)) %>%
group_by(id) %>%
mutate(n_optout = sum(choice == 0)/6 * 100) %>%
ungroup()
### Initialise code
apollo_initialise()
### Set core controls
apollo_control = list(
modelName       = "LC_2classes_wta_base_model",
modelDescr      = "LC model on first choice experiment (WTA)",
indivID         = "id",
ncores          = 3,
outputDirectory = "output"
)
# set unestimated attribute coefficients
```

```
apollo_beta=c(
asc_c1_1 = 0,
asc_c1_2      = 0,
asc_c2_1       = 0,
asc_c2_2       = 0,
asc_none      = 0,
b_trees_1     = 0,
b_redge_1     = 0,
b_fbound_1    = 0,
b_infield_1   = 0,
b_goodq_1     = 0,
b_area_1      = 0,
b_pay_1       = 0,
area_fsize_elast_1 = 0,
delta_1       = 0,
b_trees_2     = 0,
b_redge_2     = 0,
b_fbound_2    = 0,
b_infield_2   = 0,
b_goodq_2     = 0,
b_area_2      = 0,
b_pay_2       = 0,
area_fsize_elast_2 = 0,
delta_2    = 0,
female_shift_1 = 0,
female_shift_2 = 0,
grazing_shift_1 = 0,
grazing_shift_2 = 0)
apollo_fixed =
 ↪  c("asc_none","delta_1","b_infield_1","b_infield_2")

#### DEFINE LATENT CLASS COMPONENTS

```

```
apollo_lcPars = function(apollo_beta, apollo_inputs){
        lcpars = list()
        lcpars[["asc_c1"]] = list(asc_c1_1, asc_c1_2)
        lcpars[["asc_c2"]] = list(asc_c2_1, asc_c2_2)
        lcpars[["b_trees"]] = list(b_trees_1, b_trees_2)
        lcpars[["b_redge"]] = list(b_redge_1, b_redge_2)
        lcpars[["b_fbound"]] = list(b_fbound_1, b_fbound_2)
        lcpars[["b_infield"]] = list(b_infield_1, b_infield_2)
        lcpars[["b_goodq"]] = list(b_goodq_1, b_goodq_2)
        lcpars[["b_area"]] = list(b_area_1, b_area_2)
        lcpars[["area_fsize_elast"]] = list(area_fsize_elast_1,
        ↪ area_fsize_elast_2)
        lcpars[["b_pay"]] = list(b_pay_1, b_pay_2)
        lcpars[["female_shift"]] = list(female_shift_1,
        ↪ female_shift_2)
        lcpars[["grazing_shift"]] = list(grazing_shift_1,
        ↪ grazing_shift_2)
        V=list()
        V[["class_1"]] = delta_1
        V[["class_2"]] = delta_2
        classAlloc_settings = list(
        classes     = c(class_1=1, class_2=2),
        utilities   = V
        )
        lcpars[["pi_values"]] =
        ↪ apollo_classAlloc(classAlloc_settings)
        return(lcpars)
}
# GROUP AND VALIDATE INPUTS
apollo_inputs = apollo_validateInputs()
# DEFINE MODEL AND LIKELIHOOD FUNCTION
# -----------------------------------
```

```
apollo_probabilities=function(apollo_beta, apollo_inputs, functionality="estimate"){
        ### Attach inputs and detach after function exit
        apollo_attach(apollo_beta, apollo_inputs)
        on.exit(apollo_detach(apollo_beta, apollo_inputs))
        ### Create list of probabilities P
        P = list()
        ### Define settings for MNL model component that are generic across classes
        mnl_settings = list(
        alternatives = c(A=1, B=2, none=0),
        avail        = list(A=1, B=1, none=1),
        choiceVar    = choice
        )
        area_value = list()
        ### Loop over classes
        for(s in 1:2){
                area_value[[s]] = b_area[[s]] * (size / mean_size) ^ area_fsize_elast[[s]]
                ### Compute class-specific utilities
                V=list()
                V[["none"]] = asc_none
                V[["A"]]  = asc_c1[[s]] + female_shift[[s]] * (female==1) +
                grazing_shift[[s]] * grazing +
                b_trees[[s]] * (c1.type == 0) +
                b_infield[[s]] * (c1.loc == 0) +
                b_redge[[s]] * (c1.loc == 1) +
                b_fbound[[s]] * (c1.loc == 2) +
                b_goodq[[s]] * (c1.qual == 0) +
                area_value[[s]] * (c1.area==1) +
```

```
                b_pay[[s]] * c1.pay
                V[["B"]]  = asc_c2[[s]] + female_shift[[s]] * (female==1) +
                grazing_shift[[s]] * grazing +
                b_trees[[s]] * (c2.type == 0) +
                b_infield[[s]] * (c2.loc == 0) +
                b_redge[[s]] * (c2.loc == 1) +
                b_fbound[[s]] * (c2.loc == 2) +
                b_goodq[[s]] * (c2.qual == 0) +
                area_value[[s]] * (c2.area==1) +
                b_pay[[s]] * c2.pay
                mnl_settings$utilities     = V
                mnl_settings$componentName = paste0("Class_", s)
                ### Compute within-class choice probabilities using MNL model
                P[[paste0("Class_",s)]] = apollo_mnl(mnl_settings, functionality)
                ### Take product across observation for same individual
                P[[paste0("Class_",s)]] = apollo_panelProd(P[[paste0("Class_",s)]], apollo_inputs ,functionality)
        }
        ### Compute latent class model probabilities
        lc_settings  = list(inClassProb = P, classProb=pi_values)
        P[["model"]] = apollo_lc(lc_settings, apollo_inputs, functionality)
        ### Prepare and return outputs of function
        P = apollo_prepareProb(P, apollo_inputs, functionality)
        return(P)
}
```

```
# estimate MNL model and print results
model = apollo_estimate(apollo_beta, apollo_fixed,
  ↪  apollo_probabilities, apollo_inputs)
apollo_modelOutput(model, list(printPVal = TRUE))
conditionals = apollo_conditionals(model,apollo_probabilities,
  ↪  apollo_inputs)
# write to file
apollo_saveOutput(model, saveOutput_settings = list(printPval =
  ↪  TRUE))
# ----------------------------------------
conditionals <- conditionals %>%
rename(id = ID) %>%
# assign class membership based on posterior probabilitites
mutate(class1_prob = case_when(X1 <= 0.2 ~ 2,
X1 >= 0.8 ~ 1,
X1 > 0.2 & X1 < 0.8 ~ 3))
write.csv(conditionals, "lc_conditionals_CE1.csv")
#-----------------------------------------
# Joint inequality test - H0: beta_{payment} <= 0 <= beta_{area}
# High engagement class
omega <- model$varcov
print(row.names(omega))
omega <- as.matrix(omega[c(9,10), c(9,10)])
omega
beta <- model$estimate[11:12]
beta
R <- 10000
draws <- mvtnorm::rmvnorm(R, mean = beta, sigma = omega)
num = sum(draws[,1]>=draws[,2] | (draws[,1]>=0))/R
num # 0.02
# Low engagement class
omega <- model$varcov
print(row.names(omega))
```

```
omega <- as.matrix(omega[c(16,17), c(16,17)])
omega
beta <- model$estimate[20:21]
beta
R <- 10000
draws <- mvtnorm::rmvnorm(R, mean = beta, sigma = omega)
num = sum(draws[,1]>=draws[,2] | (draws[,1]>=0))/R
num         # 0
```

DCE II (WTA) Estimation

Reproduces tables 4.5 and 4.6, and tests Hypothesis II and Hypothesis III of chapter 4.

```
        ### Clear memory
        rm(list = ls())
        #install.packages("apollo")
        #install.packages("readxl")
        library(apollo)
        library(readxl)
        library(dplyr)
        library(ggplot2)
        # load trade WTA data
        database = read_excel("apollo_data_full.xlsx", sheet =
         "dce_wta") %>%
        # only keep answered choice tasks
        filter(serial_sq == 0 & !is.na(age)) %>%
        filter(hectare > quantile(hectare, probs =
         seq(0,1,0.05))[2] &
        hectare < quantile(hectare, probs = seq(0,1,0.05))[20])
         %>%
        mutate(c1.ratio_cont = case_when(c1.ratio == 0 ~ 5,
        c1.ratio == 1 ~ 10,
```

```
        c1.ratio == 2 ~ 20),
        c2.ratio_cont = case_when(c2.ratio == 0 ~ 5,
        c2.ratio == 1 ~ 10,
        c2.ratio == 2 ~ 20),
        c1.fee_cont = case_when(c1.fee == 0 ~ 5,
        c1.fee == 1 ~ 10),
        c2.fee_cont = case_when(c2.fee == 0 ~ 5,
        c2.fee == 1 ~ 10),
        irrational = case_when(task == 2 & choice == 2 ~ 1,
        .default = 0)) %>%
        group_by(id) %>%
        mutate(irrational = max(irrational),
        n_optout = sum(choice == 0)/6 * 100)
        ### Initialise code
        apollo_initialise()
        ### Set core controls
        apollo_control = list(
        modelName       = "LC_2classes_wtp_base_model",
        modelDescr      = "LC model on second choice experiment
         ↪   (WTP)",
        indivID         = "id",
        ncores          = 3,
        outputDirectory = "output"
        )
        # set unestimated attribute coefficients
        apollo_beta=c(asc_c1_1      = 0,
        asc_c1_2      = 0,
        asc_c2_1      = 0,
        asc_c2_2      = 0,
        asc_none    = 0,
        b_5to1_1     = 0,
        b_10to1_1     = 0,
        b_20to1_1     = 0,
```

```
        b_fee_1    = 0,
        b_pay_1       = 0,
        delta_1      = 0,
        b_5to1_2    = 0,
        b_10to1_2   = 0,
        b_20to1_2   = 0,
        b_fee_2   = 0,
        b_pay_2      = 0,
        delta_2   = 0)
        apollo_fixed =
        ↪  c("asc_none","delta_1","b_5to1_1","b_5to1_2")
        # -----------------------------------
        apollo_lcPars = function(apollo_beta, apollo_inputs){
                lcpars = list()
                lcpars[["asc_c1"]] = list(asc_c1_1, asc_c1_2)
                lcpars[["asc_c2"]] = list(asc_c2_1, asc_c2_2)
                lcpars[["b_5to1"]] = list(b_5to1_1, b_5to1_2)
                lcpars[["b_10to1"]] = list(b_10to1_1, b_10to1_2)
                lcpars[["b_20to1"]] = list(b_20to1_1, b_20to1_2)
                lcpars[["b_fee"]] = list(b_fee_1, b_fee_2)
                lcpars[["b_pay"]] = list(b_pay_1,
                ↪  b_pay_2)
                V=list()
                V[["class_1"]] = delta_1
                V[["class_2"]] = delta_2
                classAlloc_settings = list(
                classes      = c(class_1=1, class_2=2),
                utilities     = V
                )
                lcpars[["pi_values"]] =
                ↪  apollo_classAlloc(classAlloc_settings)
                return(lcpars)
        }
```

```
        apollo_inputs = apollo_validateInputs()
        apollo_probabilities=function(apollo_beta,
        ↪  apollo_inputs,
        ↪  functionality="estimate"){
                ### Attach inputs and detach after function exit
                apollo_attach(apollo_beta, apollo_inputs)
                on.exit(apollo_detach(apollo_beta,
                ↪  apollo_inputs))
                ### Create list of probabilities P
                P = list()
                ### Define settings for MNL model component that
                ↪  are generic across classes
                mnl_settings = list(
                alternatives = c(A=1, B=2, none=0),
                avail        = list(A=1, B=1, none=1),
                choiceVar    = choice
                )
                ### Loop over classes
                for(s in 1:2){
                        ### Compute class-specific utilities
                        V=list()
                        V[["none"]] = asc_none
                        V[["A"]]  = asc_c1[[s]] +
                        b_5to1[[s]] * (c1.ratio == 0) +
                        b_10to1[[s]] * (c1.ratio == 1) +
                        b_20to1[[s]] * (c1.ratio == 2) +
                        b_fee[[s]] * c1.fee_cont +
                        b_pay[[s]] * c1.pay
                        V[["B"]]  = asc_c2[[s]] +
                        b_5to1[[s]] * (c2.ratio == 0) +
                        b_10to1[[s]] * (c2.ratio == 1) +
                        b_20to1[[s]] * (c2.ratio == 2) +
                        b_fee[[s]] * c2.fee_cont +
```

```
                    b_pay[[s]] * c2.pay
                    mnl_settings$utilities     = V
                    mnl_settings$componentName = paste0("Class_", s)
                    ### Compute within-class choice probabilities using MNL model
                    P[[paste0("Class_",s)]] = apollo_mnl(mnl_settings, functionality)
                    ### Take product across observation for same individual
                    P[[paste0("Class_",s)]] = apollo_panelProd(P[[paste0("Class_", s)]], apollo_inputs , functionality)
            }
            ### Compute latent class model probabilities
            lc_settings  = list(inClassProb = P, classProb=pi_values)
            P[["model"]] = apollo_lc(lc_settings, apollo_inputs, functionality)
            ### Prepare and return outputs of function
            P = apollo_prepareProb(P, apollo_inputs, functionality)
            return(P)
    }
    # estimate MNL model and print results
    model = apollo_estimate(apollo_beta, apollo_fixed, apollo_probabilities, apollo_inputs)
    apollo_modelOutput(model, list(printPVal = TRUE))
```

```
conditionals = apollo_conditionals(model,apollo_probabilities, apollo_inputs)
# write to file
apollo_saveOutput(model, saveOutput_settings = list(printPval = TRUE))
# ------------------------------------
# assign respondents to classes based on posterior probabilities
conditionals <- conditionals %>%
rename(id = ID) %>%
mutate(class1_prob = case_when(X1 <= 0.2 ~ 2,
X1 >= 0.8 ~ 1,
X1 > 0.2 & X1 < 0.8 ~ 3))
write.csv(conditionals, "lc_conditionals_wta.csv")
# Joint inequality test: H0: beta_{20:1} <= beta_{10:1} <= 0
# high engagement class
omega <- model$varcov
print(row.names(omega))
omega <- as.matrix(omega[c(5,6), c(5,6)])
omega
beta <- model$estimate[7:8]
beta
R <- 10000
draws <- mvtnorm::rmvnorm(R, mean = beta, sigma = omega)
num = sum(draws[,1]>draws[,2])/R
num # 0.00
# low engagement class
omega <- model$varcov
print(row.names(omega))
omega <- as.matrix(omega[c(9,10), c(9,10)])
omega
```

```
beta <- model$estimate[13:14]
beta
R <- 10000
draws <- mvtnorm::rmvnorm(R, mean = beta, sigma = omega)
num = sum(draws[,1]>draws[,2])/R
num # 0.084
```

DCE II (WTP) Estimation

```
### Clear memory
rm(list = ls())
library(apollo)
library(readxl)
library(dplyr)
library(ggplot2)
# load trade WTP scenario data
database = read_excel("apollo_data_full.xlsx", sheet = "dce_wtp") %>%
# only keep answered choice tasks
filter(serial_sq == 0) %>%
filter(hectare > quantile(hectare, probs = seq(0,1,0.05))[2] &
hectare < quantile(hectare, probs = seq(0,1,0.05))[20]) %>%
mutate(c1.ratio_cont = case_when(c1.ratio == 0 ~ 5,
c1.ratio == 1 ~ 10,
c1.ratio == 2 ~ 20),
c2.ratio_cont = case_when(c2.ratio == 0 ~ 5,
c2.ratio == 1 ~ 10,
c2.ratio == 2 ~ 20),
c1.fee_cont = case_when(c1.fee == 0 ~ 5,
c1.fee == 1 ~ 10),
c2.fee_cont = case_when(c2.fee == 0 ~ 5,
```

```
        c2.fee == 1 ~ 10),
        irrational = case_when(task == 3 & choice == 2 ~ 1,
        .default = 0)) %>%
        group_by(id) %>%
        mutate(n_optout = sum(choice == 0)/6 * 100) %>%
        # exclude respondents who answer choice task 3
        ↪  irrationally
        filter(all(irrational == 0)) %>%
        ungroup()
        ### Initialise code
        apollo_initialise()
        ### Set core controls
        apollo_control = list(
        modelName       = "LC_2classes_wtp_base_model",
        modelDescr      = "LC model on second choice experiment
        ↪  (WTP)",
        indivID         = "id",
        ncores          = 3,
        outputDirectory = "output"
        )
        # set unestimated attribute coefficients
        apollo_beta=c(asc_c1_1      = 0,
        asc_c1_2      = 0,
        asc_c2_1      = 0,
        asc_c2_2      = 0,
        asc_none   = 0,
        b_5to1_1    = 0,
        b_10to1_1    = 0,
        b_20to1_1    = 0,
        b_fee_1    = 0,
        b_pay_1       = 0,
        delta_1      = 0,
        b_5to1_2    = 0,
```

```
        b_10to1_2    = 0,
        b_20to1_2    = 0,
        b_fee_2    = 0,
        b_pay_2        = 0,
        delta_2    = 0)
        apollo_fixed = c("asc_none","delta_1","b_5to1_1",
        ↪  "b_5to1_2")
        apollo_lcPars = function(apollo_beta, apollo_inputs){
                lcpars = list()
                lcpars[["asc_c1"]] = list(asc_c1_1, asc_c1_2)
                lcpars[["asc_c2"]] = list(asc_c2_1, asc_c2_2)
                lcpars[["b_5to1"]] = list(b_5to1_1, b_5to1_2)
                lcpars[["b_10to1"]] = list(b_10to1_1, b_10to1_2)
                lcpars[["b_20to1"]] = list(b_20to1_1, b_20to1_2)
                lcpars[["b_fee"]] = list(b_fee_1, b_fee_2)
                lcpars[["b_pay"]] = list(b_pay_1,
                ↪  b_pay_2)
                V=list()
                V[["class_1"]] = delta_1
                V[["class_2"]] = delta_2
                classAlloc_settings = list(
                classes      = c(class_1=1, class_2=2),
                utilities     = V
                )
                lcpars[["pi_values"]] =
                ↪  apollo_classAlloc(classAlloc_settings)
                return(lcpars)
        }
        apollo_inputs = apollo_validateInputs()
        apollo_probabilities=function(apollo_beta,
        ↪  apollo_inputs,
        ↪  functionality="estimate"){
                ### Attach inputs and detach after function exit
```

```
apollo_attach(apollo_beta, apollo_inputs)
on.exit(apollo_detach(apollo_beta,
 ↪ apollo_inputs))
### Create list of probabilities P
P = list()
### Define settings for MNL model component that
 ↪ are generic across classes
mnl_settings = list(
alternatives = c(A=1, B=2, none=0),
avail        = list(A=1, B=1, none=1),
choiceVar    = choice
)
### Loop over classes
for(s in 1:2){
        ### Compute class-specific utilities
        V=list()
        V[["none"]] = asc_none
        V[["A"]]  = asc_c1[[s]] +
        b_5to1[[s]] * (c1.ratio == 0) +
        b_10to1[[s]] * (c1.ratio == 1) +
        b_20to1[[s]] * (c1.ratio == 2) +
        b_fee[[s]] * c1.fee_cont +
        b_pay[[s]] *
         ↪ c1.pay
        V[["B"]]  = asc_c2[[s]] +
        b_5to1[[s]] * (c2.ratio == 0) +
        b_10to1[[s]] * (c2.ratio == 1) +
        b_20to1[[s]] * (c2.ratio == 2) +
        b_fee[[s]] * c2.fee_cont +
        b_pay[[s]] * c2.pay
         ↪
        mnl_settings$utilities     = V
```

```
                    mnl_settings$componentName = paste0("Class_", s)
                    ### Compute within-class choice probabilities using MNL model
                    P[[paste0("Class_",s)]] = apollo_mnl(mnl_settings, functionality)
                    ### Take product across observation for same individual
                    P[[paste0("Class_",s)]] = apollo_panelProd(P[[paste0("Class_", s)]], apollo_inputs , functionality)
              }
              ### Compute latent class model probabilities
              lc_settings  = list(inClassProb = P, classProb=pi_values)
              P[["model"]] = apollo_lc(lc_settings, apollo_inputs, functionality)
              ### Prepare and return outputs of function
              P = apollo_prepareProb(P, apollo_inputs, functionality)
              return(P)
        }
        # estimate MNL model and print results
        model = apollo_estimate(apollo_beta, apollo_fixed, apollo_probabilities, apollo_inputs)
        apollo_modelOutput(model, list(printPVal = TRUE))
        conditionals = apollo_conditionals(model,apollo_probabilities, apollo_inputs)
        # write to file
```

```
apollo_saveOutput(model, saveOutput_settings = list(printPval = TRUE))
# -------------------------------------------------------
# assign respondents to classes based on posterior probabilities
conditionals <- conditionals %>%
rename(id = ID) %>%
mutate(class1_prob = case_when(X1 <= 0.2 ~ 2,
X1 >= 0.8 ~ 1,
X1 > 0.2 & X1 < 0.8 ~ 3))
write.csv(conditionals, "lc_conditionals_wtp.csv")
# Joint inequality test: H0: beta_{20:1} <= beta_{10:1} <= 0
# high engagement class
omega <- model$varcov
print(row.names(omega))
omega <- as.matrix(omega[c(5,6), c(5,6)])
omega
beta <- model$estimate[7:8]
beta
R <- 10000
draws <- mvtnorm::rmvnorm(R, mean = beta, sigma = omega)
num = sum(draws[,1]>=draws[,2])/R
num # 0.004
# low engagement class
omega <- model$varcov
print(row.names(omega))
omega <- as.matrix(omega[c(9,10), c(9,10)])
omega
beta <- model$estimate[13:14]
beta
R <- 10000
```

```
        draws <- mvtnorm::rmvnorm(R, mean = beta, sigma = omega)
        num = sum(draws[,1]>=draws[,2])/R
        num # 0.16
```

DCE III Estimation

Reproduces table 5.5 and tests Hypothesis II and Hypothesis III of chapter 5.

```
        ### Clear memory
        rm(list = ls())
        library(apollo)
        library(readxl)
        library(dplyr)
        # load trade CE3 data
        database = read_excel("apollo_data_full.xlsx", sheet = "dce3") %>%
        # only keep answered choice tasks
        filter(serial_sq == 0 & !is.na(ea_respect)) %>%
        filter(hectare > quantile(hectare, probs = seq(0,1,0.05))[2] &
        hectare < quantile(hectare, probs = seq(0,1,0.05))[20]) %>%
        mutate(c1.bonus_real = c1.coord * c1.bonus,
        c2.bonus_real = c2.coord * c2.bonus)
        ### Initialise code
        apollo_initialise()
        ### Set core controls
        apollo_control = list(
        modelName       = "LC_ce3_model",
        modelDescr      = "LC model on third choice experiment",
        indivID         = "id",
        ncores          = 3,
        outputDirectory = "output"
```

```
)
# set unestimated attribute coefficients
apollo_beta=c(asc_c1_1    = 0,
asc_c1_2    = 0,
asc_c2_1    = 0,
asc_c2_2    = 0,
asc_none    = 0,
b_trees_1     = 0,
b_trees_2     = 0,
b_nocoord_1   = 0,
b_nocoord_2   = 0,
b_coord1_1 = 0,
b_coord1_2 = 0,
b_coord2_1 = 0,
b_coord2_2 = 0,
coord1_sharing_1 = 0,
coord1_sharing_2 = 0,
coord2_sharing_1 = 0,
coord2_sharing_2 = 0,
b_width10m_1 = 0,
b_width10m_2 = 0,
b_width20m_1 = 0,
b_width20m_2 = 0,
b_bonus_1     = 0,
b_bonus_2     = 0,
b_pay_1       = 0,
b_pay_2       = 0,
delta_1       = 0,
delta_2       = 0,
grazing_shift_1 = 0,
grazing_shift_2 = 0
)
```

```
        apollo_fixed = c("asc_none", "delta_1", "b_nocoord_1",
        ↪  "b_nocoord_2","b_width10m_1","b_width10m_2")
        apollo_lcPars = function(apollo_beta, apollo_inputs){
                lcpars = list()
                lcpars[["asc_c1"]] = list(asc_c1_1, asc_c1_2)
                lcpars[["asc_c2"]] = list(asc_c2_1, asc_c2_2)
                lcpars[["b_trees"]] = list(b_trees_1, b_trees_2)
                lcpars[["b_nocoord"]] = list(b_nocoord_1,
                ↪  b_nocoord_2)
                lcpars[["b_coord1"]] = list(b_coord1_1,
                ↪  b_coord1_2)
                lcpars[["b_coord2"]] = list(b_coord2_1,
                ↪  b_coord2_2)
                lcpars[["coord1_sharing"]] =
                ↪  list(coord1_sharing_1, coord1_sharing_2)
                lcpars[["coord2_sharing"]] =
                ↪  list(coord2_sharing_1, coord2_sharing_2)
                lcpars[["b_width10m"]] = list(b_width10m_1,
                ↪  b_width10m_2)
                lcpars[["b_width20m"]] = list(b_width20m_1,
                ↪  b_width20m_2)
                lcpars[["b_bonus"]] = list(b_bonus_1, b_bonus_2)
                lcpars[["b_pay"]] = list(b_pay_1, b_pay_2)
                lcpars[["grazing_shift"]] =
                ↪  list(grazing_shift_1, grazing_shift_2)
                V=list()
                V[["class_1"]] = delta_1
                V[["class_2"]] = delta_2
                classAlloc_settings = list(
                classes      = c(class_1=1, class_2=2),
                utilities    = V
                )
```

```
            lcpars[["pi_values"]] = apollo_classAlloc(classAlloc_settings)
            return(lcpars)
        }
        apollo_inputs = apollo_validateInputs()
        apollo_probabilities=function(apollo_beta, apollo_inputs, functionality="estimate"){
            ### Attach inputs and detach after function exit
            apollo_attach(apollo_beta, apollo_inputs)
            on.exit(apollo_detach(apollo_beta, apollo_inputs))
            ### Create list of probabilities P
            P = list()
            ### Define settings for MNL model component that are generic across classes
            mnl_settings = list(
            alternatives = c(A=1, B=2, none=0),
            avail        = list(A=1, B=1, none=1),
            choiceVar    = choice
            )
            ### Loop over classes
            for(s in 1:2){
                ### Compute class-specific utilities
                V=list()
                V[["none"]]  =  asc_none
                V[["A"]]  =           asc_c1[[s]] +
                grazing_shift[[s]] * grazing +
                b_trees[[s]] * (c1.type == 0) +
                b_width10m[[s]] * (c1.width == 0) +
                b_width20m[[s]] * (c1.width == 1) +
                b_nocoord[[s]] * (c1.coord == 0) +
                b_coord1[[s]] * (c1.coord == 1) +
                b_coord2[[s]] * (c1.coord == 2) +
```

```
    coord1_sharing[[s]] * (c1.coord==1) * (sharing==1) +
    coord2_sharing[[s]] * (c1.coord==2) * (sharing==1) +
    b_bonus[[s]] * c1.bonus_real +
    b_pay[[s]]*c1.pay
    V[["B"]]  =           asc_c2[[s]] +
    grazing_shift[[s]] * grazing +
    b_trees[[s]] * (c2.type == 0) +
    b_width10m[[s]] * (c2.width == 0) +
    b_width20m[[s]] * (c2.width == 1) +
    b_nocoord[[s]] * (c2.coord == 0) +
    b_coord1[[s]] * (c2.coord == 1) +
    b_coord2[[s]] * (c2.coord == 2) +
    coord1_sharing[[s]] * (c2.coord==1) * (sharing==1) +
    coord2_sharing[[s]] * (c2.coord==2) * (sharing==1) +
    b_bonus[[s]] * c2.bonus_real +
    b_pay[[s]] * c2.pay
    mnl_settings$utilities     = V
    mnl_settings$componentName = paste0("Class_", s)
    ### Compute within-class choice probabilities using MNL model
    P[[paste0("Class_",s)]] = apollo_mnl(mnl_settings, functionality)
    ### Take product across observation for same individual
    P[[paste0("Class_",s)]] = apollo_panelProd(P[[paste0("Class_", s)]], apollo_inputs ,functionality)
```

```
            }
            ### Compute latent class model probabilities
            lc_settings  = list(inClassProb = P,
             ↪ classProb=pi_values)
            P[["model"]] = apollo_lc(lc_settings,
             ↪ apollo_inputs, functionality)
            ### Prepare and return outputs of function
            P = apollo_prepareProb(P, apollo_inputs,
             ↪ functionality)
            return(P)
    }
    # estimate MNL model and print results
    model = apollo_estimate(apollo_beta, apollo_fixed,
     ↪ apollo_probabilities, apollo_inputs)
    conditionals <-
     ↪ apollo_conditionals(model,apollo_probabilities,
     ↪ apollo_inputs)
    apollo_modelOutput(model, list(printPVal = TRUE))
    # write to file
    apollo_saveOutput(model, saveOutput_settings =
     ↪ list(printPval = TRUE))
    # ------------------------------------
    conditionals <- conditionals %>%
    rename(id = ID) %>%
    mutate(class1_prob = case_when(X1 <= 0.2 ~ 1,
    X1 >= 0.8 ~ 2,
    X1 > 0.2 & X1 < 0.8 ~ 3))
    write.csv(conditionals, "lc_conditionals_ce3.csv")
    # Joint inequality tests
    # H2 class 2
    omega <- model$varcov
    print(row.names(omega))
    omega <- as.matrix(omega[c(8,10), c(8,10)])
```

```
omega
beta <- model$estimate[c(11,13)]
beta
R <- 10000
draws <- mvtnorm::rmvnorm(R, mean = beta, sigma = omega)
num = sum(draws[,1]<=draws[,2])/R
num # 0.12
# H2 class 1
omega <- model$varcov
print(row.names(omega))
omega <- as.matrix(omega[c(7,9), c(7,9)])
omega
beta <- model$estimate[c(10,12)]
beta
R <- 10000
draws <- mvtnorm::rmvnorm(R, mean = beta, sigma = omega)
num = sum(draws[,1]<=draws[,2])/R
num # 0.001
# H3 class I
omega <- model$varcov
print(row.names(omega))
omega <- as.matrix(omega[c(11,13), c(11,13)])
omega
beta <- model$estimate[c(14,16)]
beta
R <- 10000
draws <- mvtnorm::rmvnorm(R, mean = beta, sigma = omega)
num = sum(draws[,1]<=draws[,2] | draws[,2]<=0)/R
num # 0.78
# H3 class II
omega <- model$varcov
print(row.names(omega))
omega <- as.matrix(omega[c(12,14), c(12,14)])
```

```
omega
beta <- model$estimate[c(15,17)]
beta
R <- 10000
draws <- mvtnorm::rmvnorm(R, mean = beta, sigma = omega)
num = sum(draws[,1]<=draws[,2] | draws[,1]<=0)/R
num # 0.66
```

Bibliography

Abdel-Rahman, A. A. (2008). On the atmospheric dispersion and gaussian plume model. *2nd International Conference on Waste Management, Water Pollution, Air Pollution, Indoor Climate*, 31–39.

Adamowicz, W. L., Glenk, K., & Meyerhoff, J. (2014). Choice modelling research in environmental and resource economics. In *Handbook of choice modelling* (pp. 661–674). Edward Elgar Publishing.

Antoine, C. M., & Forrest, J. R. (2021). Nesting habitat of ground-nesting bees: A review. *Ecological Entomology*, *46*(2), 143–159.

Arguez, A., Durre, I., Applequist, S., Vose, R. S., Squires, M. F., Yin, X., Heim Jr, R. R., & Owen, T. W. (2012). Noaa's 1981–2010 us climate normals: An overview. *Bulletin of the American Meteorological Society*, *93*(11), 1687–1697.

Axhausen, K. W., Hess, S., König, A., Abay, G., Bates, J. J., & Bierlaire, M. (2008). Income and distance elasticities of values of travel time savings: New swiss results. *Transport Policy*, *15*(3), 173–185.

Banerjee, S., Cason, T. N., de Vries, F. P., & Hanley, N. (2017). Transaction costs, communication and spatial coordination in payment for ecosystem services schemes. *Journal of Environmental Economics and Management*, *83*, 68–89.

Banerjee, S., De Vries, F. P., Hanley, N., & Van Soest, D. P. (2014). The impact of information provision on agglomeration bonus performance: An experimental study on local networks. *American Journal of Agricultural Economics*, *96*(4), 1009–1029.

Barreca, A. I., Neidell, M., & Sanders, N. J. (2021). Long-run pollution exposure and mortality: Evidence from the acid rain program. *Journal of Public Economics*, *200*, 104440.

Bartkowski, B., Droste, N., Ließ, M., Sidemo-Holm, W., Weller, U., & Brady, M. V. (2021). Payments by modelled results: A novel design for agri-environmental schemes. *Land Use Policy*, *102*, 105230.

Bates, P. D., Savage, J., Wing, O., Quinn, N., Sampson, C., Neal, J., & Smith, A. (2023). A climate-conditioned catastrophe risk model for uk flooding. *Natural Hazards and Earth System Sciences*, *23*(2), 891–908.

Benjamin, E. O., & Sauer, J. (2018). The cost effectiveness of payments for ecosystem services—smallholders and agroforestry in africa. *Land use policy*, *71*, 293–302.

Berck, P., Moe-Lange, J., Stevens, A., & Villas-Boas, S. (2016). Measuring consumer responses to a bottled water tax policy. *American Journal of Agricultural Economics*, *98*(4), 981–996.

Berg, Å., Cronvall, E., Eriksson, Å., Glimskär, A., Hiron, M., Knape, J., Pärt, T., Wissman, J., Żmihorski, M., & Öckinger, E. (2019). Assessing agri-environmental schemes for semi-natural grasslands during a 5-year period: Can we see positive effects for vascular plants and pollinators? *Biodiversity and Conservation*, *28*, 3989–4005.

Bergé, L., et al. (2018). *Efficient estimation of maximum likelihood models with multiple fixed-effects: The r package fenmlm* (tech. rep.). Department of Economics at the University of Luxembourg.

Beychok, M. R. (2005). *Fundamentals of stack gas dispersion*. MR Beychok.

Binetti, A., Nuzzi, F., & Stantcheva, S. (2024). *People's understanding of inflation* (tech. rep.). National Bureau of Economic Research.

Block, J. B., Danne, M., & Mußhoff, O. (2024). Farmers' willingness to participate in a carbon sequestration program–a discrete choice experiment. *Environmental Management*, 1–18.

Bostian, M., Färe, R., Grosskopf, S., & Lundgren, T. (2022). Prevention or cure? optimal abatement mix. *Environmental Economics and Policy Studies*, *24*(4), 503–531.

Boulton, A. J., & Williford, A. (2018). Analyzing skewed continuous outcomes with many zeros: A tutorial for social work and youth prevention science researchers. *Journal of the Society for Social Work and Research*, *9*(4), 721–740.

Boxall, P. C., & Adamowicz, W. L. (2002). Understanding heterogeneous preferences in random utility models: A latent class approach. *Environmental and resource economics*, *23*, 421–446.

Breeze, T., Bailey, A., Potts, S., & Balcombe, K. (2015). A stated preference valuation of the non-market benefits of pollination services in the uk. *Ecological Economics*, *111*, 76–85.

Breeze, T. D., Bailey, A. P., Balcombe, K. G., Brereton, T., Comont, R., Edwards, M., Garratt, M. P., Harvey, M., Hawes, C., Isaac, N., et al. (2021). Pollinator monitoring more than pays for itself. *Journal of Applied Ecology*, *58*(1), 44–57.

Breeze, T. D., Vaissière, B. E., Bommarco, R., Petanidou, T., Seraphides, N., Kozák, L., Scheper, J., Biesmeijer, J. C., Kleijn, D., Gyldenkærne, S., et al. (2014). Agricultural policies exacerbate honeybee pollination service supply-demand mismatches across europe. *PloS one*, *9*(1), e82996.

Brian, A. D. (1952). Division of labour and foraging in bombus agrorum fabricius. *The journal of animal ecology*, 223–240.

Briggs, G. A. (1965). A plume rise model compared with observations. *Journal of the Air Pollution Control Association*, *15*(9), 433–438.

Briggs, G. A. (1982). Plume rise predictions. In *Lectures on air pollution and environmental impact analyses* (pp. 59–111). Springer.

Broadmeadow, S., Nisbet, T., Palmer, R., Webb, L., Short, C., Chivers, C.-A., Hammond, J., Lukac, M., Miller, A., Gantlett, R., et al. (2023). Incorporating technical and farmer knowledge to improve land use and management for natural flood management in lowland catchments. *Land Use Policy*, *128*, 106596.

Buntin, M. B., & Zaslavsky, A. M. (2004). Too much ado about two-part models and transformation?: Comparing methods of modeling medicare expenditures. *Journal of health economics*, *23*(3), 525–542.

Burton, M., & Rigby, D. (2009). Hurdle and latent class approaches to serial non-participation in choice models. *Environmental and Resource Economics*, *42*, 211–226.

Cai, H., Chen, Y., & Gong, Q. (2016). Polluting thy neighbor: Unintended consequences of china's pollution reduction mandates. *Journal of Environmental Economics and Management*, *76*, 86–104.

Campbell, D., Boeri, M., Doherty, E., & Hutchinson, W. G. (2015). Learning, fatigue and preference formation in discrete choice experiments. *Journal of Economic Behavior & Organization*, *119*, 345–363.

Cardwell, M. (2023). Results-based agri-environmental scheme design: Legal implications. *Environmental Law Review*, *25*(4), 260–288.

Carson, J., & Moses, H. (1969). The validity of several plume rise formulas. *Journal of the air pollution control association*, *19*(11), 862–866.

Carson, R., & Groves, T. (2007). Incentive and informational properties of preference questions. *Environmental and resource economics*, *37*, 181–210.

Chan, G., Stavins, R., Stowe, R., & Sweeney, R. (2012). The so_2 allowance-trading system and the clean air act amendments of 1990: Reflections on 20 years of policy innovation. *National Tax Journal*, *65*(2), 419–452.

Chan, N. W., & Morrow, J. W. (2019). Unintended consequences of cap-and-trade? evidence from the regional greenhouse gas initiative. *Energy Economics*, *80*, 411–422.

Chay, K., & Greenstone, M. (2003a). Air quality, infant mortality, and the clean air act of 1970.

Chay, K., & Greenstone, M. (2003b). The impact of air pollution on infant mortality: Evidence from geographic variation in pollution shocks induced by a recession. *The quarterly journal of economics*, *118*(3), 1121–1167.

Chen, S., Zivin, J. S. G., Wang, H., & Xiong, J. (2022). *Combating cross-border externalities* (tech. rep.). National Bureau of Economic Research.

Chestnut, L. G., & Mills, D. M. (2005). A fresh look at the benefits and costs of the us acid rain program. *Journal of environmental management*, *77*(3), 252–266.

ChoiceMetrics, C. (2012). Ngene 1.1. 1 user manual & reference guide. *Sydney, Australia.*

Clements, J., Lobley, M., Osborne, J., & Wills, J. (2021). How can academic research on uk agri-environment schemes pivot to meet the addition of climate mitigation aims? *Land Use Policy*, *106*, 105441.

Coase, R. H. (1960). The problem of social cost. *The journal of Law and Economics*, *56*(4), 837–877.

Cormen, T. H., Leiserson, C. E., Rivest, R. L., & Stein, C. (2022). *Introduction to algorithms*. MIT press.

Correa Ayram, C. A., Mendoza, M. E., Etter, A., & Salicrup, D. R. P. (2016). Habitat connectivity in biodiversity conservation: A review of recent studies and applications. *Progress in Physical Geography*, *40*(1), 7–37.

Costanza, R., d'Arge, R., De Groot, R., Farber, S., Grasso, M., Hannon, B., Limburg, K., Naeem, S., O'neill, R. V., Paruelo, J., et al. (1997). The value of the world's ecosystem services and natural capital. *Nature*, *387*(6630), 253–260.

Crespi, J. M., Saitone, T. L., & Sexton, R. J. (2012). Competition in us farm product markets: Do long-run incentives trump short-run market power? *Applied Economic Perspectives and Policy*, *34*(4), 669–695.

Czajkowski, M., Zagorska, K., Letki, N., Tryjanowski, P., & Was, A. (2021). Drivers of farmers' willingness to adopt extensive farming practices in a globally important bird area. *Land Use Policy*, *107*, 104223.

Dadson, S. J., Hall, J. W., & et al., A. M. (2017). A restatement of the natural science evidence concerning catchment-based 'natural' flood management in the uk. *Proc. R. Soc. A*, *473*, 1–19.

Dawson, P., & Lingard, J. (1982). Management bias and returns to scale in a cobb-douglas production function for agriculture. *European Review of Agricultural Economics*, *9*(1), 7–24.

de Bekker-Grob, E. W., Donkers, B., Jonker, M. F., & Stolk, E. A. (2015). Sample size requirements for discrete-choice experiments in healthcare: A practical guide. *The Patient-Patient-Centered Outcomes Research*, *8*, 373–384.

De Kluizenaar, Y., Aherne, J., & Farrell, E. (2001). Modelling the spatial distribution of so2 and nox emissions in ireland. *Environmental Pollution*, *112*(2), 171–182.

De Palma, A., Kuhlmann, M., Roberts, S. P., Potts, S. G., Börger, L., Hudson, L. N., Lysenko, I., Newbold, T., & Purvis, A. (2015). Ecological traits affect the sensitivity of bees to land-use pressures in european agricultural landscapes. *Journal of Applied Ecology*, *52*(6), 1567–1577.

Defra. (2022). *Environmental land management schemes: Overview*. Retrieved November 20, 2022, from https://www.gov.uk/government/publications/environmental-land-management-schemes-overview/environmental-land-management-scheme-overview

Department for Environment, Food and Rural Affairs. (n.d.). *Farm accounts in england data sets*. https://www.gov.uk/government/statistics/farm-accounts-in-england-data-sets

Desiere, S., & Jolliffe, D. (2018). Land productivity and plot size: Is measurement error driving the inverse relationship? *Journal of Development Economics*, *130*, 84–98.

Dewally, K., Bark, R., Harwood, A., & Lovett, A. (2025). Learning from the past and embracing future opportunities: Perceptions of new environmental land management schemes and private nature markets. *Journal of Rural Studies*, *119*, 103723.

Dicks, L. V., Viana, B., Bommarco, R., Brosi, B., Arizmendi, M. d. C., Cunningham, S. A., Galetto, L., Hill, R., Lopes, A. V., Pires, C., et al. (2016). Ten policies for pollinators. *Science*, *354*(6315), 975–976.

Donald, P. F., & Evans, A. D. (2006). Habitat connectivity and matrix restoration: The wider implications of agri-environment schemes. *Journal of applied ecology*, *43*(2), 209–218.

Donald, S. G., & Lang, K. (2007). Inference with difference-in-differences and other panel data. *The Review of Economics and Statistics*, *89*(2), 221–233.

Donkersley, P., Rhodes, G., Pickup, R. W., Jones, K. C., & Wilson, K. (2014). Honeybee nutrition is linked to landscape composition. *Ecology & Evolution*, *4*(21), 4195–4206.

Dottori, F., Mentaschi, L., Bianchi, A., Alfieri, L., & Feyen, L. (2023). Cost-effective adaptation strategies to rising river flood risk in europe. *Nature Climate Change*, *13*(2), 196–202.

Dubos-Paillard, E., Lavaine, E., Millock, K., et al. (2019). The effect of flood risk information on property values around paris. *European Conference on Risk Perception, Behaviour, Management and Response.*

Eden River Trust. (2025). All about eden's rivers [Accessed: October 4, 2025]. https://edenriverstrust.org.uk/edens-rivers/

Ehlers, M.-H., Huber, R., & Finger, R. (2021). Agricultural policy in the era of digitalisation. *Food policy*, *100*, 102019.

Ellerman, A. D., Joskow, P. L., Schmalensee, R., Montero, J.-P., & Bailey, E. M. (2000). *Markets for clean air: The us acid rain program*. Cambridge University Press.

Ellis, S. F., Masters, M., Messer, K. D., Weigel, C., & Ferraro, P. J. (2021). The problem of feral hogs and the challenges of providing a weak-link public good. *Applied Economic Perspectives and Policy*, *43*(3), 985–1002.

Environment Agency. (2021). *Spatial prioritisation of catchments suitable for using natural flood management* (tech. rep.). Environment Agency. https://www.data.gov.uk/dataset/2a4bcf6e-3880-4c0b-9986-6cafbec89faf/spatial-prioritisation-of-catchments-suitable-for-using-natural-flood-management

Ewald, J., Sterner, T., & Sterner, E. (2022). Understanding the resistance to carbon taxes: Drivers and barriers among the general public and fuel-tax protesters. *Resource and Energy Economics*, *70*, 101331.

Finger, R., & Möhring, N. (2024). The emergence of pesticide-free crop production systems in europe. *Nature Plants*, *10*(3), 360–366.

Forbes, H., Ball, K., & McLay, F. (2015). *Natural flood management handbook*. Scottish Environmental Protection Agency.

Fowlie, M., Greenstone, M., & Wolfram, C. (2018). Do energy efficiency investments deliver? evidence from the weatherization assistance program. *The Quarterly Journal of Economics*, *133*(3), 1597–1644.

Fowlie, M., Holland, S. P., & Mansur, E. T. (2012). What do emissions markets deliver and to whom? evidence from southern california's nox trading program. *American Economic Review*, *102*(2), 965–993.

Fowlie, M., & Muller, N. (2019). Market-based emissions regulation when damages vary across sources: What are the gains from differentiation? *Journal of the Association of Environmental and Resource Economists*, *6*(3), 593–632.

Franks, J. R. (2011). The collective provision of environmental goods: A discussion of contractual issues. *Journal of Environmental Planning and Management*, *54*(5), 637–660.

Gardner, E., Breeze, T. D., Clough, Y., Smith, H. G., Baldock, K. C., Campbell, A., Garratt, M. P., Gillespie, M. A., Kunin, W. E., McKerchar, M., et al. (2020). Reliably predicting pollinator abundance: Challenges of calibrating process-based ecological models. *Methods in Ecology and Evolution*, *11*(12), 1673–1689.

Gardner, E., Robinson, R. A., Julian, A., Boughey, K., Langham, S., Tse-Leon, J., Petrovskii, S., Baker, D. J., Bellamy, C., Buxton, A., et al. (2024). A family of process-based models to simulate landscape use by multiple taxa. *Landscape Ecology*, *39*(5), 102.

Garibaldi, L. A., Carvalheiro, L. G., Vaissière, B. E., Gemmill-Herren, B., Hipólito, J., Freitas, B. M., Ngo, H. T., Azzu, N., Sáez, A., Åström, J., et al. (2016). Mutually beneficial pollinator diversity and crop yield outcomes in small and large farms. *Science*, *351*(6271), 388–391.

Garratt, M. P., Bishop, J., Degani, E., Potts, S. G., Shaw, R. F., Shi, A., & Roy, S. (2018). Insect pollination as an agronomic input: Strategies for oilseed rape production. *Journal of Applied Ecology*, *55*(6), 2834–2842.

Garrod, G., Ruto, E., Willis, K., & Powe, N. (2012). Heterogeneity of preferences for the benefits of environmental stewardship: A latent-class approach. *Ecological Economics*, *76*, 104–111.

Gerrish, J. (2020). Fair winds: Enforcement of the good neighbor provision after wisconsin v. epa. *Ecology LQ*, *47*, 691.

Ghazoul, J. (2005a). Buzziness as usual? questioning the global pollination crisis. *Trends in ecology & evolution*, *20*(7), 367–373.

Ghazoul, J. (2005b). Response to steffan-dewenter et al.: Questioning the global pollination crisis. *Trends in Ecology & Evolution*, *20*(12), 652–653.

Glasgow, D., & Zhao, S. (2017). Has the clean air interstate rule fulfilled its mission? an assessment of federal rule-making in preventing regional spillover pollution. *Review of Policy Research*, *34*(2), 186–207.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*.

Granado-Diaz, R., Villanueva, A. J., & Colombo, S. (2024). Land manager preferences for outcome-based payments for environmental services in oak savannahs. *Ecological Economics*, *220*, 108158.

Greene, W. H., & Hensher, D. A. (2003). A latent class model for discrete choice analysis: Contrasts with mixed logit. *Transportation Research Part B: Methodological*, *37*(8), 681–698.

Griffin, R. C., & Bromley, D. W. (1982). Agricultural runoff as a nonpoint externality: A theoretical development. *American Journal of Agricultural Economics*, *64*(3), 547–552.

Groves, R. M., Singer, E., & Corning, A. (2000). Leverage-saliency theory of survey participation: Description and an illustration. *The Public Opinion Quarterly*, *64*(3), 299–308.

Gruber, J. (1994). The incidence of mandated maternity benefits. *The American Economic Review*, 622–641.

Haghani, M., Bliemer, M. C., Rose, J. M., Oppewal, H., & Lancsar, E. (2021). Hypothetical bias in stated choice experiments: Part i. macro-scale analysis of literature and integrative synthesis of empirical evidence from applied economics, experimental psychology and neuroimaging. *Journal of choice modelling*, *41*, 100309.

Halinski, R., Garibaldi, L. A., dos Santos, C. F., Acosta, A. L., Guidi, D. D., & Blochtein, B. (2020). Forest fragments and natural vegetation patches within crop fields contribute to higher oilseed rape yields in brazil. *Agricultural Systems*, *180*, 102768.

Hall, J., & Pretty, J. (2008). Then and now: Norfolk farmers' changing relationships and linkages with government agencies during transformations in land management. *Journal of Farm Management*, *13*(6), 393–418.

Halvorsen, R., & Palmquist, R. (1980). The interpretation of dummy variables in semilogarithmic equations. *American economic review*, *70*(3).

Hanley, N., & Perrings, C. (2019). The economic value of biodiversity. *Annual Review of Resource Economics*, *11*, 355–375.

Hanley, N., Wright, R. E., & Adamowicz, V. (1998). Using choice experiments to value the environment. *Environmental and resource economics*, *11*, 413–428.

Hanski, I. (1994). A practical model of metapopulation dynamics. *Journal of animal ecology*, 151–162.

Harding, T., Herzberg, J., & Kuralbayeva, K. (2021). Commodity prices and robust environmental regulation: Evidence from deforestation in brazil. *Journal of Environmental Economics and Management*, *108*, 102452.

Harstad, B., & Eskeland, G. S. (2010). Trading for the future: Signaling in permit markets. *Journal of public economics*, *94*(9-10), 749–760.

Hausman, J., & McFadden, D. (1984). Specification tests for the multinomial logit model. *Econometrica: Journal of the econometric society*, 1219–1240.

Häussler, J., Sahlin, U., Baey, C., Smith, H. G., & Clough, Y. (2017). Pollinator population size and pollination ecosystem service responses to enhancing floral and nesting resources. *Ecology and evolution*, *7*(6), 1898–1908.

Heal, G. M. (2000). *Nature and the marketplace: Capturing the value of ecosystem services*. Island Press.

Heckman, J. J., Ichimura, H., Smith, J., & Todd, P. (1996). Sources of selection bias in evaluating social programs: An interpretation of conventional measures and evidence on the effectiveness of matching as a program evaluation method. *Proceedings of the National Academy of Sciences*, *93*(23), 13416–13420.

Heo, S. W., Ito, K., & Kotamarthi, R. (2023). *International spillover effects of air pollution: Evidence from mortality and health data* (tech. rep.). National Bureau of Economic Research.

Hess, S., Hensher, D. A., & Daly, A. (2012). Not bored yet–revisiting respondent fatigue in stated choice experiments. *Transportation research part A: policy and practice*, *46*(3), 626–644.

Hess, S., & Palma, D. (2019). Apollo: A flexible, powerful and customisable freeware package for choice model estimation and application. *Journal of choice modelling*, *32*, 100170.

Hess, S., Train, K. E., & Polak, J. W. (2006). On the use of a modified latin hypercube sampling (mlhs) method in the estimation of a mixed logit model for vehicle choice. *Transportation Research Part B: Methodological*, *40*(2), 147–163.

Hintermann, B. (2017). Market power in emission permit markets: Theory and evidence from the eu ets. *Environmental and Resource Economics*, *66*, 89–112.

Holland, S. P., & Yates, A. J. (2015). Optimal trading ratios for pollution permit markets. *Journal of Public Economics*, *125*, 16–27.

Holmes, T. P., & Adamowicz, W. L. (2003). Attribute-based methods. In *A primer on nonmarket valuation* (pp. 171–219). Springer.

Holstead, K. L., Kenyon, W., Rouillard, J. J., Hopkins, J., & Galán-Díaz, C. (2017). Natural flood management from the farmer's perspective: Criteria that affect uptake. *Journal of Flood Risk Management*, *10*, 205–218. https://doi.org/DOI10.1111/jfr3.12129

Hoyos, D. (2010). The state of the art of environmental valuation with discrete choice experiments. *Ecological economics*, *69*(8), 1595–1603.

Hurley, P., Lyon, J., Hall, J., Little, R., Tsouvalis, J., White, V., & Rose, D. C. (2022). Co-designing the environmental land management scheme in england: The why, who and how of engaging 'harder to reach' stakeholders. *People and Nature*.

Image, M., Gardner, E., & Breeze, T. D. (2023). Co-benefits from tree planting in a typical english agricultural landscape: Comparing the relative effectiveness of hedgerows, agroforestry and woodland creation for improving crop pollination services. *Land Use Policy*, *125*, 106497.

Image, M., Gardner, E., Clough, Y., Smith, H. G., Baldock, K. C., Campbell, A., Garratt, M., Gillespie, M. A., Kunin, W. E., McKerchar, M., et al. (2022). Does agri-environment scheme participation in england increase pollinator populations and crop pollination services? *Agriculture, Ecosystems & Environment*, *325*, 107755.

Jaramillo, P., & Muller, N. Z. (2016). Air pollution emissions and damages from energy production in the us: 2002–2011. *Energy Policy*, *90*, 202–211.

Jauker, B., Krauss, J., Jauker, F., & Steffan-Dewenter, I. (2013). Linking life history traits to pollinator loss in fragmented calcareous grasslands. *Landscape Ecology*, *28*, 107–120.

Jia, R., Shao, S., & Yang, L. (2021). High-speed rail and co2 emissions in urban china: A spatial difference-in-differences approach. *Energy Economics*, *99*, 105271.

Johnston, R. J., Boyle, K. J., Adamowicz, W., Bennett, J., Brouwer, R., Cameron, T. A., Hanemann, W. M., Hanley, N., Ryan, M., Scarpa, R., et al. (2017). Contemporary guidance for stated preference studies. *Journal of the Association of Environmental and Resource Economists*, *4*(2), 319–405.

Jones, R. F., Kam, H., & Potter, C. (2023). Are landholders willing to collaborate under elms? promoting collaborative conservation on a landscape scale in the uk. *Journal of Rural Studies*, *103*, 103109.

Joskow, P. L. (2013). Natural gas: From shortages to abundance in the united states. *American Economic Review*, *103*(3), 338–343.

Kahn-Lang, A., & Lang, K. (2020). The promise and pitfalls of differences-in-differences: Reflections on 16 and pregnant and other applications. *Journal of Business & Economic Statistics*, *38*(3), 613–620.

Kampas, A., Melfou, K., & Aftab, A. (2013). Designing regulatory policies for complex externalities: The case of agricultural pollution. *Agricultural Economics Review*, *14*(389-2016-23495), 75–88.

Känzig, D. R. (2023). *The unequal economic consequences of carbon pricing* (tech. rep.). National Bureau of Economic Research.

Kellogg, R., & Wolff, H. (2008). Daylight time and energy: Evidence from an australian experiment. *Journal of Environmental Economics and Management*, *56*(3), 207–220.

Kenyon, W. (2007). Evaluating flood risk management options in scotland: A participant-led multi-criteria approach. *Ecological Economics*, *64*, 70–81.

Kleftodimos, G., Gallai, N., Rozakis, S., & Kephaliacos, C. (2021). A farm-level ecological-economic approach of the inclusion of pollination services in arable crop farms. *Land Use Policy*, *107*, 105462.

Kleijn, D., Baquero, R., Clough, Y., Díaz, M., De Esteban, J., Fernández, F., Gabriel, D., Herzog, F., Holzschuh, A., Jöhl, R., et al. (2006). Mixed biodiversity benefits of agri-environment schemes in five european countries. *Ecology letters*, *9*(3), 243–254.

Kleijn, D., Winfree, R., Bartomeus, I., Carvalheiro, L. G., Henry, M., Isaacs, R., Klein, A.-M., Kremen, C., M'gonigle, L. K., Rader, R., et al. (2015). Delivery of crop pollination services is an insufficient argument for wild pollinator conservation. *Nature communications*, *6*(1), 1–9.

Klein, A.-M., Vaissière, B. E., Cane, J. H., Steffan-Dewenter, I., Cunningham, S. A., Kremen, C., & Tscharntke, T. (2007). Importance of pollinators in changing landscapes for world crops. *Proceedings of the royal society B: biological sciences*, *274*(1608), 303–313.

Kling, C. L. (2011). Economic incentives to improve water quality in agricultural landscapes: Some new variations on old ideas. *American Journal of Agricultural Economics*, *93*(2), 297–309.

Krämer, J. E., & Wätzold, F. (2018). The agglomeration bonus in practice—an exploratory assessment of the swiss network bonus. *Journal for Nature Conservation*, *43*, 126–135.

Kruse, E. (2009). North carolina v. environmental protection agency. *Harv. Envtl. L. Rev.*, *33*, 283.

Kubiszewski, I., Costanza, R., Anderson, S., & Sutton, P. (2020). The future value of ecosystem services: Global scenarios and national implications. In *Environmental assessments* (pp. 81–108). Edward Elgar Publishing.

Kuhfuss, L., Préget, R., Thoyer, S., & Hanley, N. (2016). Nudging farmers to enrol land into agri-environmental schemes: The role of a collective bonus. *European Review of Agricultural Economics*, *43*(4), 609–636.

Lancaster, K. J. (1966). A new approach to consumer theory. *Journal of political economy*, *74*(2), 132–157.

Le Féon, V., Burel, F., Chifflet, R., Henry, M., Ricroch, A., Vaissière, B. E., & Baudry, J. (2013). Solitary bee abundance and species richness in dynamic agricultural landscapes. *Agriculture, Ecosystems & Environment*, *166*, 94–101.

Leelőssy, Á., Molnár, F., Izsák, F., Havasi, Á., Lagzi, I., & Mészáros, R. (2014). Dispersion modeling of air pollutants in the atmosphere: A review. *Open Geosciences*, *6*(3), 257–278.

Lepais, O., Darvill, B., O'connor, S., Osborne, J. L., Sanderson, R. A., Cussans, J., Goffe, L., & Goulson, D. (2010). Estimation of bumblebee queen dispersal distances using sibship reconstruction method. *Molecular Ecology*, *19*(4), 819–831.

Leppert, D. (2023). "no fences make bad neighbors" but markets make better ones: Cap-and-trade reduces cross-border so2 in a natural experiment. *Environmental Economics and Policy Studies*, 1–27.

Leppert, D., Dalhaus, T., & Lagerkvist, C.-J. (2021). Accounting for geographic basis risk in heat index insurance: How spatial interpolation can reduce the cost of risk. *Weather, Climate, and Society*, *13*(2), 273–286.

Liczner, A. R., & Colla, S. R. (2019). A systematic review of the nesting and overwintering habitat of bumble bees globally. *Journal of Insect Conservation*, *23*(5-6), 787–801.

Little, R. J., & Rubin, D. B. (2019). *Statistical analysis with missing data* (Vol. 793). John Wiley & Sons.

Liu, P., & Tian, X. (2021). Downward hypothetical bias in the willingness to accept measure for private goods: Evidence from a field experiment. *American Journal of Agricultural Economics*, *103*(5), 1679–1699.

Liu, Z., Banerjee, S., Cason, T. N., Hanley, N., Liu, Q., Xu, J., & Kontoleon, A. (2024). Spatially coordinated conservation auctions: A framed field experiment focusing on farmland wildlife conservation in china. *American Journal of Agricultural Economics*, *106*(4), 1354–1379.

Lloyd-Smith, P., & Adamowicz, W. (2018). Can stated measures of willingness-to-accept be valid? evidence from laboratory experiments. *Journal of Environmental Economics and Management*, *91*, 133–149.

Loft, L., Gehrig, S., Salk, C., & Rommel, J. (2020). Fair payments for effective environmental conservation. *Proceedings of the National Academy of Sciences*, *117*(25), 14094–14101.

Lonsdorf, E., Kremen, C., Ricketts, T., Winfree, R., Williams, N., & Greenleaf, S. (2009). Modelling pollination services across agricultural landscapes. *Annals of botany*, *103*(9), 1589–1600.

Loomes, G., Starmer, C., & Sugden, R. (1991). Observing violations of transitivity by experimental methods. *Econometrica: Journal of the Econometric Society*, 425–439.

Lowenberg-DeBoer, J., Behrendt, K., Godwin, R., & Franklin, K. (2019). The impact of swarm robotics on arable farm size and structure in the uk.

Lucas, A., Bull, J. C., De Vere, N., Neyland, P. J., & Forman, D. W. (2017). Flower resource and land management drives hoverfly communities and bee abundance in seminatural and agricultural grasslands. *Ecology and Evolution*, *7*(19), 8073–8086.

Maler, K.-G. (1989). The acid rain game. In *Studies in environmental science* (pp. 231–252, Vol. 36). Elsevier.

Mamine, F., Minviel, J. J., et al. (2020). Contract design for adoption of agrienvironmental practices: A meta-analysis of discrete choice experiments. *Ecological Economics*, *176*, 106721.

Marini, L., Öckinger, E., Bergman, K.-O., Jauker, B., Krauss, J., Kuussaari, M., Pöyry, J., Smith, H. G., Steffan-Dewenter, I., & Bommarco, R. (2014). Contrasting effects of habitat area and connectivity on evenness of pollinator communities. *Ecography*, *37*(6), 544–551.

Marsden, T., & Sonnino, R. (2008). Rural development and the regional state: Denying multifunctional agriculture in the uk. *Journal of Rural Studies*, *24*(4), 422–431.

Marston, C. G., O'Neil, A. W., Morton, R. D., Wood, C. M., & Rowland, C. S. (2023). Lcm2021–the uk land cover map 2021. *Earth System Science Data*, *15*(10), 4631–4649.

Matthews, Y., Scarpa, R., & Marsh, D. (2017). Stability of willingness-to-pay for coastal management: A choice experiment across three time periods. *Ecological Economics*, *138*, 64–73.

McCubbin, P. R. (2009). Cap and trade programs under the clean air act: Lessons from the clean air interstate rule and the nox sip call. *Penn St. Envtl. L. Rev.*, *18*, 1.

McFadden, D. (1974). The measurement of urban travel demand. *Journal of public economics*, *3*(4), 303–328.

Mendelsohn, R. (1980). An economic analysis of air pollution from coal-fired power plants. *Journal of Environmental Economics and Management*, *7*(1), 30–43.

Menne, M. J., Durre, I., Vose, R. S., Gleason, B. E., & Houston, T. G. (2012). An overview of the global historical climatology network-daily database. *Journal of atmospheric and oceanic technology*, *29*(7), 897–910.

Mola, J. M., Hemberger, J., Kochanski, J., Richardson, L. L., & Pearse, I. S. (2021). The importance of forests in bumble bee biology and conservation. *Bioscience*, *71*(12), 1234–1248.

Montgomery, W. D. (1972). Markets in licenses and efficient pollution control programs. *Journal of economic theory*, *5*(3), 395–418.

Montoya, D., Haegeman, B., Gaba, S., De Mazancourt, C., & Loreau, M. (2021). Habitat fragmentation and food security in crop pollination systems. *Journal of Ecology*, *109*(8), 2991–3006.

Nguyen, C., Latacz-Lohmann, U., Hanley, N., & Iftekhar, S. (2025). Conservation auctions for landscape-scale environmental management: Does spatial configuration matter for economic and ecological outcomes? *Ecological Economics*, *230*, 108509.

Nguyen, C., Latacz-Lohmann, U., Hanley, N., Schilizzi, S., & Iftekhar, S. (2022). Spatial coordination incentives for landscape-scale environmental management: A systematic review. *Land Use Policy*, *114*, 105936.

Niehoff, D., Fritsch, U., & Bronstert, A. (2002). Land-use impacts on storm-runoff generation: Scenarios of land-use change and simulation of hydrological response in a meso-scale catchment in sw-germany. *Journal of hydrology*, *267*(1-2), 80–93.

Nijkamp, P., Vindigni, G., & Nunes, P. A. (2008). Economic valuation of biodiversity: A comparative study. *Ecological economics*, *67*(2), 217–231.

Olsen, S. B. (2009). Choosing between internet and mail survey modes for choice experiment surveys considering non-market goods. *Environmental and Resource Economics*, *44*(4), 591–610.

Osborne, J. L., Martin, A. P., Shortall, C. R., Todd, A. D., Goulson, D., Knight, M. E., Hale, R. J., & Sanderson, R. A. (2008). Quantifying and comparing bumblebee nest densities in gardens and countryside habitats. *Journal of Applied Ecology*, *45*(3), 784–792.

Parkhurst, G. M., & Shogren, J. F. (2007). Spatial incentives to coordinate contiguous habitat. *Ecological economics*, *64*(2), 344–355.

Pasquill, F. (1961). The estimation of the dispersion of windborne material. *Met. Mag.*, *90*, 33.

Pattison, I., & Lane, S. N. (2012). The link between land-use management and fluvial flood risk: A chaotic conception? *Progress in Physical Geography*, *36*(1), 72–92.

PEARSON, C., et al. (2016). *Modelling the potential impact of spatially targeted natural flood management at the landscape scale for a rural uk catchment.* [Doctoral dissertation, Durham University].

Pearson, C., Reaney, S., Perks, M., Hortobagyi, B., Rosser, N., & Large, A. (2022). Identification of floodwater source areas in nepal using scimap-flood. *Journal of Flood Risk Management*, *15*(4), e12840.

Peterson, J. M., Smith, C. M., Leatherman, J. C., Hendricks, N. P., & Fox, J. A. (2015). Transaction costs in payment for environmental service contracts. *American Journal of Agricultural Economics*, *97*(1), 219–238.

Pinzon, E. (2021). Differences-in-differences in stata 17. *London Stata Conference 2021*, (21).

Pleune, J. G. (2006). Do we cair about cooperative federalism in the clean air act? *Utah L. Rev.*, 537.

Porto, R. G., de Almeida, R. F., Cruz-Neto, O., Tabarelli, M., Viana, B. F., Peres, C. A., & Lopes, A. V. (2020). Pollination ecosystem services: A comprehensive review of economic values, research funding and policy actions. *Food Security*, *12*(6), 1425–1442.

Posthumus, H., Hewett, C., Morris, J., & Quinn, P. (2008). Agricultural land use and flood risk management: Engaging with stakeholders in north yorkshire. *Agricultural Water Management*, *95*, 787–798.

Posthumus, H., & Morris, J. (2010). Implications of cap reform for land management and runoff control in england and wales. *Land Use Policy*, *27*(1), 42–50.

Potoski, M. (2001). Clean air federalism: Do states race to the bottom? *Public Administration Review*, *61*(3), 335–343.

Potts, S. G., Biesmeijer, J. C., Kremen, C., Neumann, P., Schweiger, O., & Kunin, W. E. (2010). Global pollinator declines: Trends, impacts and drivers. *Trends in ecology & evolution*, *25*(6), 345–353.

Potts, S. G., Ngo, H. T., Biesmeijer, J. C., Breeze, T. D., Dicks, L. V., Garibaldi, L. A., Hill, R., Settele, J., & Vanbergen, A. (2016). The assessment report of the intergovernmental science-policy platform on biodiversity and ecosystem services on pollinators, pollination and food production.

Powney, G. D., August, T. A., Harrower, C. A., Outhwaite, C., & Isaac, N. J. (2021). Uk biodiversity indicators 2021. *UK Government*.

Powney, G. D., Carvell, C., Edwards, M., Morris, R. K., Roy, H. E., Woodcock, B. A., & Isaac, N. J. (2019). Widespread losses of pollinating insects in britain. *Nature Communications*, *10*(1), 1018.

Qualtrics. (2020). https://www.qualtrics.com/uk/

Quick, J. C. (2014). Carbon dioxide emission tallies for 210 us coal-fired power plants: A comparison of two accounting methods. *Journal of the Air & Waste Management Association*, *64*(1), 73–79.

Reaney, S. M. (2022). Spatial targeting of nature-based solutions for flood risk management within river catchments. *Journal of Flood Risk Management*, e12803.

Reid, J., Crone, J., & Hayes, J. (2017). Geographic boundary data and their online access. In *The routledge handbook of census resources, methods and applications* (pp. 90–109). Routledge.

Ren, P., Didham, R. K., Murphy, M. V., Zeng, D., Si, X., & Ding, P. (2023). Forest edges increase pollinator network robustness to extinction with declining area. *Nature Ecology & Evolution*, *7*(3), 393–404.

Riley, M., Sangster, H., Smith, H., Chiverrell, R., & Boyle, J. (2018). Will farmers work together for conservation? the potential limits of farmers' cooperation in agri-environment measures. *Land use policy*, *70*, 635–646.

Rose, J. M., Bliemer, M. C., Hensher, D. A., & Collins, A. T. (2008). Designing efficient stated choice experiments in the presence of reference alternatives. *Transportation Research Part B: Methodological*, *42*(4), 395–406.

Rowland, C., Marston, C., Morton, R., & O'Neil, A. (2020). Land cover change 1990-2015 (25m raster, gb). *NERC Environmental Information Data Centre.*

Rural Payments Agency. (2021). *2020 crop map of england (crome)* (tech. rep.) (https://www.data.gov.uk/dataset/be5d88c9-acfb-4052-bf6b-ee9a416cfe60/crop-map-of-england-crome-2020). Rural Payments Agency.

Ruto, E., & Garrod, G. (2009). Investigating farmers' preferences for the design of agri-environment schemes: A choice experiment approach. *Journal of environmental planning and management*, *52*(5), 631–647.

Saitone, T. L., & Sexton, R. J. (2010). Product differentiation and quality in food markets: Industrial organization implications. *Annu. Rev. Resour. Econ.*, *2*(1), 341–368.

Saura, S., & Pascual-Hortal, L. (2007). A new habitat availability index to integrate connectivity in landscape conservation planning: Comparison with existing indices and application to a case study. *Landscape and urban planning*, *83*(2-3), 91–103.

Scarpa, R., Thiene, M., & Train, K. (2008). Utility in willingness to pay space: A tool to address confounding random scale effects in destination choice to the alps. *American journal of agricultural economics*, *90*(4), 994–1010.

Schennach, S. M. (2000). The economics of pollution permit banking in the context of title iv of the 1990 clean air act amendments. *Journal of Environmental Economics and Management*, *40*(3), 189–210.

Schlenker, W., & Walker, W. R. (2016). Airports, air pollution, and contemporaneous health. *The Review of Economic Studies*, *83*(2), 768–809.

Schmalensee, R., & Stavins, R. N. (2013). The so2 allowance trading system: The ironic history of a grand policy experiment. *Journal of Economic Perspectives*, *27*(1), 103–22.

Schmalensee, R., & Stavins, R. N. (2017). Lessons learned from three decades of experience with cap and trade. *Review of Environmental Economics and Policy*.

Seinfeld, J. H., & Pandis, S. N. (2016). *Atmospheric chemistry and physics: From air pollution to climate change*. John Wiley & Sons Inc.

Senapathi, D., Goddard, M. A., Kunin, W. E., & Baldock, K. C. (2017). Landscape impacts on pollinator communities in temperate systems: Evidence and knowledge gaps. *Functional Ecology*, *31*(1), 26–37.

Shapiro, J. S., & Walker, R. (2018). Why is pollution from us manufacturing declining? the roles of environmental regulation, productivity, and trade. *American Economic Review*, *108*(12), 3814–54.

Sheremet, O., Ruokamo, E., Juutinen, A., Svento, R., & Hanley, N. (2018). Incentivising participation and spatial coordination in payment for ecosystem service schemes: Forest disease control programs in finland. *Ecological Economics*, *152*, 260–272.

Shouse, K. C. (2018). The clean air act's good neighbor provision: Overview of interstate air pollution control. *Congressional Research Service R45299*.

Slonim, R., Wang, C., Garbarino, E., & Merrett, D. (2013). Opting-in: Participation bias in economic experiments. *Journal of Economic Behavior & Organization*, *90*, 43–70.

Spicer, E. A., Swaffield, S., & Moore, K. (2021). Agricultural land use management responses to a cap and trade regime for water quality in lake taupo catchment, new zealand. *Land use policy*, *102*, 105200.

Stanley, D. A., Gunning, D., & Stout, J. C. (2013). Pollinators and pollination of oilseed rape crops (brassica napus l.) in ireland: Ecological and economic incentives for pollinator conservation. *Journal of insect conservation*, *17*(6), 1181–1189.

Stavins, R. N. (1995). Transaction costs and tradeable permits. *Journal of environmental economics and management*, *29*(2), 133–148.

Stavins, R. N. (2003). Experience with market-based environmental policy instruments. In *Handbook of environmental economics* (pp. 355–435, Vol. 1). Elsevier.

Steffan-Dewenter, I., Potts, S. G., & Packer, L. (2005). Pollinator diversity and crop pollination services are at risk. *Trends in ecology & evolution*, *20*(12), 651–652.

Steffan-Dewenter, I., & Tscharntke, T. (1999). Effects of habitat isolation on pollinator communities and seed set. *Oecologia*, *121*, 432–440.

Stranlund, J. K., & Chavez, C. A. (2000). Effective enforcement of a transferable emissions permit system with a self-reporting requirement. *Journal of Regulatory Economics*, *18*(2), 113–131.

Sutton, O. (1947). The problem of diffusion in the lower atmosphere. *Quarterly Journal of the Royal Meteorological Society*, *73*(317-318), 257–281.

Tait, M. (2009). A remedy even the plaintiffs don't like, the dc circuit's vacatur of the clean air interstate rule. *Mo. Envtl. L. & Pol'y Rev.*, *16*, 552.

Tanguy, M., Dixon, H., Prosdocimi, I., Morris, D., & Keller, V. (2021). Gridded estimates of daily and monthly areal rainfall for the united kingdom (1890-2019)[ceh-gear]. *Environmental Information Data Centre*, *10*.

Thompson, S. (2022). *Cereal and oilseed rape areas in england at 1 june 2022* (tech. rep.) (https://www.gov.uk/government/statistics/cereal-and-oilseed-rape-areas-in-england/cereal-and-oilseed-rape-areas-in-england-at-1-june-2022). Defra.

Tietenberg, T. H. (1990). Economic instruments for environmental regulation. *Oxford review of economic policy*, *6*(1), 17–33.

Timberlake, T. P., Vaughan, I. P., & Memmott, J. (2019). Phenology of farmland floral resources reveals seasonal gaps in nectar availability for bumblebees. *Journal of Applied Ecology*, *56*(7), 1585–1596.

Tobin, J. (1958). Estimation of relationships for limited dependent variables. *Econometrica: journal of the Econometric Society*, 24–36.

Train, K., & Weeks, M. (2005). *Discrete choice models in preference space and willingness-to-pay space*. Springer.

Tyllianakis, E., Martin-Ortega, J., Ziv, G., Chapman, P. J., Holden, J., Cardwell, M., & Fyfe, D. (2023). A window into land managers' preferences for new forms of agri-environmental schemes: Evidence from a post-brexit analysis. *Land Use Policy*, *129*, 106627.

UK Department for Environment, Food and Rural Affairs. (2022). *Agriculture in the uk evidence pack: 2022 update* (tech. rep.) (https://assets.publishing.

service.gov.uk/government/uploads/system/uploads/attachment_data/file/1106562/AUK_Evidence_Pack_2021_Sept22.pdf). Defra.

UK Environment Agency. (n.d.). Lidar composite digital terrain model (dtm) 10m [Accessed: 2024-12-29].

Ulveling, E. F., & Fletcher, L. B. (1970). A cobb-douglas production function with variable returns to scale. *American Journal of Agricultural Economics*, *52*(2), 322–326.

Upcott, E. V., Henrys, P. A., Redhead, J. W., Jarvis, S. G., & Pywell, R. F. (2023). A new approach to characterising and predicting crop rotations using national-scale annual crop maps. *Science of The Total Environment*, *860*, 160471.

U.S. EPA. (1989). *User's guide to the complex terrain dispersion model plus algorithms for unstable situations (ctdmplus)* (tech. rep.). U.S. Environmental Protection Agency.

U.S. EPA. (2019). Review of the primary national ambient air quality standards for sulfur oxides. https://www.federalregister.gov/documents/2019/03/18/2019-03855/review-of-the-primary-national-ambient-air-quality-standards-for-sulfur-oxides

Van Buuren, S., & Groothuis-Oudshoorn, K. (2011). Mice: Multivariate imputation by chained equations in r. *Journal of statistical software*, *45*, 1–67.

Vatn, A., & Bromley, D. W. (1994). Choices without prices without apologies. *Journal of environmental economics and management*, *26*(2), 129–148.

Velázquez, B., Buffaria, B., & Commission, E. (2017). About farmers' bargaining power within the new cap. *Agricultural and Food Economics*, *5*, 1–13.

Villanueva, A. J., Barreiro-Hurlé, J., & Rodríguez-Entrena, M. (2025). Impact of consequentiality in willingness to accept: Evidence from a choice experiment with land managers. *Agricultural and Food Economics*, *13*(1), 60.

Villanueva, A. J., Glenk, K., & Rodríguez-Entrena, M. (2017). Protest responses and willingness to accept: Ecosystem services providers' preferences towards incentive-based schemes. *Journal of Agricultural Economics*, *68*(3), 801–821.

Vossler, C. A., Doyon, M., & Rondeau, D. (2012). Truth in consequentiality: Theory and field evidence on discrete choice experiments. *American Economic Journal: Microeconomics*, *4*(4), 145–171.

Waxman, H. A. (1991). An overview of the clean air act amendments of 1990. *Envtl. L.*, *21*, 1721.

Wei, Y., Gu, J., Wang, H., Yao, T., & Wu, Z. (2018). Uncovering the culprits of air pollution: Evidence from china's economic sectors and regional heterogeneities. *Journal of Cleaner Production*, *171*, 1481–1493.

Welling, M., Zawojska, E., & Sagebiel, J. (2022). Information, consequentiality and credibility in stated preference surveys: A choice experiment on climate adaptation. *Environmental and Resource Economics*, *82*(1), 257–283.

Weninger, T., Scheper, S., Lackóová, L., Kitzler, B., Gartner, K., King, N., Cornelis, W., Strauss, P., & Michel, K. (2021). Ecosystem services of tree windbreaks in rural landscapes—a systematic review. *Environmental Research Letters*, *16*(10), 103002.

Wilkinson, M., Addy, S., Quinn, P., & Stutter, M. (2019). Natural flood management: Small-scale progress and larger-scale challenges. *Scottish Geographical Journal*, *135*(1-2), 23–32.

Winfree, R., Williams, N. M., Gaines, H., Ascher, J. S., & Kremen, C. (2008). Wild bee pollinators provide the majority of crop visitation across land-use gradients in new jersey and pennsylvania, usa. *Journal of applied ecology*, *45*(3), 793–802.

Wingfield, T., Macdonald, N., Peters, K., Spees, J., & Potter, K. (2019). Natural flood management: Beyond the evidence debate. *Area*, *51*(4), 743–751.

Wintrup, J. (2022). Promising careers? a critical analysis of a randomised control trial in community health worker recruitment in zambia. *Social Science & Medicine*, *299*, 114412.

Wooldridge, J. (2007). What's new in econometrics? lecture 10 difference-in-differences estimation. *NBER Summer Institute, available at: www.nber.org/WNE/Slides7–31–07/slides/10/iffindiffs.pdf, accessed April*, *9*(2011), 85.

World Health Organization. (2006). *Air quality guidelines: Global update 2005: Particulate matter, ozone, nitrogen dioxide, and sulfur dioxide*.

Xepapadeas, A. (1992a). Optimal taxes for pollution regulation: Dynamic, spatial and stochastic characteristics. *Natural Resource Modeling*, *6*(2), 139–170.

Xepapadeas, A., et al. (1997). *Advanced principles in environmental policy*. Edward Elgar Publishing Ltd.

Xepapadeas, A. P. (1992b). Environmental policy design and dynamic nonpoint-source pollution. *Journal of environmental economics and management*, *23*(1), 22–39.

Xiao, Y., Li, X., Cao, Y., & Dong, M. (2016). The diverse effects of habitat fragmentation on plant–pollinator interactions. *Plant Ecology*, *217*, 857–868.

Yadav, M. L., & Roychoudhury, B. (2018). Handling missing values: A study of popular imputation packages in r. *Knowledge-Based Systems*, *160*, 104–118.

Zannetti, P. (2013). *Air pollution modeling: Theories, computational methods and available software*. Springer Science & Business Media.

Zereyesus, Y. A., & Featherstone, A. M. (2017). Empirical analysis of profit maximization and cost minimization behaviour of kansas farms. *Applied Economics Letters*, *24*(17), 1255–1258.

Zereyesus, Y. A., Featherstone, A. M., & Langemeier, M. R. (2021). Are kansas farms profit maximizers? a stochastic additive error approach. *Agricultural Economics*, *52*(1), 37–50.

Zhang, J., & Adamowicz, W. L. (2011). Unraveling the choice format effect: A context-dependent random utility model. *Land Economics*, *87*(4), 730–743.

Zheng, S., Cao, J., Kahn, M. E., & Sun, C. (2014). Real estate valuation and cross-boundary air pollution externalities: Evidence from chinese cities. *The Journal of Real Estate Finance and Economics*, *48*(3), 398–414.

Zulian, G., Maes, J., & Paracchini, M. L. (2013). Linking land cover data and crop yields for mapping and assessment of pollination services in europe. *Land*, *2*(3), 472–492.