Applications of Quantitative Methods in Environmental Economics: Econometrics, Simulation Modelling and Experiments

Gregory Steven Smith

For the Degree of Doctor of Philosophy

University of East Anglia

School of Environmental Sciences

January 2016

ABSTRACT

In Part I of this thesis we employ novel econometric techniques to explore elicitation anomalies in contingent valuation (CV). According to standard assumptions regarding preferences, changes in the way values are elicited in CV questions should be decision-irrelevant. That responses are observed to systematically differ according to elicitation format has, therefore, called the CV method into question. One possible explanation lies in the proposition that respondents are uncertain about their preferences and that their uncertainty precipitates systematically different responses to different question formats. We test this hypothesis using data from a split-sample CV survey. We analyse our data using an innovative application of a semi-parametric estimator more commonly used for duration modelling in the medical sciences but find that uncertainty alone cannot explain away common elicitation anomalies.

In Part II we employ simulation modelling and experimental techniques to investigate payment for ecosystem services (PES) schemes that involve multiple buyers. In Chapter 2, we explore opportunities for buyers in PES scheme to realise Pareto-improving outcomes through spatial coordination in their independent purchases of changes to land-management practices. We develop a simulation environment imitating a heterogeneous agricultural landscape and using techniques of integer-linear programming solve for outcomes under different institutional arrangements. Our simulations allow us to explore how gains from negotiated or fully-cooperative purchasing differ across different configurations of landscape and buyer objectives. In Chapter 3, we investigate negotiation as a multiple-purchaser ecosystem service procurement mechanism. We design and conduct novel three-person bargaining experiments in which two potential buyers can negotiate not only between each other but also with a seller of ecosystem services. We find that negotiated deals can be reached that are mutually advantageous to all parties. In all treatment scenarios presented, the vast majority of groups are able to reach agreements; in addition, these agreements are reached relatively quickly.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

INTRODUCTION

This thesis focuses on applying some of the core quantitative techniques of environmental economics—econometrics, simulation modelling and experiments—to examine three problems. Part I of this thesis uses econometric techniques to analyse uncertain responses to contingent valuation surveys. Part II of this thesis explores two related topics: Chapter 2 uses simulation modelling techniques to assess the opportunities and barriers for forming multiple-purchaser payment for ecosystem services (PES) schemes; Chapter 3 uses experimental economics techniques to examine the potential workings of a negotiated multiple-purchaser PES scheme.

In the first chapter we utilise novel econometric techniques to explore elicitation anomalies in contingent valuation (CV) surveys. The CV method is commonly used to elicit value estimates for non-market goods, and in particular environmental non-market goods, however the validity of the method has been questioned as responses are observed to systematically differ according to the elicitation format—so called elicitation anomalies. According to standard assumptions regarding preferences, changes in the way values are elicited in CV questions should be decision-irrelevant. Asking respondents of CV surveys to value non-market environmental goods is often a complex task in that it is unlikely that individuals will have previously considered the trade-off between the provision of such goods and money. One explanation which has been proposed (Ready, et al. 1995; Welsh and Poe 1998; Ready, et al. 2001; Flachaire and Hollard 2007) is that respondents are uncertain about their preferences and that their uncertainty precipitates systematically different responses to different question formats. In other words it is the process of requiring individuals to express values in CV surveys as if they had well-defined certain preferences that leads to elicitation anomalies. Testing this hypothesis is the first key contribution of Chapter 1.

The chapter is structured around two common CV elicitation anomalies: (i) the disparity between values elicited using dichotomous choice and those elicited using open-ended CV questions and (ii) starting point bias or anchoring on an initial bid. Values elicited using dichotomous choice (DC) questions have been shown to be consistently higher than those elicited using open ended (OE) questions (Brown, et al. 1996; Champ, et al. 1997; Vossler, et al. 2003). Likewise, it is well-established that when eliciting values using a series of two or more DC questions (for example, the double-bound DC format) responses to later questions are anchored on the initial DC bid (Whitehead 2002; Flachaire and Hollard 2006, 2007).

To assess for evidence of elicitation anomalies when respondents are allowed to express uncertainty we use a large contingent valuation dataset from Suffolk, UK, collected in 2004. Individuals in the study participated in a valuation exercise comprising three tasks: initially respondents were allocated to a treatment group in which they received either a standard open ended or single-bounded dichotomous question; subsequently, they answered a follow-up question in which they could state their level of certainty; finally, all respondents completed a novel payment-ladder style question to establish the range of values over which they were certain and uncertain. The third task, the novel payment-ladder, is closely linked to the multiple-bounded-dichotomous-choice (MBDC) elicitation method; Mahieu et al. (2014) provide a recent summary of MBDC studies. A key difference in our dataset is that respondents state their level of certainty to a semi-continuous payment ladder ranging from £1 to £500 with £1 increments. This allows for more precision when locating the highest values which respondents are certain they would be willing to pay and also the lowest value they are certain they are not willing to pay. If there exists a gap between these values then that forms the range of values a respondent is uncertain about paying—their uncertainty range.

The second key contribution of Chapter 1 is the development of a novel econometric method to investigate responses to CV questions when respondents are allowed to express uncertainty ranges as part of their CV response. In particular we model the size (or width) and location of the uncertainty range. The estimator we describe is adapted from the duration modelling techniques used in the medical literature to statistically analyse the ‘time to event’; for examples see: (Frydman 1995; Commenges 2002; Frydman and Szarek 2009) and for a review of duration modelling see Klein and Moeschberger (1997). Our econometric method is not the first to analyse uncertainty ranges from CV surveys but so far no consensus has yet emerged on the most appropriate method. In Chapter 1 we review the alternative methods, including the Random Valuation Model developed in Wang (1997), a probability based model developed in Evans et al. (2003) and the Latent Threshold Estimator developed in Kobayashi et al. (2012). Our method is closest to the Latent Threshold Estimator in that it analyses the transition between different states of certainty, however it does this in a radically different way by adapting the multi-state duration models used in the medical literature to analyse the progression along a WTP scale rather than a progression through time. This allows statistical analysis over the full range of the WTP value distribution without requiring restrictive parametric assumptions.

As far as we are aware this is the first time that this form of multi-state duration modelling has been used in economic analysis. We employ our estimator in Chapter 1 to test the hypothesis that elicitation anomalies commonly observed in CV studies are the result of asking respondents with uncertain preferences to answer as if those preferences were precisely-defined. Our econometric model allows us to simultaneously explore how the width and the location of the uncertainty range are influenced by the CV elicitation method used (DC or OE) and the initial bid offered if in a DC group. If the expectation of procedural invariance is supported by our data when respondents are allowed to express uncertain preferences, then the prognosis for the CV method is rather encouraging; by allowing for the possibility of uncertainty, CV can elicit preferences that are procedurally invariant and conform to many of the expectations of standard economic theory.

In Chapter 2 of this thesis we develop a sophisticated framework of simulation modelling methods to investigate multiple-buyer payments for ecosystem services (PES) schemes.[1] Recently there has been a significant increase in interest in creating markets, payments or regulations to encourage the production of ecosystem services (Salzman 2005; Engel, et al. 2008; Wunder 2008; Kemkes, et al. 2010; Kinzig, et al. 2011; Defra 2013; Quick, et al. 2013). In Chapter 2, we focus on modelling a voluntary scheme in which landowners are compensated for the ecosystem services they produce—a PES scheme. In PES schemes, the landowner is paid to produce ecosystem

[1] The motivation for studying multiple-purchaser PES schemes in this thesis stems primarily from interest expressed by stakeholders in collaborative work undertaken by the University of East Anglia, The Department of Environment Food and Rural Affairs (Defra) and a number of water companies in the UK, see Defra reports Day and Couldrick (2013) and Day et al. (Forthcoming).

services, commonly this involves the farmer undertaking an alternative land-management practice, those alternative land uses are costly; possibly requiring additional expenditure in land management or resulting in a lower yield of agricultural output. At the same time, those changes often deliver flows in one or more ecosystem service. Importantly, those flows may accrue to different beneficiary groups each of whom might be prepared to contribute to payments made through the scheme. Furthermore, those flows depend on the spatial pattern of land use—which landowners are in the PES scheme and what land-management change they are undertaking.

Designing a method which can find the spatial pattern of land-use that produces the maximum amount of ecosystem service(s) is a complex task but one that has been previously studied. Recently, Polasky et al. (2014) study this from the point of view of a regulator with limited knowledge of landowner's costs, proposing a special type of auction which incentivises the landowners to truthfully reveal their costs, this allows the regulator to select the optimal spatial pattern of land-use. In addition, the conservation biology literature contains many examples that focus specifically on biodiversity, for a review see Williams et al. (2005). Previously the majority of PES literature has largely concentrated on the single-purchaser problem; Chapter 2 differs markedly by focusing on the issue of PES mechanism design when the activity incentivised through the scheme benefits multiple independent groups—multiple-purchaser PES schemes.

Of course, multiple-purchaser PES schemes are a subset of all PES schemes. In some situations single-purchaser schemes may be more appropriate, in other situations a single buyer, such as the government, may act on behalf of multiple beneficiary groups. However, certainly in the UK, there is interest in developing multiple-purchaser PES schemes, this is captured in the following quote from the Department for Environment Food and Rural Affairs (Defra) *"There is a need to explore new means to aggregate demand from beneficiaries and mobilise funding solutions" "These approaches … draw in multiple sources of funding and strengthen the overall economic case for action"* p23. (Defra 2013).

There exists a wide variety of potential buyers of ecosystem services, realistically any group or organisation that benefits from an increase in ecosystem service flows is a potential buyer, for example, national or local governments, NGO's, environmental groups or private companies, such as water companies. Historically, governments have been considered the main buyers of ecosystem services, predominately through agri-environmental schemes (FAO 2007). The largest schemes are currently in the US and EU. The US, in 2013, spent just less than $6 billion on conservation programmes, with approximately a third of the spending on the largest scheme—the Conservation Reserve Program (USDA 2014). Agri-environmental schemes were introduced into European policy in the late 1980s (Regulation (EEC) No 797/85) and since 1992 have been compulsory for member states. Initially agri-environmental schemes were included as an "accompanying measure" to the Common Agricultural Policy (CAP) reform and later became a dedicated regulation (Regulation (EEC) No 2078/92). Since that time, member states have been required to introduce agri-environmental measures throughout their land; with the aim to limit risks to the environment and promote biodiversity and preserve cultural landscapes.[2, 3] EU expenditure on agri-environment measures for 2007 - 2013 amounts to nearly €20 billion (European Commission 2014) and in England agri-environmental spending is over £400 million per year (Natural England 2014).

Nevertheless, governments are not the only agents interested in purchasing ecosystem services. It is increasingly recognised that risks to the environment create risks to business; either directly—through the reliance on ecosystem services as inputs to production—or indirectly—through markets or supply chains (TEEB 2012). While for many in the private sector opposition remains to the concept of paying for something that they have not paid for before, for others there is the realisation of potential business benefits. Some private sector companies have direct incentives to protect a natural business input; for example, Vittel in France (Perrot-Maître 2006) and water companies such as Wessex Water, United Utilities and South West Water in the UK (Defra 2013). Other private sector companies may be more interested in offsetting some of their environmentally damaging activities by paying for improvements elsewhere through carbon offsetting or biodiversity offsetting. Still other private sector

---

[2] Member states are required to implement the European regulations into Rural Development Programmes and they are currently drawing up new programmes to begin in 2015 with the previous programmes having ended on 31st December 2013.

[3] The Rural Development Programmes in 2015 will be based on the latest regulation (Regulation (EU) 1305/2013) which repeals (Regulation (EC) 1698/2005).

companies or NGOs may be interested in eco-certification or labelling to improve brand image. Whatever the motivation, it is clear that funding from sources other than governments has the potential to increase further the procurement of ecosystem services and thus the justification for exploring multiple purchaser PES schemes becomes stronger.

In Chapter 2 we explore multiple-purchaser PES schemes by focusing on the issue of spatial coordination on the demand side of the market; that is to say, the question of which beneficiary buys land-management changes on which land parcels. Introducing multiple buyers adds complexity to finding optimal spatial patterns of land-use, moreover, it introduces unique problems such as opportunities for free-riding on other investments[4]. We start with a simple motivating example, in that we assume that while all buyers may be interested in incentivising the same type of land-management change (for example, reducing the intensity of agricultural activity or taking land out of production altogether), it is not necessarily the case that each would choose for those changes to be sited in the same locations. As an example, imagine the differing spatial preferences for a water company interested in paying for land-management changes on land that is likely to lead to water quality improvements (for example, land close to water courses or land with direct drainage into water courses); and a biodiversity buyer, for example the government, who might be interested in creating large contiguous areas of habitat by paying for land-management changes on land close to established reserves, following the principles set out in the Lawton report (Lawton, et al. 2010). The first key contribution of Chapter 2 is to develop a general framework of methods that can incorporate the spatial purchasing decision of multiple PES buyers.

The framework needs to be capable of incorporating different buyers' objectives, for example objectives for different ecosystem service benefits, and include different constraints on those objectives. In addition we need to model how the buyers might come together in a PES purchasing institution. To do that we develop four example decision making institutions: in the first the buyers are independent and make their decisions simultaneously and without regard for the actions of the other buyer; in the second the buyers are independent and make their decision sequentially where the second buyer to decide is aware of the first buyer's purchasing decisions; in the third

[4] Multiple buyers can also lead to the issue of collusion, however this is not studied in this thesis.

the buyers make their buying decisions strategically as the outcome of a process of negotiation; in the fourth the buyers make their decisions cooperatively.

The framework of methods developed to model the buyer's decision making allows us to identify situations in which we might expect a multiple purchaser PES scheme to be practical—this is the second key contribution of Chapter 2. To compare the solutions from the four decision making institutions (outlined above) we use the concept of Pareto efficiency, that is solutions that can make one buyer better off without making another buyer worse off. We explore multiple purchaser PES institutions in two simulation environments. In the first simulation we investigate the effect that correlation in the production of ecosystem services has on the efficiency for the multiple buyers using our four PES purchasing institutions. In the second simulation we investigate a more complex and perhaps more realistic situation, in which we model a catchment landscape comprising agricultural land parcels and a river system. In that simulation we imagine two buyers, one (a water quality buyer) whose benefits rely on the spatial heterogeneity of the landscape—sites closer to the river were more beneficial to a water quality buyer—and another (a biodiversity buyer) whose benefits rely on the spatial interdependency and configuration of the landscape—connected habitats provide more benefits to the biodiversity buyer. Modelling a buyer with spatial interdependency in their benefits necessarily creates a non-linear problem, we show how our framework of methods is capable of creating solutions even for spatially interdependent benefits by forming a linearised version of the buyer's decision problem.

The two simulation environments show how the general framework of methods can be used to assess the opportunities for realising Pareto-improving outcomes through a PES scheme when multiple independent groups stand to benefit from changing landowners' land-management practices. In addition, the method we develop allows us to identify optimal patterns of land use across a spatial landscape, potentially providing a useful tool for both ecosystem services buyers and policy makers—this is the third key contribution of Chapter 2. The decision making problems of the buyers are modelled in such a way as to be solvable by linear integer programming methods allowing for exact optimal solutions to be found over a reasonably large and heterogeneous landscape.

In Chapter 3, we design a novel economic experiment to examine the potential workings of a negotiated multiple-purchaser ecosystem service scheme.

The motivation for this chapter stems primarily from a collaborative project between the University of East Anglia, South West Water, Defra, and Westcountry Rivers Trust. The report, Day and Couldrick (2013), shows a pilot ecosystem service procurement scheme conducted in the River Fowey catchment area. The scheme distributed funds for capital investment on farms to improve water quality and was funded by South West Water's Upstream Thinking Initiative. It explores and contrasts a negotiated scheme ('advisor-led mechanism') with a competitive reverse auction. The authors conclude that the advisor-led mechanism, in which farm advisors go out to visit and negotiate directly with farmers, is recommended for small scale schemes, where the farm advisors have good local knowledge, and known target farms are likely to yield positive outcomes. In contrast, competitive auction mechanisms are recommended for large scale schemes where the buyers have little local knowledge.

The procurement of ecosystem services is possible through a number of different mechanisms, such as fixed price mechanisms, competitive bidding or negotiation. The choice of mechanism for PES schemes will depend on the specific circumstances of any particular scheme. The most appropriate mechanism, in some situations, such as when the number of bidders is high, may be a competitive bidding scheme, in other situations, such as when the details are particularly complex, a negotiated scheme may be recommended (Bajari and Tadelis 2001; Bajari, et al. 2009). Since the exchanges transacted in PES schemes are often complex, negotiation between buyers and sellers may play an important role in certain PES mechanisms; accordingly, our experimental investigation focuses on an exchange process facilitated through the multilateral bargaining of buyers and sellers.

Bilateral negotiated ecosystem service procurement schemes have been successfully implemented both in real world schemes, Perrier-Vittel (Perrot-Maître 2006) and United Utilities UK (Smith, et al. 2013), and in laboratory experiments (Bruce and Clark 2010b,2012). In Chapter 3 of this thesis, we extend the literature on negotiation as an ecosystem service procurement mechanism by moving beyond bilateral negotiation to consider multilateral negotiation. In order to provide clarity and also to keep the experiment computationally manageable for the participants, our experiments involve just three parties to those negotiations—two buyers and one seller. The two potential buyers can negotiate not only between each other but also with a seller of ecosystem services to reach a mutually beneficial outcome. The aim of Chapter 3 is therefore to gather insights as to whether negotiated multiple purchaser PES schemes might be achievable and to explore the factors shaping the division of gains from negotiations between multiple purchasers and sellers in such a scheme.

The experiments are structured as non-cooperative alternating bargaining, in which two buyers alternate in proposing how much each buyer should pay and therefore also how much the seller receives should a deal be agreed, the seller acts as a veto player, able to reject any unsatisfactory deal. If negotiations fail then each participant receives their default payment, for the buyers this is comparable to purchasing their next best alternative, for the seller it is comparable to receiving their normal income. The experiment is conducted over a maximum of 15 rounds of negotiation, although each time a participant rejects an offer there is an increasing risk (presented clearly to the participants) that negotiations will fail and therefore no deal will be agreed.

We use this experimental framework to investigate a number of complexities of the negotiating environment that might typically arise in a PES scheme. First, the degree to which the buyers offer (and the seller accepts) an amount over and above the sellers costs. By including a seller in the negotiation process, buyers not only have to negotiate between themselves regarding how much each might contribute but they must also ensure that that offer is satisfactory for the seller. Second, the degree to which *asymmetry in the gains* enjoyed by the two buyers from a successful transaction affects the outcome of negotiations. Here we imagine that one buyer would benefit more from a PES scheme being implemented. Third, the degree to which *asymmetry in the income* of the two buyers (irrespective of their gains from the transaction) affects the outcome of negotiations. Here we imagine that one buyer might be a large, wealthy organisation and that the relatively less wealthy buyer might be more inclined to free ride on the wealthy buyer's contribution to the PES scheme. Fourth, how negotiations differ under conditions of *incomplete information*. Here we imagine that differences in knowledge exist between the buyers and sellers, for example the seller may know the costs for supplying the environmental output but the buyers might not. Finally, how negotiations evolve when the benefits enjoyed by the buyers from the transaction are not known for sure but are *stochastic* in nature. Here we imagine that the buyers are paying for the seller to undertake an action and therefore they are not entirely certain of the actual environmental output that will be produced, this could be due to unpredictable phenomenon such as weather patterns. This is a common situation for PES schemes but has received relatively little attention in the experimental economics literature. Chapter 3 explores these five issues within our experimental design to not only establish whether participants can successfully negotiate multilateral agreements in such a purchasing setting but also to explore how the gains from successfully-negotiated exchanges are partitioned both between the purchasers and between the purchasers and sellers.

PART I

CHAPTER 1

Uncertainty and elicitation anomalies in contingent valuation: an analysis using a state-dependent semiparametric estimator

1.1. Introduction

Economic research often proceeds under the assumption that individuals hold precisely-defined preferences over all bundles of consumption goods. It is not at all evident however, that this is the case; evidence suggests that individuals' valuations, even for familiar market goods, are uncertain (Roselius 1971; Heiman, et al. 2001; Jin, et al. 2005). The existence of uncertainty in preferences is also evident in various market institutions; for example, in the money-back guarantees offered by retailers that allow customers the opportunity to try the good in their daily routine before deciding whether to keep or return it and also in second-hand markets that allow customers to sell unwanted or bad fitting items. The valuation of non-market goods is often additionally complex in that it is unlikely that individuals will have ever previously considered the trade-off between the provision of those goods and money. Moreover, it is rarely the case that individuals are sufficiently well-informed regarding the benefits of such goods that they could hope to express the value in some single willingness-to-pay (WTP) amount.

Typically, however, attempts to estimate WTP for non-market goods using contingent valuation (CV) make no adjustment for uncertainty; despite evidence to the contrary (Kahneman and Knetsch 1992; Dubourg, et al. 1994; Ready, et al. 1995; Champ, et al. 1997; Dubourg, et al. 1997; Wang 1997; van Kooten, et al. 2001; Ariely, et al. 2003; Akter, et al. 2008; Hanley, et al. 2009). The possibility exists, therefore, that subjects with uncertain preferences may provide unanticipated patterns of response to standard CV questions. Likewise, analysts wrongly assuming certainty in preferences, may interpret those responses incorrectly; for example, in construing CV responses as providing evidence of so-called elicitation anomalies.

The central purpose of this paper is to investigate the claim made by numerous authors, (Ready, et al. 1995; Welsh and Poe 1998; Ready, et al. 2001; Flachaire and Hollard 2007) that commonly observed elicitation anomalies in CV—for example, differences in WTP between open ended (OE) and dichotomous choice (DC) formats, and differences in WTP according to the initial bids in repeated DC formats—arise as a result of asking individuals with uncertain preferences to express their value as if those preferences were precisely-defined[5]. Certainly, the possibility exists that subjects with uncertain preferences may provide unanticipated patterns of response when presented with standard CV questions.

In Section 1.2 we discuss the rationale behind suggestions that uncertainty could explain anomalies in CV; specifically, we look at the effect of requiring individuals with uncertain preferences to answer CV questions in a certain or precise manner. The fundamental position underpinning this argument is that individuals, in the presence of uncertainty, may adopt contrasting heuristics in answering CV questions posed in different ways. The first objective of this chapter is to assess this hypothesis.

To assess the hypothesis we require a CV elicitation technique which allows for the expression of uncertainty. One such method that has been widely applied is the *multiple bounded discrete choice* (MBDC)[6] method (Welsh and Bishop 1993; Welsh and Poe 1998; Alberini, et al. 2003; Evans, et al. 2003; Vossler, et al. 2004; Kobayashi, et al. 2010). The MBDC method presents respondents with an ordered list of bids[7]. For each bid, respondents report on a certainty scale their likelihood of being willing to pay that amount. Accordingly, the MBDC method typically presents respondents with a multiple-bounded choice across bid amounts, consistent with the payment card

[5] Another possible explanation for the observed elicitation anomalies is that there is something specific about the CV method that fails to encourage respondents to accurately or truthfully reveal their preferences. Research has focused on the idea that certain formats of CV elicitation encourage strategic (Carson et al., (2001) or ill-considered responses (Poe and Vossler (2009), Hutchinson et al. (2007).

[6] The *multiple bounded discrete choice* method is also known as the *multiple bounded uncertainty choice* method.

[7] There is mixed evidence that the MBDC method may itself lead to elicitation anomalies. Vossler et al. (2004) assess the MBDC method for bid design effects; specifically, three arrays of bids were varied according to how many high (or low) bids were included, with the maximum and minimum bid kept constant throughout all three arrays. They found no statistical difference between the WTP values elicited from the three samples. Dubourg et al. (1997) and Roach et al. (2002) in similar analyses had differing maximum bids in their bid arrays and specifically tested for range effects; they both found that groups offered a bid array with a higher maximum value had significantly higher WTP estimates. This contrary evidence raises doubt about the procedural invariance properties of the MBDC method.

method[8], combined with a polychotomous choice from a scale of certainty ranging from definitely yes to definitely no. An alternative is the payment ladder approach from Hanley et al. (2009), as in the MBDC method they present respondents with a multiple-bounded choice across bid amounts but with just two choices: “I would definitely pay that amount” and “I would definitely NOT pay that amount”. For each respondent, both the MBDC elicitation method and the payment card method from Hanley et al. (2009) provide data recording a range of values for which that respondent is certain they would pay, we label this the *certainty range*, and a range of values over which they are certain they would not pay. Between those two there may exist a range of values over which they are uncertain—we label this the *uncertainty range*[9].

In this study respondents undertake a single valuation exercise split into three tasks, the final stage of that exercise is an MBDC task, prior to this the respondents are allocated to certain treatment groups, those treatment groups each receive either a single-bounded DC question or an OE question with all groups then answering a follow-up question on their certainty. If, as has been hypothesised, the respondents underlying preferences are uncertain and procedurally invariant then the particular treatment group in which the respondent is assigned should not influence the uncertainty ranges expressed in the MBDC part of the exercise. Of course, individual uncertainty ranges are likely to vary due to individual characteristics such as their experience of the good in question (Ackerberg 2003 and Czajkowski et al. 2015) but on average those differences should be randomised across the treatments. To analyse the procedural invariance of the respondents’ uncertainty ranges we measure the uncertainty range over two key parameters: location and width. We denote the location as how far up the WTP scale the uncertainty range is, and the width as the size (or precision) of the uncertainty range.

[8] The payment card method allows respondents to state the maximum bid they would be willing to pay, this can be expanded to the *multiple-bounded* format. In the multiple bounded format the respondent answers the question “would you be willing to pay?” for each of the *k* bid amounts. If implemented with a “yes”/”no” response to each bid the method reveals the interval within which the WTP exists. Welsh and Bishop (1993) took the multiple bounded format and incorporated polychotomous responses in each of the *k* bid amounts to give the multiple bounded discrete choice method.

[9] Referred to as the “value gap” in Hanley et al. (2009).

The second objective of this chapter is to explore a new method for the econometric analysis of CV data with an uncertainty range. While the analysis of MBDC data has grown in sophistication, previous attempts at modelling such data have all been based on strong parametric assumptions; for example, Wang 1997; Alberini, et al. 2003; Evans, et al. 2003; Kobayashi, et al. 2012. We propose the use of a radically different semiparametric estimator based on the multi-state duration models used in the medical statistics literature (Commenges 2002; Frydman and Szarek 2009). The estimator we describe is a three-state duration-dependent Markov model which allows us to simultaneously explore how the width and location of the uncertainty range are influenced by different treatments. As far as we are aware this is the first time that this form of multiple-state duration modelling has been used in economic analysis.

The chapter is organised as follows. In Section 1.2 we expand on the argument that uncertainty might explain elicitation anomalies in CV data, review previous empirical evidence in this area of study, and develop the central hypothesis of the research. In Section 1.3 we review the literature on methods for analysing CV data with an uncertainty range, specifically from MBDC surveys, and justify using our semi-parametric estimator. In Section 1.4 we outline our semi-parametric estimator which is used to analyse the invariance of the uncertainty ranges to external cues. In Section 1.5 we describe the design of a CV survey experiment developed to explore the central hypothesis of uncertainty as an explanation of elicitation anomalies in CV surveys. In Section 1.6 and 1.7 we present the results of our empirical analysis and consider the implications of our findings and Section 1.8 concludes and presents some closing remarks.

1.2. Uncertainty as an explanation of elicitation anomalies in CV

It has been hypothesised that the underlying uncertainty in individuals' preferences may explain elicitation anomalies in CV studies (Ready, et al. 2001; Flachaire and Hollard 2007). In this chapter, we focus specifically on two such anomalies: the divergence in values between OE and DC elicitation methods and starting point bias[10].

[10] Here we focus on the OE-DC disparity and starting point bias but other elicitation anomalies have plausible explanations through the lens of uncertain preferences: for example, the disparity between WTP and WTA (Dubourg et al., 1994; Horowitz and McConnell 2002). Elicitation anomalies could

It is well established in the literature that DC methods of elicitation invariably report higher estimates of WTP when contrasted with OE methods of elicitation (Boyle, et al. 1996; Brown, et al. 1996; Ready, et al. 1996). A number of authors have argued that uncertainty is an explanation for that observation (Ready, et al. 1995; Welsh and Poe 1998; Ready, et al. 2001; Flachaire and Hollard 2007). At the core of those arguments is the conjecture that in the presence of uncertainty, respondents interpret OE and DC questions very differently. In particular, when faced with an OE question respondents are believed to report a value that they are reasonably certain they would pay. In contrast, when presented with a DC question offering a bid amount in their range of uncertainty, respondents are believed to react as if the question is asking them whether there is some possibility they would pay that amount. In the words of Flachaire and Hollard (2007), *"anomalies come from the fact that, when uncertain, respondents tend to answer yes. Indeed, if the bid belongs to his range of acceptable values, a respondent answers yes..."* (p. 192). Ready et al. (2001) report empirical findings that support that assertion. In their data, they observe that the norm response for respondents that are unsure is to say "yes" when answering DC questions. In contrast, when answering an OE question respondents will tend to state that they are not prepared to pay that amount.

To illustrate, imagine an individual with the uncertain preferences shown in Figure 1.1 At values below £$U^{L,}$ on the WTP scale, the individual is certain she would pay; at values between £U^L and £U^H the individual is uncertain about paying (uncertainty range), and finally, at values above £U^H the individual is certain she would not pay. When asked to state WTP in an OE valuation task, Ready et al. (2001) predict the respondents will answer towards the lower end of their uncertainty range. On the other hand, if the respondent had instead been offered a DC format question and presented with a bid amount within their uncertainty range the respondent would tend to say yes.

be explained through respondents with uncertain preferences answering questions in the former frame in a risk averse manner and questions in the latter frame in a risk seeking manner.

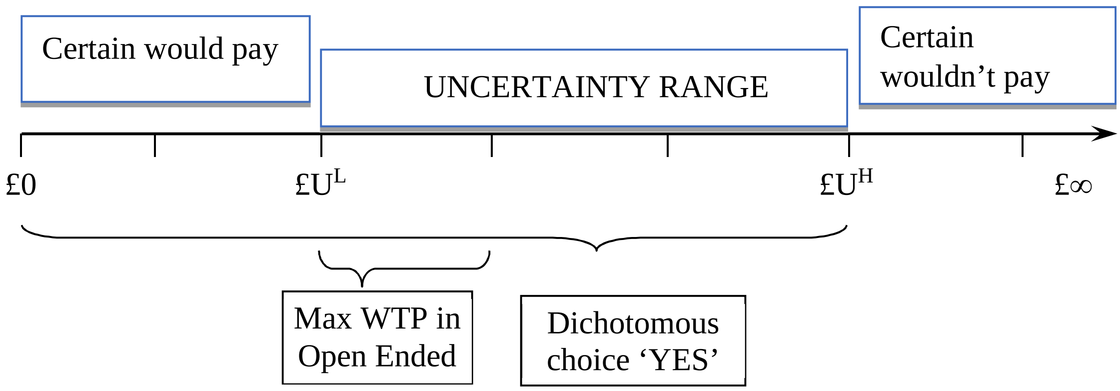

Figure 1.1. WTP Scale showing the uncertainty range and the predictions of answering different elicitation methods when the respondent is uncertain.

The central tenet of the Ready et al. (2001) conjecture is that elicitation anomalies are observed because individuals use different heuristics to deal with their uncertainty in responding to OE as compared to DC format questions. The variance in WTP from the two elicitation methods may therefore be explained simply by allowing for the possibility of uncertain preferences.

A second common elicitation anomaly is *starting point bias*; commonly observed in CV studies which ask a series of DC questions. A widely documented result is that the bid value offered in the first DC question systematically influences the response to subsequent valuation questions. While several possible explanations for starting point bias have been proposed[11], a plausible possibility is that this too arises from preference uncertainty in which respondents adopt a simplifying heuristic.

[11] A number of interpretations have been proposed, such as the initial value signalling the cost or alternatively acting as an anchor (McFadden 1994); (Herriges and Shogren 1996); Bateman et al., (2009); (Flachaire and Hollard 2007). The general explanation for anchoring is that the initial value creates the possibility, at least momentarily, that the valuation being estimated is near to the initial value. It was famously shown in Tversky and Kahneman (1974) work in which respondents’ answers on the subject of the number of African countries in the United Nations were significantly related to a number randomly generated in front of the respondents on a spinning wheel from 1 to 100. Anchors are particularly prevalent in situations when the source of the anchor is perceived as knowledgeable and trustworthy and the recipient is low in knowledge (Kahneman and Tversky 1982; Vanexel, et al. 2006). The situation described is very common in CV studies as the respondent has very little experience of the good or may be very uncertain of their WTP and could see the source of the bid as an ‘expert’.

To illustrate, imagine two individuals with identical but uncertain preferences shown in Figure 1.2. One individual is initially offered a low bid, denoted $B^L{}_1$, and one is offered a high bid denoted $B^H{}_1$. $B^L{}_1$ is comfortably within the certain would pay range and so the individual would answer *yes* they would pay. Conversely, $B^H{}_1$ is comfortably within the certain would not pay range and so that individual would answer *no*. Both individuals are then asked a second valuation question for the same value B_2 which is within their uncertainty ranges. Observe that the individual initially offered $B^L{}_1$ is coming from a state of *certainly would pay* to a state of uncertainty; whereas the individual initially offered $B^H{}_1$ is coming from a state of *certainly wouldn't pay* to uncertainty. For the individual coming up from the low bid, a natural reaction might be to reason that, 'I was previously certain I would pay and so answered *yes*. Now I am not certain that I would pay so to signal that change in state I'll answer *no*'. The reverse is true for the other individual, having previously answered with certainty that they would not pay the high bid amount a way to signal that there is now a possibility they might pay would be to answer *yes*. In other words, by adopting a simplifying heuristic to deal with their change in certainty the individuals may want to express that B_2 puts them in a different state of certainty by reversing their answers to the initial question.

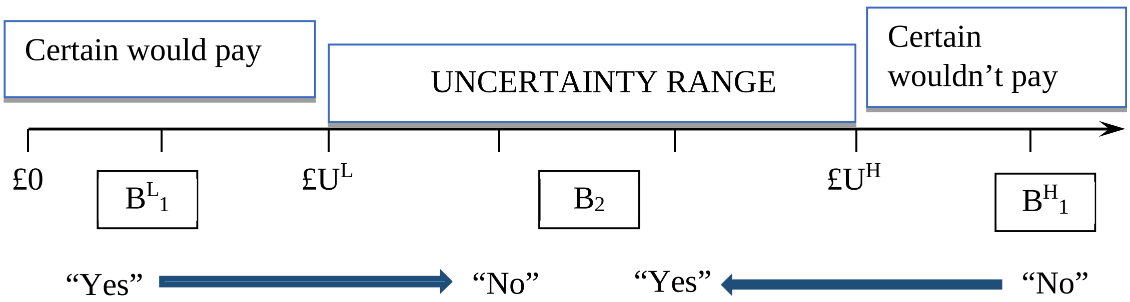

Figure 1.2. WTP scale showing the uncertainty range and the predictions of answering iterative CV questions when the respondent is uncertain.

Again the conjecture is that starting point bias is not the consequence of shifting preferences; but that respondents interpret DC questions differently when moving to a state of uncertainty from different 'directions'. Answers to the follow up question are therefore dependent upon which state they were in previously. An important prediction resulting from this hypothesis is that individuals' preferences do not change but they express the transition from a state of certainty to a state of uncertainty in their responses to multiple DC elicitation questions.

As we have shown, both the DC-OE anomaly and starting point bias might plausibly be explained through uncertain preferences. The key prediction of that explanation is that the elicitation procedure does not shift around the underlying preferences but instead the elicitation procedure might lead respondents to express those preferences differently in the presence of uncertainty. This chapter tests for such patterns by eliciting uncertainty ranges using the MBDC method and observing if value-irrelevant details of the elicitation procedure lead to variation in the uncertainty ranges. If we can show that the responses from the MBDC method are invariant to external cues in the elicitation procedure, then this would provide two particularly useful results. It would not only add support for using the MBDC method as the core method in CV survey design, but in addition, it would provide strong evidence to support the hypothesis that common CV elicitation anomalies can be explained by uncertain preferences. If the expectation of preferences that are uncertain but also procedurally invariant is supported by our data, then the prognosis for the CV method is rather encouraging; by allowing for the possibility of uncertainty, CV methods can elicit preferences that conform to many of the expectations of standard economic theory.

1.3. Modelling CV data with uncertainty ranges

Here, we review the current modelling techniques used to analyse uncertain response data in CV studies. The common assumption is that individuals do not hold fixed values for environmental goods and services, rather an individual's value, v_i, is considered to be a random variable with a continuous probability density function, $g_i(v_i)$. Figure 1.3 depicts a hypothetical individual valuation probability density function, with mean $E(v_i)$, *certainty thresholds* s_1 and s_2 and the difference between $E(v_i)$ and s_1 and s_2 denoted by a_i and b_i respectively. The certainty thresholds are defined here as the WTP value at which respondents switch their state of certainty about paying for a particular bid. Two thresholds are of particular importance for this discussion, the threshold in which respondents of CV methods switch from a state of ‘certainly would pay’ to a state of ‘uncertainty’ and the threshold from a state of ‘uncertainty’ to a state of ‘certainly would not pay’.[12]

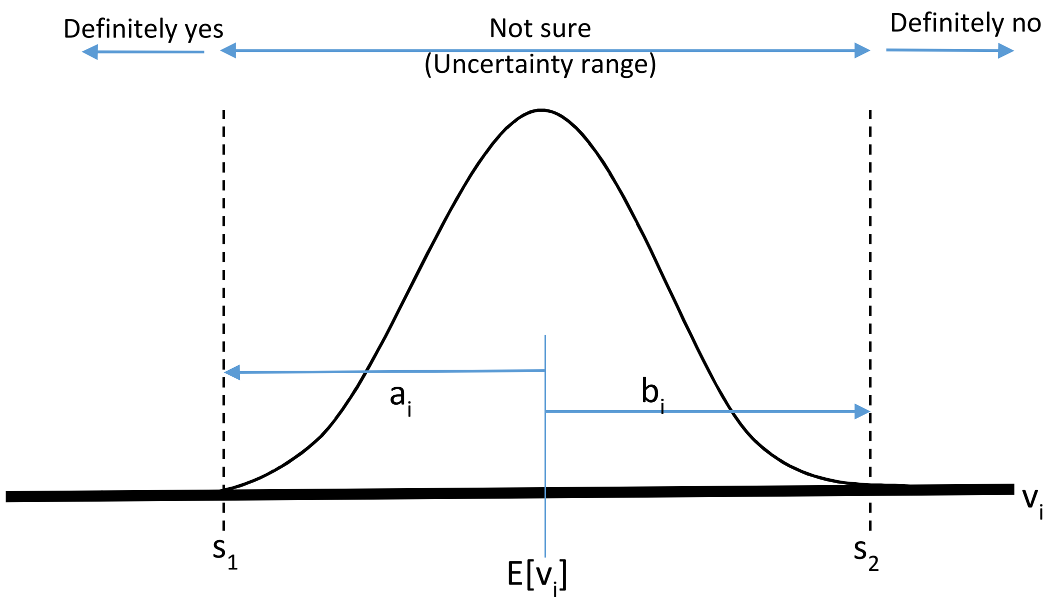

Figure 1.3. Random valuation model for an individual’s probability density function with certainty thresholds and response categories.[13]

Wang (1997) first introduced a method of modelling uncertain responses to CV surveys in his random valuation model. The random valuation model posits that a respondent answers “yes” only if their value is sufficiently large relative to the bid amount, $E(v_i) - a_i > bid$; “no” only if their value is sufficiently small relative to the bid, $E(v_i) + b_i < bid$; and “not sure” if their value lies in between, $E(v_i) - a_i < bid < E(v_i) + b_i$. For other response categories used in MBDC surveys, such as probably yes or probably no, similar boundary expressions are derived in Alberini et al. (2003) who adapt the Wang (1997) model to include five response categories.

[12] Traditionally in MBDC studies there are four thresholds, one that separates definitely yes from probably yes (PY), one from PY to not sure (NS), one from NS to probably no (PN) and one from PN to definitely would not pay; for simplicity we assume PY, NS and PN are contained within the uncertainty range.

[13] Increased knowledge or experience of the good is likely to reduce the width ($a_i + b_i$) of the individual’s uncertainty range, in addition, a change in income could shift the location of the uncertainty range (Ackerberg 2003 and Czajkowski et al. 2015).

Both Wang (1997) and Alberini et al. (2003) look to gain inference on the population distribution of $E(v_i)$ using the random valuation model. However, doing so comes at the cost of imposing restrictive assumptions on the location of $E(v_i)$ relative to the certainty thresholds, since underlying the random valuation model is an ordered Probit model. The ordered Probit model produces one less estimate coefficient than the number of parameters in the model, as such, one identification restriction is required. One such restriction imposed by Wang (1997) and some of the sub-models in Alberini et al. (2003) is symmetry around $E(v_i)$. This assumption restricts $E(v_i)$ to lie precisely between the thresholds s_1 and s_2 such that $a_i = -b_i$. An alternative restriction is to set a/b equal to a constant. Both restrictions impose the assumption that all respondents have the same relationship between the location of $E(v_i)$ and their certainty thresholds.

An alternative method of modelling uncertain responses to CV is the probability based estimator used in Evans et al. (2003). They use psychological studies to justify the mapping of categorical MBDC responses to certain survival probabilities. For example, imagine a respondent who states they "probably would pay" when presented with a bid. The probability based estimator attaches certain probabilities with various verbal probability terms, so the interpretation of the term "probable" (that the event occurs) could be 0.75; therefore, the probability that respondent $i's$, value v, lies above that bid is $Pr(v_i > bid)$ = 0.75 for any respondent. As noted by Hanley et al. (2009), the polychotomous choices in the MDBC method rely on the researcher interpreting how different respondents consider terms such as "probable" and "likely", models such as Evans et al. (2003) effectively assume that all respondents interpret these terms in the same way. Given the wide variety of unobserved forces that have the potential to affect an individual's valuation density function it seems unlikely that they would share common factors across individuals such as the same density families, let alone identical probabilities.

Finally, Kobayashi et al. (2012) propose an alternative method, the Latent Threshold Estimator. The estimator, rather than focusing on the expectation of the underlying individual valuation distribution, instead focuses on the certainty thresholds. Each threshold is modelled as a linear function and a normally distributed additive error term. The authors state that the model parameters could be estimated using maximum likelihood techniques but instead use a Bayesian approach highlighting the computational challenges of the applying maximum likelihood techniques and the small sample sizes. The Bayesian approach requires all priors to be specified and the authors use standard multivariate normal and inverse Wishart priors. The Latent Threshold Estimator models uncertain responses to CV questions and can estimate certainty threshold means and variances without requiring restrictive assumptions to the individual valuation distribution, as such, each individual valuation probability density function has its own expectation and variance.

Our semiparametric estimator, similar to Kobayashi et al. (2012), models the certainty thresholds that respondents change to different levels of certainty about paying for the good. By analysing thresholds instead of the expectation of the valuation distribution both models can assess correlation between the widths of certainty ranges and uncertainty ranges. For example, it may be the case that respondents with a narrow certainty range will also have a relatively narrow range of values over which they are uncertain. Likewise, respondents that have a large certainty range may also have a large uncertainty range. In Kobayashi et al. (2012) the certainty thresholds are modelled using normally distributed error terms; our estimator can be viewed as an alternative to the Latent Threshold Estimator with the parametric assumptions removed.

1.4. A semiparametric estimator for uncertain WTP data

A key contribution of this work is to propose a radically different semiparametric estimator based on the multi-state duration models used in the medical statistics literature (Commenges 2002; Frydman and Szarek 2009). Duration modelling deals with the statistical analysis of data recording ‘time to event’; most commonly, this is time to death or illness in medical science, and time to failure in engineering, although there are a wide variety of other uses (Klein and Moeschberger 1997).

In the medical literature the objective is to explore the progressions of illness over time and to identify durations spent in different stages of an illness and how those durations relate to each other. The estimator we describe is similar to the three-state duration-dependent Markov model developed by Frydman (1995) to analyse data on the progression of HIV/AIDS. In the context of uncertain WTP data the progression we are interested in is across money amounts and through states of certainty. As illustrated in Figure 1.4, moving up money amounts respondents transition between three states, certainly would pay, uncertain whether would pay or not, certainly would not pay.

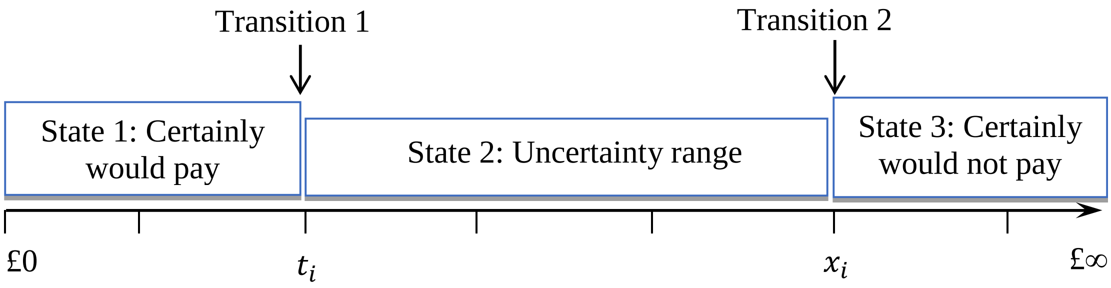

Figure 1.4. Three state duration dependent Markov process.

To construct our econometric model, we assume that each respondent, i, knows the highest amount they certainly would pay, an amount we label t_i, and the lowest amount they certainly would not pay, an amount we label x_i. The gap between these two values defines their uncertainty range, the width of which (in money amounts) we define as w_i such that $x_i = t_i + w_i$.

Accordingly, at the heart of our econometric model is a calculation of the probability of observing a respondent reporting intervals of the width t_i and w_i. We write that probability as

$$\Pr[t_i, w_i] = \Pr[t_i] \Pr[w_i \mid t_i] \qquad (1.1)$$

Observe that we allow for the possibility that the width of the uncertainty range w_i may be dependent on the width of the 'certainly would pay' interval t_i.

In the MBDC exercise respondents reveal information on their preferences over a finely spaced grid defined by the M bid points;

$$0 = b_0 \le b_1 \le b_2 \le \cdots \le b_M \le b_{M+1} = \infty$$

Accordingly, our data are discrete in nature identifying only the interval between bid points in which t_i and x_i fall. We shall refer to the bid interval between bid point b_{m-1} and b_m as B_m. We assume that bids are equally spaced along the WTP scale such that each interval B_m $(m = 1, 2, \dots, M + 1)$ is of the same width.

Now imagine that individual i indicates that they are certain they would pay each of the first n_{it} bid amounts. Subsequently, they report that they are in a state of uncertainty over the next n_{iw} bid amounts. Accordingly, for all bid amounts, b_m, where $m > n_{it} + n_{iw}$ they are certain they would not pay. For the purposes of developing our estimator, we summarize that discrete data using the following dummy variables;

- δ_{ij}^{t} $(j = 1,2, \ldots, M + 1)$ is a set of dummy variables identifying the certainly would pay range, where $\delta_{ij}^{t} = 1$ if respondent i stated that they certainly would pay b_j (such that $\delta_{ij}^{t} = 1$ for all $j = 1,2, \ldots, n_{it}$ intervals) and $\delta_{ij}^{t} = 0$ otherwise.
- d_{ij}^{t} $(j = 1,2, \ldots, M + 1)$ is a dummy variable indicating the bid interval within which t_i must fall. It is identified as the bid interval after the highest bid amount that respondent *i* indicated they certainly would pay.

The notation is a little different for the state of uncertainty. In particular, we are now concerned with the number of bid intervals over which a respondent reports a state of uncertainty, while, for the time being we ignore the fact that individuals may enter this state at different bid levels. For the purpose of clarity we use k to index the uncertainty range, where $k = 1,2, \ldots, K$ and K is the greatest number of bid intervals in the uncertainty range observed in the data. Accordingly;

- δ_{ik}^{w} $(k = 1,2, \ldots, K)$ is a set of dummy variables identifying the uncertainty range, where $\delta_{ik}^{w} = 1$ if respondent i stated that they were uncertain (such that $\delta_{ik}^{w} = 1$ for all $k = 1,2, \ldots, n_{iw}$ intervals) and $\delta_{ik}^{w} = 0$ otherwise.
- d_{ik}^{w} $(k = 1,2, \ldots, K)$ is a dummy variable indicating the bid interval within which w_i must fall. It is identified as the first bid interval before the bid amount that respondent *i* indicated they certainly would not pay.

The Model:

Our model adopts the maximally flexible parameterisation of $\Pr[t_i]$ in which a set of parameters $p_j (j = 1, 2, \ldots, M + 1)$ are estimated that capture the probability of respondent *i* having a certainly would pay range that ends in interval *j*. Accordingly;

$$\Pr[t_i] = \prod_{j} {p_j}^{d_{ij}^{t}} \tag{1.2}$$

Following Frydman (1995), we parameterise $\Pr[w_i]$, that is the probability that respondent i has an uncertainty range of width w_i, using the hazard function. In particular, we specify the hazard function using the logistic form;

$$h_k^w(t_i) = \frac{\lambda_k\, e^{\beta t_i}}{1 + \lambda_k\, e^{\beta t_i}}\ (k = 1,2,\dots,K) \tag{1.3}$$

Where $h_{ik} = h_k^w(t_i)$ represents the probability that respondent i transitions from a state of uncertainty to a state of 'certainly would not pay' after k intervals of uncertainty. Observe that the hazard is expressed with maximal flexibility through the estimation of a set of parameters λ_k $(k = 1, 2, \dots, K)$ that define the baseline hazard. At the same time, we allow for the width of the state of 'certainly would pay', t_i, to influence the hazard through the parameter β. For example, with a positive β the hazard is increasing with t_i, in other words, longer ranges of certainly would pay are associated with shorter ranges of uncertainty. Conversely, with a negative β the hazard is decreasing with t_i, in other words, longer ranges of certainly would pay are associated with longer ranges of uncertainty. From (1.3) and the earlier dummy variable definitions we derive;

$$\Pr[w_i \mid t_i] \;=\; \prod_{k=1}^{K} {h_{ik}}^{d_{ik}^w} \prod_{k=1}^{K} (1 - h_{ik})^{\delta_{ik}^w} \qquad (1.4)^{14}$$

From (1.1), (1.2) and (1.4) we obtain the loglikelihood;

$$lnL(\boldsymbol{p}, \boldsymbol{\lambda}, \beta) = \sum_{i=1}^{N} \left[\sum_j {d_{ij}}^t \ \ln p_j + \sum_{k=1}^{K} [{d_{ik}}^w \ ln(h_{ik}) \ + \ {\delta_{ik}}^w \ ln(1 - \ h_{ik})] \right] \tag{1.5}$$

Where $p = [p_1 p_2 \dots p_M], \lambda = [\lambda_1 \lambda_2 \dots \lambda_K]$

14 The hazard function is defined as the ratio of the probability density function $P(x)$ to the survival function $S(x)$, $h(x) = \frac{P(x)}{S(x)}$; therefore $P(x)$ may be obtained by multiplying the hazard function by the survival function, $P(x) = h(x) * S(x)$.

Maximising (1.5) with respect to $(\boldsymbol{p}, \boldsymbol{\lambda}, \beta)$, subject to the following constraints, $p_j \geq 0$ $(j = 1,2, \dots, M+1)$, $\sum p_j = 1$, $\lambda_k \geq 0$ $(k = 1,2, \dots, K)$, results in the following estimating equations (derived in Appendix A1);

$$p_j = \frac{n_j}{N} \quad (j = 1,2, \dots, M+1) \tag{1.6}$$

$$\sum_{i=1}^{N} d_{ir}^{w} = \sum_{i=1}^{N} h_{ir}^{w}(d_{ir}^{w} + \delta_{ir}^{w}) \quad r = 1,2, \dots, K \tag{1.7}$$

$$\sum_{i=1}^{N} t_i \sum_{k=1}^{K} d_{ik}^{w} \left(1 - h_{ik}(\lambda_k, \beta)\right) = \sum_{i=1}^{N} t_i \sum_{k=1}^{K} \delta_{ik}^{w} \, h_{ik}(\lambda_k, \beta) \tag{1.8}$$

Where n_j is the number of respondents with t_i in the interval j and N is the total sample size. Notice that in this specification we have a closed form solution for $p_j(j = 1,2, \dots, M+1)$, but not for the λ_k's and β. To estimate those parameters we use the self-consistency algorithm suggested by Frydman (1995). The algorithm steps are as follows,

0. Calculate p_j from (1.6)
1. Choose initial values λ_k $(k = 1,2, \dots, K)$ and β, which we denote λ_k^0 $(k = 1,2, \dots, K)$ and β^0, where the superscript 0 indicates the initial iteration of the algorithm.
2. Calculate new values for λ_k $(k = 1,2, \dots, K)$, which we denote λ_k^s, where s indexes the iteration of the algorithm such that in the first iteration $s = 1$. From (1.7) we obtain the estimating equation (derived in appendix A1);

$$\lambda_k^s = \frac{\lambda_k^{s-1} \;\; n_r^w}{\sum_{i=1}^{N} \left[\; h_{ik} \; (d_{ik}{}^{w} + \delta_{ik}{}^{w}) \; \right]} \quad (k = 1,2, \dots, K) \tag{1.9}$$

Where $n_r{}^w$ is the number of people who fail in a particular interval r (such that, $n_r{}^w = \sum_{i=1}^{N} d_{ir}{}^w$ $(r = 1,2, \dots, K)$).

3. Calculate new value for β^s from (1.8). Accordingly, one has to solve a non-linear equation. To do so we apply the Newton-Raphson method

where convergence is achieved when the change in β^s falls below a certain threshold, ε.

4. Stop when $|\lambda_k^s - \lambda_k^{s-1}| < \varepsilon$ and $|\beta^s - \beta^{s-1}| < \varepsilon$, otherwise return to step 2 and iterate.

As demonstrated by Frydman (1995) the self-consistent algorithm returns maximum likelihood estimates of the parameters of the model.

1.5. Experimental design

Survey respondents in our experiment each faced a valuation exercise made up of three tasks, see Figure 1.5 for the progression of the tasks. In Task 1 respondents were randomly allocated by an unseen process into one of eight treatment groups, seven of the eight groups received a single bounded DC question at a specific bid level and the other group received an OE question. The DC bid amounts were chosen according to two criteria: that they represented reasonable values suggested by prior focus group testing, and that they produced results which could be unambiguously tested against our hypotheses. Accordingly, five DC bid levels of £5, £30, £60, £100 and £150 provide a range of bid levels which vary in terms of the absolute value; in addition, the original survey included two bid levels that were designed to test if respondents answer CV questions differently because the bid levels are not round numbers. The effect of this spurious accuracy in the bid levels on CV responses is not explored in this chapter but the two additional treatment groups of £28.70 and £31.30 are included in the subsequent analyses where appropriate.

Three-step valuation exercise

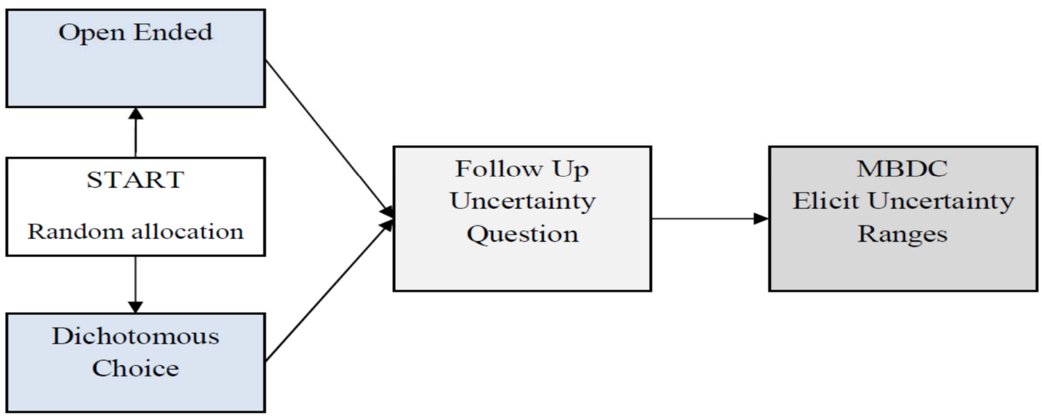

Figure 1.5. The valuation exercise.

Task 2 and Task 3 were completed by all respondents regardless of their treatment group. Following the procedure of Li and Mattsson (1995) and Ready et al. (2001), Task 2 presented respondents with a follow-up question that required them to state the level of certainty they attached to their DC or OE answer from task 1. Five responses were available:

- I definitely *would* pay the amount of money.
- I probably *would* pay the amount of money.
- I am *not sure* if I would pay this amount of money.
- I probably *would not* pay the amount of money.
- I definitely *would not* pay the amount of money.

Task 3 uses a novel version of the MBDC method[15] to establish the values over which respondents are certain and uncertain. The standard format of an MBDC question is to

---

[15] Our MBDC task was undertaken as follows:

If definite in task 2: Just now you said that you would/(would not) pay £ … for the enlarged beaches. I'll indicate your answer by placing a tick in the "Definitely Yes"/("Definitely No") box next to that amount.

Now consider the higher/(lower) amounts on the list. (Pass list and clipboard to respondent). Starting with £ … (next highest/(lowest) amount), work down/(up) the list considering each of these amounts in turn until you reach an amount that there's a possibility you would not/(would) pay, however small. Again, looking at the card decide which category best describes your response to that amount and tick the corresponding box on the list.

Continue working down/(up) the amounts on the list, ticking one box for each amount. Stop once you reach an amount that you definitely would not/(would) pay.

If Not Sure/probably in task 2: Just now you said that you "probably would"/"uncertain whether would"/"would not" pay £ … for the enlarged beaches. I'll indicate your answer by placing a tick in the "probably yes"/"not sure"/"probably no" box next to that amount.

Now consider the next amount down on the list. Still looking at the card, if the amount was (£next highest amount) which of the categories on the card best describes your response to that amount (tick in appropriate box next to amount on valuation sheet).

Work down the amounts on the list, ticking one box for each amount. Stop once you reach an amount that you Definitely Would Not Pay.

Now I'd like you work up the amounts on the list. Starting at (£next lowest amount) tick one box for each amount and stop once you reach an amount that you Definitely Would Pay.

have five to ten bids in which the respondent states their certainty to paying these bounds using the standard polychotomous choice options. In our application those handful of bounds were replaced by a semi-continuous range of bids. The values on the MBDC card ranged from £1 to £500, increasing in £1 increments. The large number of bids was presented over two pages which were shown to the respondents in advance.[16, 17]

It has been shown that the range of bids on a payment card can systematically influence responses to MBDC questions (Dubourg, et al. 1997; Roach, et al. 2002). Accordingly, we held the range of bids constant for all respondents. Moreover, the traditional format for MBDC questions is to space bids along a logarithmic scale, for example, see Ready et al. (2001) and Welsh and Poe (1998). As a result, more precise information is provided on the location of low WTP amounts than of high. In contrast, our design using £1 increments across the whole range of bids ensures high precision regarding the uncertainty range at all levels of WTP.

The MBDC data for each respondent reveals their maximum *definitely would pay* amount, the value at which they transition to *probably would pay*, *not sure*, and *probably would not* pay, and also a minimum value for their *definitely would not pay* amount. We label the uncertainty range as the range between the lower bound (the maximum definitely would pay amount) and the upper bound (the minimum would not pay amount). This method provides the exact size (within £1) of respondents' uncertainty ranges. Accordingly, we can test these uncertainty ranges for movement in both location and width over the whole range of bids in the MBDC design.

In the context of a carefully designed split-sample experiment we aim to test the hypothesis that it is the process of requiring individuals to express values in CV surveys as if they had well-defined certain preferences that leads to elicitation anomalies. To test this we elicit uncertain CV responses and test for invariance to the nature of the elicitation procedure using three key tests: (i) is the location of the uncertainty range invariant to the absolute value of the bid of a prior DC question; (ii) is the width of the uncertainty range invariant to the bid of a prior DC question; (iii)

[16] The survey is included in Appendix A3.

[17] The full MBDC payment card in shown in Appendix A4.

are the location and the width of the uncertainty range invariant to whether the respondent received a prior OE or DC question?

1.6. Implementation

Our specific case study concerns potential improvements in coastal protection (extending the size of the beach through the installation of more groynes) in the town of Southwold in Suffolk, UK. The data was originally collected for use as part of dissertation projects at the University of East Anglia in 2004, the data has not previously been published in any peer reviewed source. Personal interviews were conducted by four interviewers at three locations close to areas that would receive the additional coastal protection if the project were to go ahead. The proposal was described by the interviewer who also presented respondents with maps and visual representations of the site before and after the construction of additional coastal protection. Survey respondents were informed that the existing defences would be maintained by government funding but that additional improvements would require funding through an increase in general taxation.

Respondents were first asked questions regarding the frequency with which they visited the beaches, as well as their reasons for visiting and how far away they lived. Subsequently, they were presented with the information on the coastal-protection proposal before proceeding to complete Task 1 (answering either an OE or DC CV question) and Task 2 (the uncertainty follow-up question) of the value-elicitation procedure. Task 3 (MBDC elicitation) began by introducing respondents to the MDBC card listing the bid levels from £1 to £500 and the certainty scale associated with each. The interviewer then translated a respondent's answers from Task 1 and Task 2 onto the MBDC card. For example, if the respondent answered that they were probably sure (Task 2) they would pay £10 (Task 1) then the interviewer ticked that particular box on the MBDC card. Respondents were asked to proceed from that point in completing the MBDC card. If they were sure they would pay that initial amount then they were asked to work up the card marking their certainty against each amount until they reached amount they were certain they would not pay. The reverse was true if they were certain they would not pay the initial bid amount. If a respondent was unsure as to whether they would pay the initial bid amount then they were first asked to work up the card to an amount they were certain they would not pay and then down the card to identify the highest amount they were certain they would pay. The final part of the survey elicited socio-economic details.

Individuals in the study were randomly allocated to either an OE group or one of seven DC treatment groups. The total sample was 952 respondents, of that 36 are classified as unusable for the subsequent analysis undertaken in this chapter. The exclusion of observations is mainly due to incomplete MBDC tasks. For example, a number of respondents only stated a single ‘not sure’ figure and no values for any of the other polychotomous choice options. In addition to this, 4 respondents ticked that they were certain they would pay all the way up to £500 (the upper limit of the payment card). This data, although possibly very important for total WTP estimates in standard CV analyses, fails to provide any information about the location or width of the range of values to which the respondent is uncertain and so is ignored for the purposes of this study. As such, the total usable sample was 916 respondents, with the OE group containing 272 respondents and the DC groups each containing between 85 and 95 respondents.

Table 1.1 provides summary details of the socioeconomic composition of each treatment group. As can be seen from the final column, no significant differences are observable across socioeconomic characteristics in the eight treatments, suggesting that the randomisation to treatment groups was largely successful.

Variable	**Sample means (standard deviations for continuous variables)**								**Test of difference in groups (*p*-value)**
	Open ended	**£5**	**£28.7**	**£30**	**£31.3**	**£60**	**£100**	**£150**	
Gender (1=Male)	0.60	0.63	0.55	0.52	0.59	0.52	0.66	0.47	**0.108**[a]
Employed % (1 =yes)	0.57	0.63	0.67	0.54	0.52	0.65	0.54	0.66	**0.160**[a]
Nature/green group member (1=yes)	0.42	0.34	0.44	0.37	0.43	0.39	0.32	0.41	**0.592**[a]
Number of people in household	2.82 (1.3)	2.64 (1.3)	2.79 (1.2)	2.69 (1.3)	2.51 (1.1)	2.91 (1.3)	2.72 (1.3)	2.82 (1.2)	**0.446**[b]
Age (Years)	52.96 (14.5)	51.44 (14.5)	52.06 (12.3)	55.48 (14.3)	54.03 (14.5)	53.77 (12.9)	53.52 (15.0)	53.33 (13.0)	**0.648**[b]

Income per household (£/month)	2856 (1353)	2626 (1266)	2845 (1372)	2690 (1261)	2892 (1333)	3122 (1350)	2729 (1253)	3122 (1369)	**0.209[b]**
Sample size	272	91	85	95	91	93	95	94	Total = 916

Table 1.1. Comparison of descriptive statistics across samples.

[a] *p*-value calculated from χ^2 test of equality of proportions across multiple groups.

[b] *p*-value calculated from ANOVA F-test of equality of means across multiple groups.

1.7. Results

Elicitation effects

A number of studies have reported that DC questions generate larger WTP estimates than OE questions, for a review see Brown et al. (1996). Figure 1.6 demonstrates a similar pattern in our results. OE responses are summarised through the survivor function of that data (calculated using the Kaplan-Meier estimator) which plots the percentage of respondents whose initial WTP (from Task 1) is greater than or equal to each bid level used in the DC treatments. Equivalent data for each DC treatment group is plotted on the graph; in this case, each point illustrates the percentage of respondents in a group stating they would pay the DC bid level. What is immediately evident from Figure 1.6 is the fact that the implied distribution of values from the DC treatment groups greatly exceeds that volunteered by the OE treatment group. This observation is confirmed statistically through a series of two-sample proportional tests using the Fisher-Exact method (*p*-values 0.000 to 0.016) with the results reported in table 1.2.

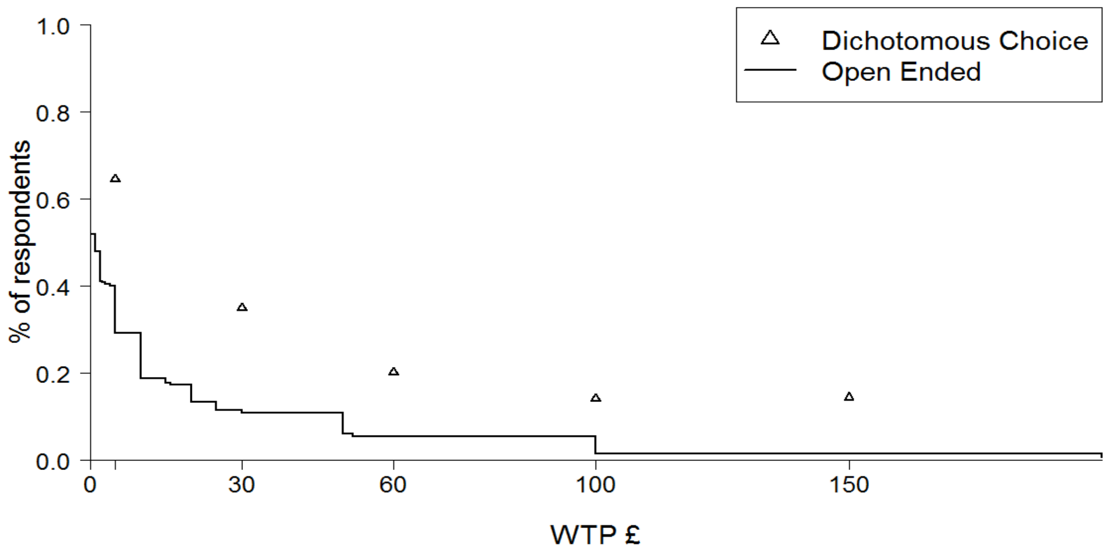

Figure 1.6. Empirical survivor function of OE treatment group and acceptance rate for the DC treatment groups at discrete bid levels.

Bid level	Original			Recoded OE and Recoded DC		
	OE (higher:lower)	DC (yes:no)	Fisher Exact test	CS OE (higher:lower)	CS DC (yes:no)	Fisher Exact test
5	106:166	58:33	0.000	99:173	36:55	0.617
28.7	29:243	47:38	0.000	23:249	31:54	0.000
30	29:243	32:63	0.000	23:249	15:80	0.051
31.3	27:245	41:50	0.000	21:251	18:73	0.003
60	13:259	18:75	0.000	9:263	5:88	0.360
100	13:259	12:83	0.016	8:264	5:90	0.334
150	2:270	11:83	0.000	3:269	6:88	0.011

Table 1.2. Open Ended, Dichotomous Choice and Certainty Standardised (CS) results with two sample proportional tests

Convergence of DC and OE after recoding to the same certainty level

Now consider the hypothesis of Ready et al. (2001) who conjecture that respondents process OE and DC questions differently in the face of uncertainty; in a DC setting a respondent may state that they are willing to pay a bid amount lying in their uncertainty range, but submit a WTP value from the bottom of that range in response to an OE question. We test that hypothesis by recoding the DC and OE responses in Task 1. Following Ready et al. (2001), if the respondent subsequently expressed a state of certainty less than “definitely would pay” in Task 2, then “yes” answers to Task 1 DC questions are recoded to “no”. Ready et al. (2001) contrast their DC treatment with a payment card elicitation method; they recode to the “definitely would pay” level of certainty by requiring those respondents who state a lower level of certainty to point to a number on the payment card that they definitely would pay. In our survey, we did not require our OE respondents to express a new open ended value, instead we gathered equivalent information through the MBDC method in Task 3. As such, for those respondents in the OE treatment group who stated a certainty level of less than “definitely would pay” we recoded their value to the highest amount they indicated in the MBDC method that they “definitely would pay”.

As illustrated in Figure 1.7 and reported in table 1.2, once responses from the DC and OE treatment groups have been standardised to the ‘definitely would pay’ level of certainty a significant gap still remains between acceptance rates for DC questions and the survivor function for OE responses. A series of pointwise comparisons using Fisher-exact proportion tests reveal statistically significant differences between the two certainty-standardised data sets at WTP amounts of £28.70, £31.30 and £150 with marginally significant differences at £30. Contrary to the findings of Ready et al. (2001) our data suggest that respondents to DC questions continue to indicate higher levels of WTP than respondents to OE even once the levels of certainty in responses has been standardised.

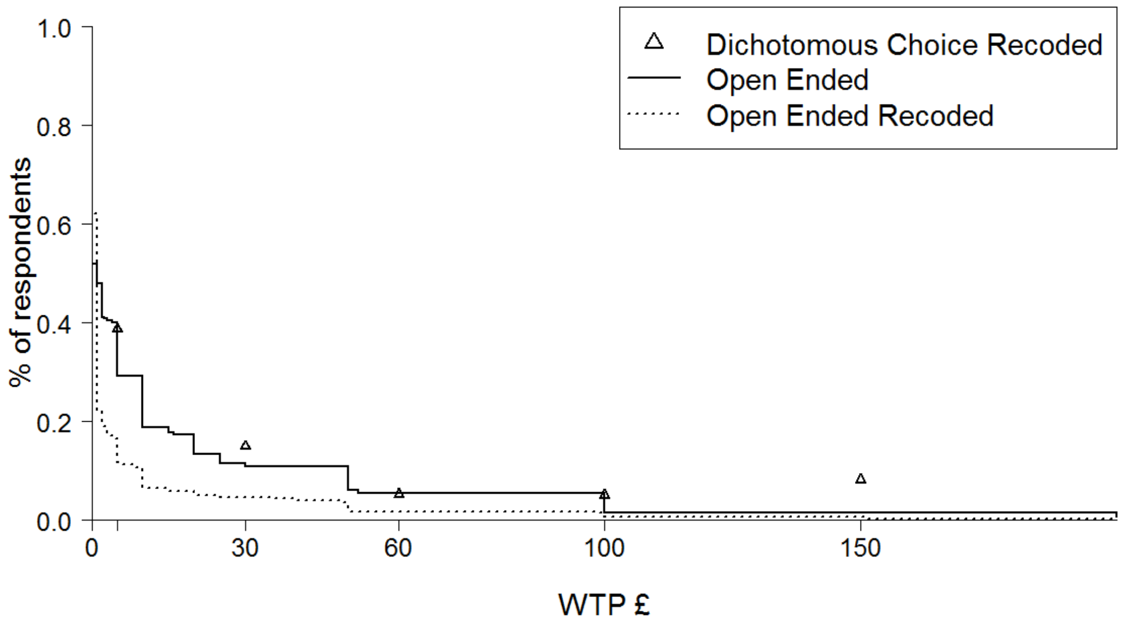

Figure 1.7. Empirical Survivor function OE and DC responses after recoding to “definitely would pay” level of certainty.

A second pattern of response that stems from the Ready et al. (2001) hypothesis is that respondents presented with a DC bid amount lying within their uncertainty range are more likely to respond “yes”. In their empirical application Ready et al. (2001) document 11 respondents that classed their level of certainty with respect to their response to a DC question as being “not sure”. Of those 11, nine respondents (82%) had opted to answer “yes” to the DC question, a result that Ready et al. (2001) claim supports their hypothesis. However, data from our experiment displayed in Table 1.3 presents a contradictory result; only 12% of the 25 respondents that classed their level of certainty in answering a DC question as being “Not Sure” opted to answer “yes” to that question. One important difference may be the use of language in the follow-up polychotomous choice certainty question. Our method used the phrase "definitely sure" and "probably sure" whereas (Ready et al. 2001) used "95% sure" and "more likely", importantly though, both studies have an "unsure" group which should produce consistent results.

Follow up certainty response	Open-ended	DC "yes"	DC "no"	DC % "yes" [a]
Definitely would pay	208 (75%)	119	0	100%
Probably would pay	63 (23%)	116	0	100%
Not Sure	4 (1%)	**3**	**22**	**12%**
Probably would not pay	0 (0%)	1	29	3.3%
Definitely would not pay	2 (1%)	0	385	0%

Table 1.3. Follow-up certainty responses for DC and OE groups.[18]

[a] calculated from 'DC yes'/('DC yes' + 'DC No')

Overall, our data replicate the standard finding from the CV literature in which DC format questions elicit responses implying higher WTP than responses to OE format questions. Our data provide little support for the Ready et al. (2001) hypothesis in that responses remain significantly different even when answers under the two formats were compared at the same level of respondent certainty. In addition, our data contradict the finding of Ready et al. (2001) in that we find no propensity for respondents to answer "yes" when a DC question falls in their uncertainty range. Indeed, our data suggest the opposite tendency with a large majority of respondents in those circumstances opting to answer "no".

Multiple Bounded Discrete Choice – Uncertainty Ranges

[18] All 952 respondents are used for this analysis as all respondents completed the first two tasks. Removing the 36 respondents with incomplete data in the third (MBDC) task results in 18.75% of "not sure" respondents answering "yes" to the DC question in the first task (3 out of 19).

If the patterns of response identified from standard CV elicitation can be explained through the existence of underlying uncertain preferences then the central question becomes whether those uncertain preferences are themselves influenced by the elicitation procedures. We examine responses to the MBDC elicitation from Task 3 of our valuation experiment to explore whether the uncertainty ranges identified in that task are invariant to the nature of the standard CV elicitation question presented to them in Task 1.

Figure 1.8 summarises our data from the MBDC exercise in graphical form; each bar shows the average uncertainty range for each treatment group. The OE group is represented by the lower most horizontal bar with the other seven DC treatment groups above. Visually, the uncertainty range for the £5 DC treatment group is lower to that from the OE treatment group and, for successively larger DC bid level treatment groups, those uncertainty ranges shift up the WTP scale and span a seemingly wider range.

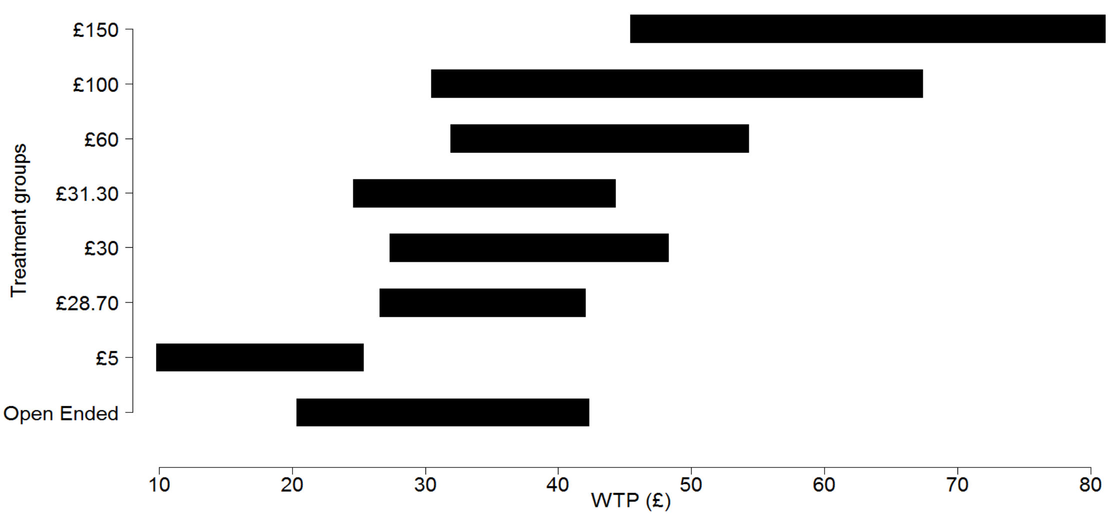

Figure 1.8. Mean Uncertainty ranges of DC and OE treatment groups.

To explore those patterns more formally, we statistically compare the lower and upper bounds and the widths of the of uncertainty range across treatment groups. More specifically, we compare the means of the highest value that respondents in each treatment group ‘certainly would pay’, the lowest value respondents ‘certainly would not pay’ and the mean width of the uncertainty range. This is summarised in Table 1.4.

Mean uncertainty ranges from MBDC (95% confidence intervals) Only positive WTP [19]									Difference in means of multiple groups (*p*-values)
Treatment groups	**Open ended**	**£5**	**£28.70**	**£30**	**£31.30**	**£60**	**£100**	**£150**	
Highest 'certainly would pay MBDC (£)	20.3 (14.8-25.8)	9.8 (5.8-13.8)	26.6 (21.6-31.6)	27.3 (17.4-37.3)	24.6 (19.5-29.7)	31.9 (24.8-39.0)	30.5 (19.5-41.4)	45.4 (30.7-60.1)	0.000[a]
Lowest 'certainly would not pay' MBDC (£)	42.3 (33.9-50.6)	25.3 (17.2-33.4)	42.0 (36.4-47.6)	48.3 (34.5-62.1)	44.3 (38.9-49.6)	54.3 (38.3-70.3)	67.4 (49.0-85.7)	81.1 (57.0-105.2)	0.000[a]
Width of uncertainty range (£)	22.0 (16.9-27.1)	15.5 (10.4-20.6)	15.4 (12.1-18.7)	21.0 (14.2-27.7)	19.7 (14.4-25.0)	22.4 (9.4-35.3)	36.9 (23.2 50.7)	35.7 (23.3-48.0)	0.001[a]
Obs	**143**	**59**	**54**	**48**	**48**	**36**	**42**	**48**	

Table 1.4. MBPC responses for OE and DC treatments.

[a] *p*-value calculated from ANOVA F-test of equality of means across multiple groups

The final column in Table 1.4 contains the *p*-values of an F-test for equality of means between the multiple treatment groups. We see significant difference at the 99.9% confidence level showing that at least one of the treatment groups has a different mean to another treatment group. In addition, we also test for difference in the width of the uncertainty range and observe significant difference between at least one of the treatment groups when compared to the other treatment groups (*p*-value 0.001).

To better understand the patterns of difference in the location and width of the uncertainty ranges of different treatment groups we employ the semi-parametric estimator described in Section 1.4. Our strategy is to parameterise the two durations in our model as functions of treatment group. More specifically we define a set of dummy variables q_0 to q_7 with q_0 defining the OE treatment, and q_1 to q_7 defining the seven

[19] Figures in Table 1.4 are based on only those respondents who stated a positive WTP, this data therefore excludes all respondents who stated they would not be willing to pay anything for the project. The analysis in Table 1.4 using all responses is qualitatively identical to those described here.

different initial bids in the DC treatments (such that, q_1 = £5 DC treatment and q_7 = £150 DC treatment) and use those to parameterise the two 'WTP state durations' in our model; duration in a state of 'certainly would pay' (t) and duration in a state of 'uncertainty'(w). Specifically, we use those treatment group dummy variables to parameterise the probability of transitioning from a state of 'certainly would pay' to a state of 'uncertainty' as defined by the hazard function h_j^t $(j = 1, 2, \dots, M+1)$ and the probability of transitioning from a state of 'uncertainty' to a state of 'certainly would not pay' as defined by the hazard function h_k^w $(k = 1, 2, \dots, K)$. The latter parameterisation is a straightforward extension of equation (1.3):

$$h_k^w(t_i, \boldsymbol{q}_i) = \frac{\lambda_k\ e^{\beta t_i + \beta_0 q_0 + \cdots + \beta_7 q_7}}{1 +\ \lambda_k\ e^{\beta t_i + \beta_0 q_0 + \cdots + \beta_7 q_7}} \quad (k = 1,2,\dots,K) \tag{1.10}$$

where, as before, λ_k $(k = 1,2,\dots,K)$ define the baseline hazard, t_i is the maximum amount that respondent i was certain they would pay, and $\boldsymbol{q}_i = [q_{0i}\ \ q_{1i}\ \dots\ q_{7i}]$ is a vector of dummy variables identifying respondent i's treatment group and $\boldsymbol{\beta} = [\beta\ \ \beta_0\ \ \beta_1\ \dots\ \beta_7]$ is a vector of parameters to be estimated.

In a similar vein, we define the probability of transitioning out of a state of certainly would pay as:

$$h_j^t(\boldsymbol{q}_i) = \frac{\phi_j\ e^{\alpha_0 q_{0i} + \dots + \alpha_7 q_{7i}}}{1 +\ \phi_j\ e^{\alpha_0 q_{0i} + \dots + \alpha_7 q_{7i}}} \quad (j = 1,2,\dots,M+1) \tag{1.11}$$

Where ϕ_j $(j = 1,2,\dots,M+1)$ define the maximally flexible baseline hazard for that transition, $\boldsymbol{q}_i$ identify treatment group for respondent i and $\boldsymbol{\alpha} = [\alpha_0\ \ \alpha_1\ \dots\ \alpha_7]$ are parameters to be estimated. As per equation (1.2), the probability of observing a particular maximum 'certainly would pay' quantity, t_i, can be calculated from the hazard function as follows;

$$\Pr[t_i|\boldsymbol{q_i}] = \prod_j h_j^t(\boldsymbol{q_i})^{{d_{ij}}^t} \tag{1.12}$$

For the purposes of identification, we set $\beta_0 = 0$ and $\alpha_0 = 0$ such that OE elicitation forms our comparator treatment group.

Test 1: Is the location of the uncertainty range invariant to the absolute value of the initial bid amount?

Table 1.5 reports parameters of the model estimated using the self-consistency algorithm described in Section 1.4. The first two columns of Table 1.5 report the parameters associated with the probability of transitioning from a state of 'certainly would pay' to one of 'uncertainty'. Observe that the probability of transition for each of the DC treatment groups is statistically significantly different from that of OE treatment group at greater than the 95% confidence level. In the case of the £5 DC treatment group the parameter exhibits a positive sign indicating that individuals offered an initial bid amount of £5 had significantly higher transition hazards than OE respondents and therefore had a higher probability of exiting the 'certainly would pay' interval at lower WTP amounts. The other treatment group parameters all have a negative sign indicating that individuals offered an initial bid amount of £28.70 or more had significantly lower transition hazards than OE respondents. In other words, respondents offered £28.70 or more as an initial bid level had a higher probability of exiting the 'certainly would pay' interval at higher amounts relative to the OE respondents.

Transition Hazard parameters	Certainty to uncertainty transition – with parameters for initial bid effect $h_j^t(\boldsymbol{q}_i)$	Transition Hazard parameters	Uncertainty to certainly would not pay transition – with parameters for initial bid effect $h_k^w(t_i, \boldsymbol{q}_i)$
		Entry point certainly would not pay β	-0.0088*** (0.0018)
Open ended α	Base case	Open ended β	Base case
£5 α	0.6248*** (0.1699)	£5 β	0.2616 (0.1581)
£28.70 α	-0.3328* (0.1618)	£28.70 β	0.2739 (0.1613)
£30 α	-0.3730* (0.1733)	£30 β	-0.0467 (0.1691)
£31.30 α	-0.3847* (0.1693)	£31.30 β	-0.0272 (0.1675)

£60 α	-0.6052** (0.1905)	£60 β	0.0717 (0.1954)
£100 α	-0.4757** (0.1836)	£100 β	-0.2851 (0.1816)
£150 α	-0.8063*** (0.1721)	£150 β	-0.1258 (0.1836)
Log likelihood	-1388		-1671
Obs	473		473

$\lambda's\ and\ \phi's$ are not reported.

Significance levels: *0.05, **0.01, ***0.001

Table 1.5. Semi parametric estimator output.

Our findings provide evidence that the maximum amount respondents indicate they are ‘certain they would be willing to pay’ in a MBDC exercise is not invariant to the nature of a preceding elicitation question. Relative to the OE sample, the start of the uncertainty range significantly shifted down for respondents previously asked about paying a low (£5) initial bid amount, and the start of the uncertainty range significantly shifted up for respondents previously offered a high (£28.70 or more) initial bid amount. These results are consistent with an anchoring effect in that the initial bid offered in the DC task significantly influenced the uncertainty ranges elicited in the MBDC task.

To visualise the data we select a subset of our data. We include only those subjects who in task 1 were initially offered a standard DC bid amount of either £5 or £150. We are therefore left with a total of 185 observations with 91 observations in the £5 group and 94 observations in the £150 group. We summarise the results graphically in Figure 1.9 using a survivor function representation of the WTP state distributions. In Figure 1.9 we illustrate the distribution for the ‘certain would pay’ state contrasting the £5 treatment group on the left with that of the £150 treatment group on the right. Clearly, those offered the £150 initial bid amount in the DC task were more likely to express certainty over being willing to pay higher bid amounts in the MBDC task.

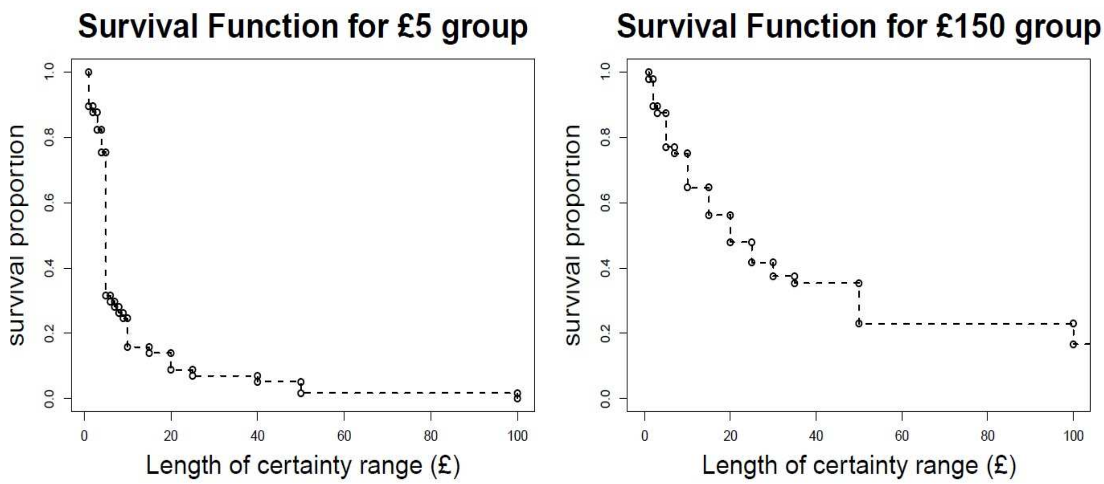

Figure 1.9. Survival functions for the certainty range, including only those individuals who gave a positive WTP.[20]

Test 2: Is the width of the uncertainty range invariant to the initial bid amount?

The parameter estimates reported in the final two columns of Table 1.5 are those associated with the width of a respondent's 'uncertainty' range. The β parameter allows for the width of that range to depend on the level of WTP at which a respondent entered a state of uncertainty. Since the model is parameterised in terms of the hazard function, the highly significant and negative β reveals that the higher up the WTP scale the individual enters uncertainty the smaller their transition probability is for exiting uncertainty. In other words, respondents who state higher certain WTP amounts in the MBDC exercise also exhibit wider uncertainty ranges.

Now consider the parameters estimated on the treatment group dummy variables for the transition from 'uncertainty' to 'certainty would not pay'. The results show clear evidence that, having controlled for the WTP-level effect captured by β, there are no statistically significant treatment group effects. This result is consistent across all treatment groups including those offered the precise bid amounts of £28.70 and £31.30. Accordingly, for our sample, the width of the uncertain ranges elicited using the MBDC method are invariant to the nature of the Task 1 valuation exercise.

[20] A proportions test shows no statistical difference between treatment groups for the number of respondents who were not willing to pay anything for the environmental change.
Number of zero WTP respondents = 34 out of 91 for £5 treatment group and 46 out of 94 for £150 treatment group, P-value = 0.112.

Figure 1.10 compares the distribution of widths of uncertainty ranges across the same two treatment groups as Figure 1.9. As we have seen our model shows a significant WTP-level effect; that is to say, the width of the uncertainty range depends on the level of certain WTP. Accordingly, the distributions in Figure 1.10 are presented using the mean maximum certain WTP amounts (t) for those two treatment groups; £9.80 for the £5 treatment group and £45.40 for the £150 group. Again we observe a difference in the distributions with that for the £150 group identifying a substantially increased density of respondents with large uncertainty ranges.

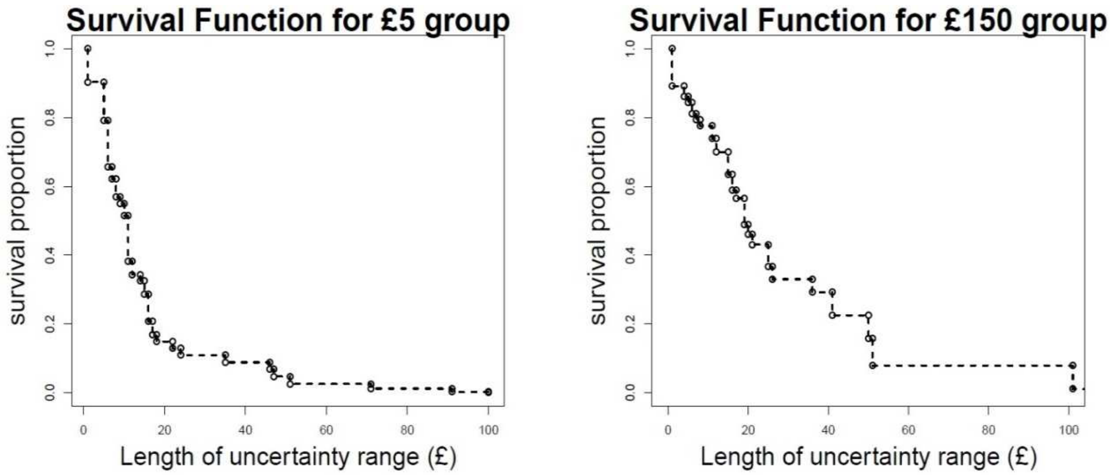

Figure 1.10. Survival functions for the uncertainty range with t_i equal to the mean t_i for each treatment, t = £9.80 for £5 and t = £45.40 for £150.

Finally, Figure 1.11 compares the distribution of the width of uncertainty ranges while holding the WTP-level effect constant. In particular, we plot the distributions assuming an identical maximum certain WTP amount (t) of £10. Observe that the two distributions now appear very similar suggesting that treatment group has little independent effect on the uncertainty range width. In summary, the graphs visually confirm the results that the location of the uncertainty range can be shifted by the initial bid amount but the width of the uncertainty range is not independently affected.

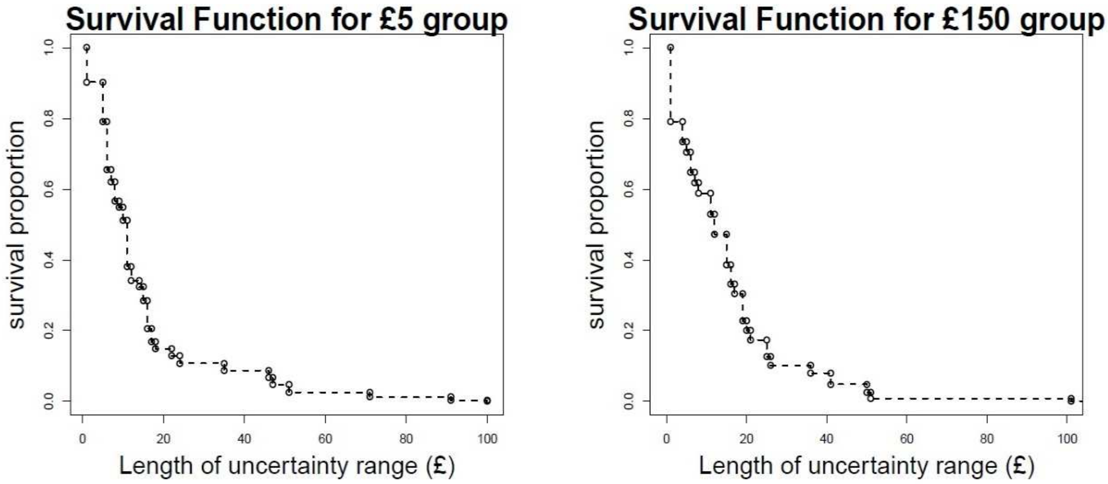

Figure 1.11. Survival functions for the uncertainty range using an identical entry bid value t = £10, including only those individuals who gave a positive WTP.

Test 3: Are the location and the width of the uncertainty range invariant to the OE treatment versus the DC treatment?

Finally, we can consider the general differences between the DC sample as a whole and the OE sample. Using the results presented in Table 1.5 it is clear that the important characteristic is the absolute value of the initial DC bid amount. The absolute value of the initial bid has an anchoring effect on the respondents in such a way that it carries through to a subsequent MBDC task and is expressed in those elicited uncertainty ranges; however, this result is only seen in the location of the uncertainty range and not the width of the uncertainty range. We see that the differences in the width of the uncertainty ranges are captured by the effect from the different width of the 'certainly would pay' interval.

Overall, we conclude that it is the absolute value of the bid level from a previous DC question that causes respondents to express different uncertain preferences to those previously offered an OE question. Our results show, visually in Figure 1.8 and numerically in table 1.4, that OE responses were closest to the £28.70, £30 and £31.30 DC bid levels. In addition, those previously offered a £5 bid level expressed a downwards shift in their uncertain preferences and those offered £60 or higher expressed an upwards shift in their uncertain preferences compared to the OE respondents. Therefore, for our sample, the uncertainty ranges elicited using the MBDC can be shifted up or down by the particular nature of a prior valuation question.

1.8. Conclusion and discussion[21]

The main focus of this chapter is to evaluate the hypothesis that individuals hold preferences that are uncertain but in all other respects comply with the standard assumptions of economic theory. Results from our empirical experiment comparing responses to a standard single bounded DC and standard OE questions confirmed findings in the existing literature on CV anomalies; in particular, the DC treatment implies higher WTP than the OE treatment. Our data provide little support for the Ready et al. (2001) hypothesis that those differences can be explained through differences in the certainty with which respondents answer these two different elicitation methods. When respondents' answers were recoded to a comparable level of certainty, significant differences in WTP were still observable in our data. In addition, our data contradicts one of the key findings of Ready et al., (2001) as we found that the most frequent response for those individuals who stated they were "not sure" to their initial answer was to say "no" in the DC format.

Uniquely, our experiment investigated the uncertainty of preferences by following-up on an initial value-elicitation task with a MBDC exercise. Our null hypothesis being that uncertain preferences should be invariable in response to value-irrelevant details of that initial elicitation task. Our data show clear evidence that the location (though not width) of the uncertainty range is significantly influenced by the elicitation procedures. Respondents updated the location of the start of their uncertainty range to higher amounts when offered a high DC bid amount (£28.70 to £150) and to lower amounts when offered a low DC bid (£5) relative to the OE method. We conclude that, for our sample, respondents to CV surveys readily express uncertain preferences and their responses can be shifted around by the elicitation procedure. Specifically, the results are consistent with an anchoring effect in that the initial bid offered in the DC task significantly influenced the uncertainty ranges elicited in the MBDC task.

[21] Further concluding remarks on all three chapters, in which we highlight potential future extensions, can be found at the end of this thesis.

A major contribution of this work is the development of a semi-parametric estimator for the analysis of uncertain preferences. The estimator we describe is commonly used in the medical literature and can be described as a three-state duration-dependent Markov model. It allows analysis of both the certainty range and the uncertainty range. We conclude that the semiparametric estimator presented here is an interesting and potentially fruitful technique for analysing uncertain WTP data. We have shown an example of how the estimator can be used to analyse MBDC data. Moreover, we test correlation in the width of the certainty and uncertainty ranges and independently test the effect of a prior DC question on the uncertainty range; analysis that would be very difficult without specifically modelling the thresholds between an individual's state of certainty about paying for a good.

We conclude that our modelling technique makes very few assumptions about the characteristics of the data. We believe that by extending the analysis of MBDC data outside of the previously used parametric methods we have broadened the potential for analysis on uncertain valuation preferences. As far as we are aware this is the first time that this form of multiple-state duration modelling has been used in economic analysis.

PART II

CHAPTER 2

OPTIMAL PATTERNS OF LAND USE USING SIMULATION MODELLING: ECOSYSTEM SERVICES AND MULTIPLE PURCHASERS

2.1. Introduction

The primary aim for landowners of agricultural land is to grow market goods such as food, at the same time agricultural land provides, or has the potential to provide, a wide array of non-market ecosystem goods and services (UK National Ecosystem Assessment 2011). Since those ecosystem services often have the characteristics of public goods, the landowners have little motivation to produce more or even preserve such services. As such, to deliver ecosystem services landowners will likely require some form of external incentive (Kemkes, et al. 2010). Mechanisms instituting such incentives, including regulation and direct payments, are on the rise around the world. The particular focus of the work described in this chapter is on Payment for Ecosystem Services (PES) schemes (FAO 2007; Engel, et al. 2008). A recent summary of the literature can be found in (Schomers and Matzdorf 2013), a review which documents the rapid growth in PES over recent years.

A particular feature of PES schemes which seek to incentivise land-management practices on agricultural land is that they regularly deliver multiple ecosystem benefit flows. A frequently cited example of a change that leads to multiple benefit flows is the planting of a riparian buffer, in which a strip of land along a watercourse is planted with vegetation, usually trees. The riparian buffer can improve water quality, reducing sediment, nitrate and phosphate runoff, while simultaneously sequestering carbon and providing habitat for wildlife (Salzman 2010). Indeed, agriculture can be managed to deliver a whole suite of ecosystem services including hydrological and climate regulation services, food and water production services, pollination services and cultural and recreational benefits (Millennium Ecosystem Assessment 2005). As such, reducing intensity of agricultural practices will likely deliver improvements in a variety of different ecosystem services resulting in benefits that will accrue to a variety of different groups. This chapter focuses on the issue of PES mechanism design when the activity incentivised through the scheme benefits multiple groups each of whom might be prepared to contribute to payments made through the scheme; that is to say, the design of multiple-purchaser PES mechanisms. That focus differs markedly from the majority of the PES literature that has largely concentrated on the single-purchaser problem.

While all purchasers may be interested in incentivising the same type of land-use change (for example, reducing the intensity of agricultural activity or taking land out of production altogether), it is not necessarily the case that each would choose for those changes to be instituted in the same locations. For example, a purchaser interested in improving water quality might gain most by reducing agricultural activity along water courses. Similarly, a purchaser interested in biodiversity outcomes may benefit most by instituting the same changes around previously established nature reserves (Lawton, et al. 2010). The first example represents a case where the ecosystem service benefits delivered by undertaking an activity in a certain location are determined solely by the characteristics of that location. The second example represents a case of spatial interdependency whereby the ecosystem service benefits of an activity in a location are determined in part by the activities undertaken in neighbouring locations (Goldman, et al. 2007). Recently, spatial targeting has been recognised as an important feature of PES mechanism design; see discussions of the agglomeration bonus in Wätzold and Drechsler (2014) and Banerjee et al. (2014), and policy relevance see the new Countryside Stewardship land-management scheme announced in the new Rural Development Programme for England (Defra 2014) and the Natural Capital Committee recommendations (Natural Capital Committee 2015).

While the literature has focused on the spatial purchasing decision for a single buyer, even when multiple benefit flows exist, in this chapter we focus on the issue of spatial coordination on the demand side of the market; that is to say, the question of which beneficiary buys land-management changes on which land parcels. As described subsequently we present a framework of methods for exploring potential efficiency gains from multiple purchaser PES schemes. For instance, if the buyers act independently and both adopt a PES scheme, the landowners may receive double payments for the change in land-management practice (Woodward 2011). Alternatively, the buyers may consider it the responsibility of the other buyer to pay, leading to free riding behaviour. Further adverse effects may occur when one buyer's decisions on the location of the land-management change adversely affect the benefit flows received by another buyer. In addition to adverse effects, synergies could result from multiple buyers working together, such as cost savings or greater overall ecosystem service flows (Venter, et al. 2009). Here we focus on the spatial coordination of the buyers' decisions but it is important to note that the framework of methods outlined in this chapter can easily be applied to considering the potential efficiency gains when the purchasers are paying for different land-management actions or even when the benefits from different actions vary according to the location. The first key contribution of this chapter is to develop a general framework of methods that can incorporate the spatial purchasing decision of multiple buyers.

Using our framework of methods we explore four multiple purchaser decision making institutions— three non-cooperative and one cooperative. In the first, we assume that the multiple purchasers act in complete independence and implement PES schemes simultaneously—*independent and simultaneous*. In the second, we assume that the multiple purchasers are independent but make their decisions sequentially—*independent and sequential*. In the third, we assume that the multiple purchasers enter into negotiation with each other—*negotiated*. Finally, we explore a fully cooperative decision making problem in which we assume the multiple buyers give up power over their decisions to a trusted third party—*cooperative*.

The second key contribution of this chapter is to employ the framework of methods to allow us to identify situations in which we might expect a multiple purchaser PES scheme to be practical. To compare the solutions from the four decision making institutions (outlined above) we use the concept of Pareto efficiency, that is solutions that can make one buyer better off without making another buyer worse off. We explore the non-cooperative and cooperative decision making problems in two simulation environments to assess the opportunities for realising Pareto-improving outcomes through a PES scheme when multiple independent groups stand to benefit from changing farmers' land-management practices. In both simulation environments we create spatial heterogeneity in the benefit flows for the multiple buyers and in the second simulation we include spatial interdependence in the benefit flows. Modelling a buyer with spatial interdependency in their benefits necessarily creates a non-linear problem, we show how our framework of methods is capable of creating solutions even for spatially interdependent benefits by forming a linearised version of the buyer's decision problem.

In addition to allowing us to identify situations in which the multiple purchasers might be practical, the method we develop allows us to identify optimal patterns of land use across a spatial landscape, potentially providing a useful tool for both ecosystem services buyers and policy makers—this is the third key contribution of this chapter. The decision making problems of the buyers are modelled in such a way as to be solvable by linear integer programming methods allowing for exact optimal solutions to be found over a reasonably large and heterogeneous landscape.

In Section 2.2 we introduce the key literature upon which this chapter builds, in particular we focus on the existing literature on multiple purchasers of ecosystem services and the literature on modelling the spatial decision making of buyers of ecosystem services. In Section 2.3 we set out our motivating example and build up a model that can be used to describe the spatial decision making of purchasers of ecosystem services, including showing how to model the costs, benefits and the motivations of the different buyers using integer programmes. In Section 2.4 we show how incorporating multiple buyers into a single PES scheme adds a level of complexity and then go on to develop four multiple purchaser PES institutions. In Section 2.5 and 2.6 we present two simulation environments in which we provide insights into and draw conclusions about the potential for multiple purchaser schemes using comparisons of the solutions gained from the non-cooperative and cooperative decision problems. In addition, we explore how the correlation between the production of ecosystem services affects the potential for Pareto-improvements. Finally, through a more realistic simulation environment, we present an example showing how the framework of methods can be used by policy makers to find optimal land-use patterns and in Section 2.7 we provide concluding remarks.

2.2. Literature review

Most PES schemes are run as monopsonies (Salzman 2009). Where multiple purchaser schemes have been successfully implemented it has generally been coordinated through a single organisation operating as a monopsony buyer, such as New York City's water authority acting on behalf of all of its customers by paying for a PES scheme in the Catskills catchment (Daily and Ellison 2002), or in Costa Rica where a PES scheme in which landowners are paid to protect forests developed through allowing new buyers to fit payments within an existing payment framework[22] (Sánchez-Azofeifa, et al. 2007; Pagiola 2008). A single purchaser PES scheme will

[22] The scheme in Costa Rica is financed through several sources: such as a fossil fuel sales tax, hydro-electric companies, the World Bank and the Global Environment Facility.

only be efficient when a sole individual or organisation are the only beneficiary and potential buyer of the ecosystem service(s)—in other words the single purchaser captures all the benefit and has monopsony power[23] (Kemkes, et al. 2010). If we consider that multiple groups benefit from the ecosystem services produced then excluding multiple buyers from the scheme may create inefficiencies. Such inefficiencies could arise from underutilisation of the potential available funds, for example beneficiaries free riding on the investment of a single purchaser, or from not maximising the welfare of all beneficiaries due to the single purchaser not fully considering the outcomes for each beneficiary or only considering their own welfare when making decisions about implementing a PES scheme.

There have been relatively few studies on the provision of ecosystem services from land-use change with multiple purchasers. Nevertheless, several studies have estimated the trade-offs from land-use change with a single decision maker, for example the trade-off between goods such as timber or agriculture and species conservation (Nalle, et al. 2004; Polasky, et al. 2005; Polasky, et al. 2008) or carbon storage and biodiversity (Nelson, et al. 2008; Venter, et al. 2009). In solving the problem the authors have assumed that a single budget exists and with that single budget a decision maker solves for the efficient outcomes, however with the trade-offs in the goods studied it is easy to imagine multiple buyers, for example, separate buyers of REDD carbon reduction objectives and biodiversity objectives in the Venter et al. (2009) paper. Assuming a single decision maker fails to account for complexities that arise when multiple economic agents with differing objectives participate within a single scheme, such as one buyer free-riding on another buyer's investment, or one seller receiving double or stacked payments.

In the context of a multiple-purchaser scheme, Woodward (2011) investigates the specific issue of the "stacking" of payments. Stacking refers to the practice of a landowner receiving multiple separate payments under different schemes or from different buyers as a result of a single land-management change delivering multiple ecosystem service improvements. The majority of current PES schemes either do not allow stacking or dissuade stacking through requiring each payment to generate

[23] In such situations, the transaction costs are likely to be low as coordination or negotiation need not occur between buyers and therefore basic Coasean rules are likely to lead to an efficient outcome.

additionality (the principle that landowners should not receive payments for benefits that would have occurred without their actions) (Salzman 2009; Woodward 2011). As such, sellers of ecosystem services cannot benefit from simultaneously selling in more than one market, this can be to the detriment of social net benefits as shown in Woodward (2011). In this chapter, we move away from considering the efficiency of the scheme from the seller's perspective, or society's perspective, to considering efficiency from the point of view of the multiple buyers.

The research reported in this chapter seeks to explore multiple purchaser PES schemes using methods of simulation modelling. A key element of that undertaking is in developing models of the different buyers' purchasing decisions. To that end we imagine a spatially heterogeneous landscape consisting of a large number of land parcels each managed by a separate seller in the PES scheme. The costs of paying for a change in land-management activity differ from parcel to parcel as does the improvements in ecosystem service provision delivered by that change. A buyer's choice problem is to purchase land-management changes through the PES on that set of land parcels that deliver the greatest net gains subject to constraints imposed, for example, by a limited budget.

In the context of a single purchaser, there is a long history of developing quantitative methods for spatially selecting land parcels to maximise biodiversity (Kirkpatrick 1983). The approach adopted in that literature generally involves mathematical programming. To optimise biodiversity outcomes a quantitative measure is required as an objective function; typically either a measure of species richness (the representation of all species from a list of target species) or a representation of habitat requirements has been used (Pressey, et al. 1997). Armsworth et al. (2012), in a recent application, quantify biodiversity by conducting a survey of the density and richness of bird species and then regress those measures against farm management variables to represent the response of biodiversity to farm management responses to agri-environmental policy scenarios across a spatial landscape. A review of the techniques used to incorporate spatial objectives over biodiversity outcomes in the conservation biology literature is given in Williams, ReVelle and Levin (2005). The vast majority of studies in that literature focus on a single, budget constrained purchaser; see Sarkar et al. (2006) for a review. We draw on the conservation biology literature and apply mathematical programming methods to model buyers of ecosystem services as independent agents, both when they make their decisions simultaneously and sequentially.

One of the key extensions required for our research is to go beyond the modelling of PES buyers as independent agents and explore how the purchasing behaviour of a buyer in a PES scheme might adapt to, and influence the purchasing behaviour of other buyers. One form such interactions might take is that of bargaining in which the purchasers negotiate with each other regarding which parcels each should purchase. Indeed, in this chapter we explore negotiation between buyers through applying techniques of non-cooperative game theory; particularly by applying Rubinstein's alternate-bargaining game (Rubinstein 1982). Similar problems have been studied in both mathematical programming under the name bi-level programming and in game theory as the Stackelberg game (Vallée and Başar 1999; Sinha, et al. 2014). In the Stackelberg game negotiations are assumed to proceed through multiple rounds of offer and counter-offer until agreement is reached. One problem encountered in modelling a multiple-round bargaining game where choices are patterns of purchases across a large spatial landscape is that those games constitute complex combinatorial choice problems. In this research, we address the problem of identifying solutions to such problems through the application of genetic algorithms. Genetic algorithms are a branch of evolutionary computation, which solve optimisation problems by imitating natural selection, selecting the 'fittest' solutions for 'breeding' in the next generation. We utilise the genetic algorithm over alternative search methods due to the computational efficiency of the genetic algorithm. A detailed account of genetic algorithms is given in Haupt and Haupt (2004).

A final element of the research in this chapter is to explore outcomes when the buyers act cooperatively. For cooperative decision making we use multi-objective programming techniques. A number of techniques are available to solve multi-objective optimisation problems, the one used in this chapter is the ε-constraints method proposed by (Chankong and Haimes 1983). In brief, the method proceeds by maximising the objective of one agent whilst introducing a constraint that the objective of the second must reach at least a certain level ε. By varying ε it is possible to trace out the whole production possibilities frontier. This method is utilised in an ecosystem services context in Polasky et al. (2005). Tóth et al. (2009) provide another example in their study on open space in Chicago, USA; in which their objectives are to jointly minimise costs and maximise the area of open space. An alternative approach to multi-objective optimisation problems is through the weighted-sum or scalarisation technique. With that approach the objective function for the maximisation is constructed as the sum of the objectives of the two agents. The two elements of that sum are weighted by scalars, $0 \geq w_1 \geq 1$ and $0 \geq w_2 \geq 1$ such that $w_1 + w_2 = 1$. The trade-off curve between the two objectives can be traced out by optimising the combined objective function while varying $w_1$over the range 0 and 1. This technique is used by Snyder et al. (2007) to maximise grassland habitat while at the same time minimising the pairwise distance between new habitat and existing reserves. Similarly, Venter et al. (2009) use the technique of assigning and then varying the weights of the two objectives for a study on jointly reducing carbon emissions and improving biodiversity.

2.3. Spatial targeting in PES schemes

Our motivating example concerns an agricultural landscape. That landscape comprises a large number of independent land parcels each managed by a farmer whose primary objective is to maximise profits from the production of food. Alternative land-management practices are possible, including taking land parcels (which might be a field or entire farm) out of agricultural production. For farmers, however, those alternatives are costly: they may require additional expenditure or result in a lower yield of agricultural output. At the same time, alternative land-management practices can deliver ecosystem services beneficial to one or many groups. A payment scheme in which the beneficiaries compensate the farmer for the costly land-management change (and is beneficial for all parties involved) can be described as Pareto-improving. The focus of this section is on the complexities that arise in realising such Pareto-improving outcomes through a PES scheme when multiple independent groups stand to benefit from changing farmers' land-management practices.

To address that question we simplify our analysis by assuming that farmers have an observable reservation price (perhaps greater than cost) at which they are prepared to adopt some particular change in management practices on a land parcel. Moreover, to maintain tractability we concentrate on the case where only one alternative land-management practice exists; to fix ideas, let us assume that that alternative might be taking a land parcel out of agriculture. In particular, we focus on the question of issues of spatial coordination on the demand side of the market; however the methods developed in this chapter could just as easily be applied to address the question of coordination between multiple buyers in determining which of several different changes in land-management practice to fund.

For the purposes of this chapter we imagine an agricultural landscape consisting of N land parcels. For each land parcel j in that landscape we assume that the farmer has a binary choice; carry on with normal production or undertake an alternative land-management practice. We use the decision variable x_j to denote the land-use choice on each land parcel. If a farmer carries on with normal agricultural production on their parcel of land then $x_j = 0$, however if that farmer agrees to undertake an alternative land-management practice then $x_j = 1$, such that $x_j \in \{0,1\}$. Building on that notation we denote a landscape configuration by the vector $\boldsymbol{x} = [x_1\ x_2\ \ldots\ x_N]$.

Each buyer makes a choice as to which land parcels to fund. We use X to denote the decision set; that is, the set of all land parcels that a PES buyer could convert to an alternative land-management practice by meeting the reservation price of that parcel's landowner for conversion. When the buyers are not choosing simultaneously, another PES purchaser may have already paid for management changes on one or more land parcels, therefore our formal definition of the decision set is given by $X = \{x_j : x_j = 0\}$. We describe a particular choice for buyer A by N-vector $\boldsymbol{x}_A$ where element $x_{A,j} = 1$ if buyer A chooses to fund land-management changes in land parcel j (such that that parcel must be an element of X) and $x_{A,j} = 0$ otherwise (either because that parcel is not in X or because buyer A chose not to fund changes in that parcel). The choice vector $\boldsymbol{x}_B$ is defined analogously for buyer B. Subsequently we shall define maximisation problems that identify optimal funding choices for each buyer. We will denote the solutions to those problems by the vectors $\boldsymbol{x}_A^*$ and $\boldsymbol{x}_B^*$. The vector $\boldsymbol{x}_A^*$ ($\boldsymbol{x}_B^*$) is 1 for all j where buyer A (buyer B) funds the land-management change and zero otherwise; this differs from $\boldsymbol{x}$, which is 1 for all parcels that have been converted across the whole landscape.

2.3.1. Costs

The creation of any PES scheme requires the exchange of information between buyers and sellers, one important piece of information is the cost to a farmer of adopting an alternative land-management practice. It can be assumed that farmers know that cost more accurately than the purchasers; they know up to some degree of certainty the opportunity cost of an alternative land-management change and also the reservation price they are willing to accept to undertake the change (Salzman 2005). A number of payment mechanisms have been used to facilitate the exchange of cost information, we briefly discuss a selection of these mechanisms and the complexities that result from considering spatial coordination with multiple purchasers.

A common payment mechanism is the fixed price mechanism in which the buyer posts a price that they are willing to pay to secure changes in land-management practices. Fixed price schemes often use a field-level approach in which farmers are paid for conservation on a field by field basis or alternatively paid per hectare or per metre. The Countryside Stewardship scheme in England is an example of a fixed price mechanism in which landowners can select from a number of land-management options each offered at a fixed price. From a buyer's point of view, a fixed price mechanism leads to a number of inefficiencies. Due to a single price being set across the whole landscape, fixed price mechanisms overpay landowners, in addition, because it fails to differentiate between the levels of ecosystem services provided, some landowners may be excluded that could have provided substantial benefits because their reservation price was higher than the fixed price. Furthermore, fixed price mechanisms fail to account for the additional benefits that might arise from funding spatial patterns, for example concentrating participation in the scheme into one location (so called agglomeration benefits). With regards to the latter, mechanisms such as the agglomeration bonus have been proposed which attempt to communicate the value of certain spatial configurations to the farmers with the expectation that the farmers can then coordinate their land-management practices in such a way as to be most beneficial to the buyers (Parkhurst, et al. 2002). However, such schemes may prove costly (particularly in terms of the cost of coordination) when scaled up to large landscapes.

A further payment mechanism that allows differentiation on price is provided by a reverse auction. For example, the Bushtender project in Australia (Stoneham, et al. 2003) and the River Fowey Upstream Thinking initiative in the UK (Day and Couldrick 2013). Reverse auctions ask landowners to declare a price at which they are prepared to participate in the PES scheme. With that information the buyer can choose which farmers to accept. Through combining the reverse auction with additional information on the expected ecosystem service benefits, spatial coordination is possible; for example, the buyer can weigh up the cost of the bids and the potential benefits that would arise from accepting a certain spatial pattern of bids.

In this chapter we begin from the assumption that the cost that must be paid by a purchaser to achieve some particular landscape configuration, is already known to the buyer. We do not investigate more complex institutional arrangements where farmers can negotiate payment, or where there are choices across different possible land-management practices. For example, let us assume that all farmers have submitted bids to a reverse auction. In that case, each cost c_j is known and independent such that a buyer's costs from a PES scheme can be represented by the following simple form:

$$c(x) = \sum_{j=1}^{N} c_j x_{A,j} \tag{2.1}$$

For the purpose of clarity in the simulations we make the further assumption that the cost for each land parcel is uniform across the whole landscape such that:

$$c(x) = \sum_{j=1}^{N} c x_{A,j} \tag{2.2}$$

2.3.2. Benefits

Despite knowing the reservation price of each farmer potential purchasers may still face a complex choice as to the particular set of land parcels to select for funding through the PES scheme. In particular, the production of many ecosystem services is spatially heterogeneous. As such, the choice of which land parcels to fund must at least consider the trade-off between the cost of funding land-management changes on a particular land parcel and the benefits of the ecosystem flows that arise from those changes. Moreover the level of flows of ecosystems services realised from funding change on one land parcel may depend on whether changes are also instituted on neighbouring or nearby parcels, a feature we describe as spatial interdependence. In that case, the particular spatial configuration of land parcels entering the PES scheme is vital in determining the overall benefit flows.

A further complicating factor concerns possible constraints on purchasers' actions. For example, they may be working within a certain budget or need to achieve a target level of benefit.

In order to capture these complexities in our simulation modelling we need to integrate the heterogeneity and interdependence of benefit flows from land parcels as well as the existence of constraints into the mathematical programmes with which we represent a PES purchaser's choice problem.

Spatial heterogeneity and interdependency of the benefits

The spatial production of ecosystem services can be thought of in terms of two key components, spatial heterogeneity and spatial interdependence. Spatial heterogeneity refers to the uneven nature of potential ecosystem service production across the landscape. For example, the production of carbon storage by planting trees is relatively spatially homogeneous, although the carbon storage potential of trees could depend on spatial characteristics such as altitude, soil type, exposure, latitude the primary determinant is likely, in most cases, to be the amount of trees planted[24]. In contrast, the production of water quality is spatially heterogeneous, such that converting certain land parcels (possibly locations close to a watercourse or on a steep slope) to an alternative land-management practice produces more benefits than other land parcels.

Spatial interdependence, on the other hand, refers to the relationship between land-management practices on one parcel of land and the productive capacity of other land parcels. For example, this could be a quantity interdependence, such that the aggregate abundance of a particular land use affects the benefits from converting another land parcel, or a configuration interdependence, such that locating certain land-management practices close together or far apart affects the benefits of converting another land parcel. An example of configuration interdependence can be found in the creation of a large contiguous natural habitat which delivers greater biodiversity benefits than the creation of a series of separate natural habitats of the same total area. A contrasting example concerns the location of a natural site used for recreation. The

[24] It is also possible to think of getting greater benefits from carbon storage in the long term by creating a large contiguous forested area; a large single area may be more resilient to environmental shocks, see Laurance et al. (2002).

closer that site is to an already existing site offering a similar recreational experience the less benefit the site is likely to provide.

In table 2.1 we use the two spatial production components—spatial heterogeneity and interdependence—to identify four classes of ecosystem service production. In the first, production is independent of location in the landscape (spatially homogenous); in the second, production differs across locations in the landscape (spatially heterogeneous); in the third, production is dependent on the aggregate abundance of alternative land-management practices (spatially interdependent - quantity); and in the fourth, production is dependent on the spatial configuration of the land use (spatially interdependent - configuration).

Spatial dependency of ecosystem service benefits
Spatially **homogeneous**
Spatially **heterogeneous**
Spatially **interdependent** (quantity)
Spatially **interdependent** (configuration)

Table 2.1. Spatial dependency of ecosystem service schemes.

For the purpose of illustration, we visualise how the land parcels which give the buyers the most benefit may change depending on the nature of production of ecosystem service benefits. We imagine that a particular agricultural landscape exists and that running through the landscape is a river flowing from north to south. Figure 2.1 shows such a landscape which has been partitioned into cells; each cell represents a land parcel with some cells representing agricultural land (red) and other cells river (blue). Each of the four panels in Figure 2.1 shows an example of the cells which give a buyer the most benefit but the variation results from assuming the ecosystem service benefits are dependent on one of the four spatial dependency classifications defined in table 2.1. Panel A shows an example for spatially homogeneous ecosystem service benefits, panel B shows an example for spatially heterogeneous ecosystem service benefits, panel C shows an example for spatially interdependent ecosystem services benefits based on the quantity, and panel D shows an example for spatially interdependent ecosystem service benefits based on the configuration.

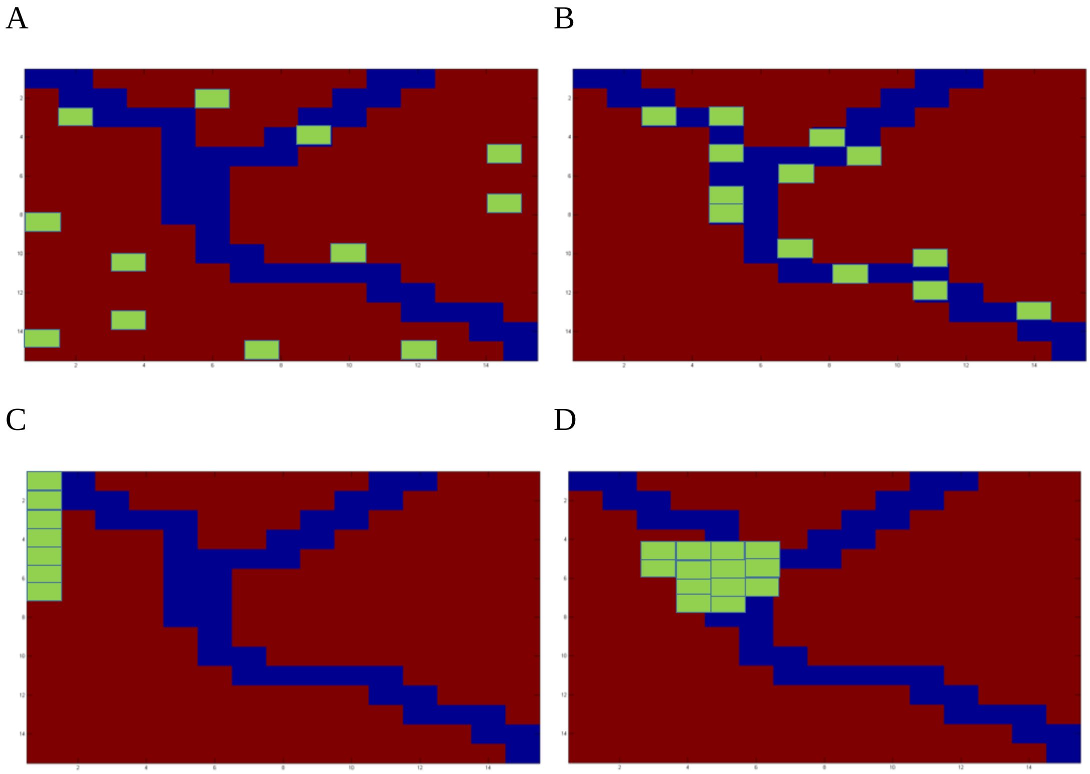

Figure 2.1. Example spatial land-use patterns.

Figure 2.1: In each picture above, the green shaded cells represent land parcels in which a buyer can receive their highest amount of benefits, the blue cells represent a river and red cells are agricultural parcels of land. (A) Shows a potential carbon storage buyer in which the ecosystem service benefits are spatially homogeneous such that each land parcel is of equal benefit and the buyer may choose based on the lowest cost. (B) Shows a potential water quality buyer in which the ecosystem service benefits are spatially heterogeneous such that land parcels close to the river are more beneficial. (C) Shows a buyer in which the ecosystem service benefits are spatially interdependent such that the benefits change based on the abundance of converted land parcels. (D) Shows a potential biodiversity buyer in which the ecosystem service benefits are spatially interdependent such that when land parcels close together are converted additional benefits accrue.

The buyers' benefits

We denote the benefits that a buyer gains from a particular landscape configuration by the function $b(\boldsymbol{x})$. For spatially independent ecosystem production that function takes the simple form:

$$b(x) = \sum_{j=1}^{N} b_j x_j \tag{2.3}$$

where b_j represents the benefits from land-management changes on each land parcel. The ecosystem service of carbon sequestration produced through the planting of trees could be represented through a benefit function such as (2.3). The benefits for a buyer looking to secure water quality improvements might also take a similar form in which land-management changes on land parcels near to water courses deliver high b_j.

For spatially interdependent ecosystem services the structure of the benefit function becomes more complex. We take biodiversity as an example of an ecosystem service whose production can be described as spatially interdependent and briefly discuss potential benefit functions used in the conservation biology literature. That literature indicates that the spatial interdependence of biodiversity production is often too complex to represent in a form suitable for inclusion in a mathematical model (Williams, et al. 2005). Rather than explicitly modelling the biodiversity production function, a common alternative is to use a proxy based on the pattern of land use (Nalle, et al. 2002; Polasky, et al. 2005; Polasky, et al. 2008).

Different desired landscape configurations call for different mathematical representations of the benefit function. Certain functional favour a single large area adjacent land parcels adopting alternative land-management practice, others a number of smaller areas, others can provide connectivity or a certain shape to the land parcels selected. Önal and Briers (2005), for example, use mathematical programming to minimise a benefits function calculated as the sum of the distances between neighbouring land parcels:

$$b(x) = \sum_{j=1}^{N} \sum_{k>j} d_{jk} x_j \, x_k \tag{2.4}$$

where x_j and x_k are pairs of binary land parcels and d_{jk} is the distance between them. Minimising equation (2.4) achieves clustering of the land parcels converted to an alternative land-management practice, since selecting parcels separated by the smallest distance d_{jk} adds the least to the objective value.

A number of alternative approaches exist; for example, prioritising proximity to particular spatial features by minimising the sum of pairwise distances to the spatial feature (Onal and Briers 2002); prioritising compactness of the selected land parcels to minimise fragmentation and boundaries (Tóth and McDill 2008); or forcing the selected land parcels to be fully connected (so you can walk between them) (Williams and Snyder 2005).

Accordingly, there are many options available for the benefit functions of spatially homogeneous, spatially heterogeneous and spatially interdependent ecosystem services. For the simulations presented in this chapter, we take the benefit function in equation (2.3) to represent the buyer of a spatially independent ecosystem service and equation (2.4) to represent the buyer of a spatially interdependent ecosystem service.

2.3.3. The purchasers' problems

With defined cost and benefit functions all that remains is to bring those together to form a choice problem which a purchaser is assumed to solve in selecting the optimal set of land parcels in which to fund changes in land-management practice. Since the decision involves a series of binary choices for each land parcel, that optimisation takes the form of an integer programme.

In particular, we assume that buyers of ecosystem services seek to maximise the benefits that receive from ecosystem service flows whilst being constrained by the financial costs of paying farmers to convert to alternative land-management practices. In general form, a single purchaser maximisation problem can be formulated as follows:

$$\begin{gathered}\max_{x_j \in X} F(\boldsymbol{x}), \\ s.t. \, G(\boldsymbol{x}) \leq 0\end{gathered} \tag{2.5}$$

where $F(x)$ is the objective function and $G(x)$ is the constraint set (both equality and inequality constraints). As we are considering an integer programme one constraint in $G(\boldsymbol{x})$ will be to define each x_j in $\boldsymbol{x}$ as a binary integer, that is to say, $x_j = 0$ (not in the PES scheme) or $x_j = 1$ (land parcel j is in the PES scheme is converted to an alternative land-management practice). As such $x_j \in \{0,1\}$ is a constraint in all subsequent programmes.

To model the buyers decisions as an integer programme requires us to consider that the buyers may have different motivations: some may want to protect a key input into their business, for example, water companies or hydroelectric power companies (Day and Couldrick 2013); others may want to improve ecosystem services for the benefits to society or local communities; still others may want to offset environmentally harmful activity elsewhere (TEEB 2012). Differences in the motivations of the purchasers lead to differences in how a buyer's integer programme should be constructed. To differentiate, we classify the different buyers based on three characteristics—budget constrained, target constrained and profit maximising. For budget constrained buyers we imagine a potential purchaser wants to achieve the most ecosystem services they can with a limited budget. For target constrained buyers we imagine a potential purchaser wants to achieve a target level of ecosystem service and they want to achieve that using the minimum expenditure. For profit maximising buyers, we imagine a potential purchaser wants to buy ecosystem services up to the quantity which maximises their profit. Table 2.2 summarises the potential purchasers of ecosystem services in terms of costs and benefits, providing representation of the objectives and example purchasers. As before $b(x)$ represents the benefit function for the buyer from the PES scheme and $\sum_{j=1}^{N} c_j x_{A,j}$ the sum of the independent costs of land parcels in the scheme; $\bar{T}$ is the target amount of ecosystem service benefits and $\bar{B}$ is the budget.

Ecosystem service	Simplest representation of objective	Potential purchasers
Budget constrained	$\max_{x_j} b(x)$ $s.t. \sum_{j=1}^{N} c_j x_{A,j} \leq \bar{B}$	National, regional and local Governments, NGOs
Target constrained	$\min_{x_j} \sum_{j=1}^{N} c_j x_{A,j}$ $s.t. b(x) \geq \bar{T}$	National, regional and local Governments, NGOs, Private companies
Profit maximising	$\max_{x_j} \quad b(x) - \sum_{j=1}^{N} c_j x_{A,j}$	Private companies, offsetting

Table 2.2. Types of buyers and their objectives.

Solving a buyer's integer programme means finding the combinations of x_j's that maximise their objectives; in other words, finding which combination of land parcels when changed to an alternative land-management practice leads to the highest objective value for the buyer. We denote this optimal combination of land parcels as $\boldsymbol{x}^*$ such that $\boldsymbol{x}^* = [x_1^*, x_2^*, \ldots, x_N^*]$.

In the current literature, the most commonly studied objective is one in which decisions are constrained by a limited budget, particularly relating to payments for biodiversity (Williams, et al. 2005). We assume budget constrained buyers for the simulation in this chapter but in Appendix B1 we discuss both target constrained and profit maximising buyers. A profit maximising programme requires maximising the difference between buyers' benefits and costs, to do this requires the benefits to be measured in monetary terms and assuming diminishing marginal benefits means that the problem becomes inherently non-linear. Appendix B1 shows how even a non-linear profit maximising problem can be included in our framework of methods.

2.4. Multiple-purchaser problems

In this section, we consider another tier of complexity in a buyer's spatial decision problem, the existence of another buyer. In doing so, we examine outcomes under increasing levels of sophistication and coordination in the interactions of buyers in their purchasing behaviour. We focus on a case where there are only two beneficiaries of ecosystem services in the locality therefore only two potential buyers of ecosystem services—buyer A and buyer B. In particular, we present four decision making problems for buyer A and buyer B: in the first the buyers are independent and make their decisions simultaneously and without regard for the actions of the other buyer; in the second the buyers are independent and make their decision sequentially where the second buyer to decide is aware of the first buyer's purchasing decisions; in the third the buyers make their buying decisions strategically as the outcome of a process of negotiation; in the fourth the buyers make their decisions cooperatively. For clarity of results we assume that the buyers do not face transaction costs, apart from the bargaining delay costs which form part of the model describing the negotiated purchasing decision.

Independent and simultaneous decision making

The first decision making problem we consider involves two independent buyers making simultaneous choices. In essence this is a simple extension of the single purchaser problem represented in equation (2.4) in which two buyers act without knowledge of the existence of the other buyer. Buyer A and buyer B's problems are represented as follows:

Buyer	Problem	Solution Vectors
Buyer A	$\max_{x_j \in X} F_A(x)$ $s.t.\ G_A(x) \leq 0$	x_A^*
Buyer B	$\max_{x_j \in X} F_B(x)$ $s.t\ G_B(x) \leq 0$	x_B^*

where $F_A(x)$ denotes the objective function, $G_A(x)$ the constraint set and x_A^* the solution vector of buyer A, and $F_B(x)$ denotes the objective function $G_B(x)$ the constraint set and x_B^* the solution vector of buyer B.

One immediate insight is that when buyers' decisions are made independently and simultaneously, there is nothing to stop $x_{A,j}^* = x_{B,j}^*$ such that both buyers choose to fund changes in the same land parcel. In other words, when buyers do not consider each other's purchasing behaviour, they may both elect to pay the same farmer who will enjoy "stacked" payments. Clearly, from the buyers' point of view efficiency gains are possible from alternative purchasing choices.

Independent and sequential decision making

The second decision making problem we consider involves independent buyers making sequential choices. Such a situation may occur when one buyer, the first mover, choses to act independently in a locality funding changes in land-management practices. At a later time, another buyer, the second mover, aware of the actions of the first mover chooses to invest in the same locality. A general form representation of this decision problem is:

Buyer	Problem	Solution Vectors
Buyer A first mover	$\max_{x_j \in X} F_A(x)$	x_A^*

	$s.t.\ G_A(x) \leq 0$	
Buyer B second mover	$\max_{x_j \in Y} F_B(x)$ $s.t\ G_B(x) \leq 0$	x_B^*

where $F_A(x)$ and $G_A(x)$ represent the objective function and constraint set of buyer A (assuming buyer A is the first mover) and $F_B(x)$ and $G_B(x)$ represent the objective and constraint set of buyer B (assuming buyer B is the second mover).

The important difference between the simultaneous problem and the sequential problem is that in the sequential problem the second mover is advantaged from knowing which land parcels have been funded by the first mover. As such, the second mover can avoid stacking payments for those land parcels and instead use their budget to select alternative land parcels. We denote this in the problem by maximising over $x_j \in Y$, where $Y \subset X$ such that $Y = \{x_j : x_{A,j}^* = 0\}$, in words, Y is the set of x_j that have not been selected in the first mover's solution. A simple result that could be concluded from this setup is that if a buyer understands the benefits of moving second then we would expect to see free riding behaviour in which buyers wait for the other buyer to move first.

Negotiated decision making

The third non-cooperative decision making problem we consider is strategic negotiation between buyers. Consider how a strategic buyer would act in our independent and sequential problem. A strategic player would not, when moving first, simply choose the land parcels that are most valuable to them. Rather they would consider which land parcels the other buyer will choose, given any particular purchasing pattern of the first buyer. The optimal land parcels to purchase for the first buyer will therefore be to avoid purchasing sites that, despite providing buyer A with high benefit, would be purchased by Buyer B in their subsequent choice. Accordingly, in our simple setup, strategic buyers gain an advantage by moving first.

Things get more complex when we imagine a situation in which buyers negotiate with each other in reaching a binding agreement over which land parcels each will fund. To explore negotiated outcomes we consider a form of strategic non-cooperative bargaining famously modelled in Rubinstein's alternating bargaining theory (Rubinstein 1982). In that model, bargaining proceeds via a structured non-cooperative game in which two players make alternate offers to one another. In our case, those offers would concern the land parcels that the offering buyer would choose to fund if they were given the strategic advantage of buying first. For simplicity, we assume those decisions are made with perfect information regarding the purchasing preferences of the second buyer. If an offer is accepted then agreement is reached. Alternatively, an offer can be rejected in which case the second buyer is given the option of making an offer. The negotiation might be played out over a fixed number of rounds of offers or over an infinite horizon.

The simplest form of such a bargaining institution is the one-round ultimatum game in which one player, the leader, makes an offer that the other player, the follower, can either accept or reject. If the follower rejects the offer then both players get nothing. The subgame-perfect-equilibrium for the ultimatum game is one in which the leader should make a proposal in which they get all the benefits and the follower should accept that because they can do no better. In the context of our two PES buyers, a one-round negotiation analogous to the ultimatum game can be represented by the following decision problems:

Buyer	Problem	Solution Vectors
Buyer A Leader	$\max_{x_j \in X} F_A\left(\boldsymbol{x}_A + \left(\max_{x_j \in \{x_j : x_{A,j}=0\}} F_b(\boldsymbol{x})\ s.t.\ G_B(\boldsymbol{x}) \leq 0\right)\right)$	x_A^*
Buyer B Follower	$\max_{x_j \in Y} F_B(\boldsymbol{x})$ $s.t.\ G_B(\boldsymbol{x}) \leq 0$	x_B^*

Observe how the leader's problem takes the form of a bi-level programme, that is to say, a mathematical programme that itself contains an optimisation problem. In this literature, this form of problem is often referred to as a Stackelberg game. In our case, the leader perfectly anticipates the optimisation problem of the follower, knowledge that they exploit in choosing which land parcels to include in their proposal. Since this is a simple one round negotiation, the follower has no choice but to accept that proposal and (since we assume perfect information) choose a set of land parcels to fund themselves which perfectly matches the prediction of the leader. The key point to note is that in this strategic setting a buyer's funding proposal anticipates how that proposal will affect the choices made by the other buyer.

A simple extension to the ultimatum game would be to allow for a second round of negotiation in which the second player can refuse the first player's initial offer and propose their own counter-offer. In this case, the second player gets to make the final proposal and would therefore be able to claim all the benefits. Indeed, in this form of bargaining game it is always the player entitled to make the final offer in the negotiation who stands to be most advantaged.

To increase realism, the usual assumption is that the number of bargaining rounds is unlimited since no player would agree to participate in a bargaining institution in which the other player was privileged with last mover advantage. Moreover, bargaining itself is considered a costly endeavour; each time a player rejects an offer they delay the reaching of an agreement and delay costs are experienced by all buyers since a further round of negotiation is required.

To understand how delay costs affect negotiations, consider a simple two round negotiation; the first player makes a proposal and the second player has the option to refuse and make a final counter-offer. With no delay costs, the first player can do no better than make an offer that optimises the second player's outcome; any other offer will be rejected by the second player allowing them to achieve that same outcome with their counter offer. When that rejection is associated with a delay cost, however, things are a little different. The first player knows that if the negotiation goes to an extra round of negotiation the benefits are reduced by the amount of the delay cost. As a result, they can make a first round offer which claims that delay cost for themselves at the expense of second player's payoff. Indeed, the subgame-perfect equilibrium in any finite length bargaining procedure is determined by the delay costs to the players. If there are unlimited rounds of negotiation then, with equal delay costs, the equilibrium will tend towards an equal split of the benefits.

In our simulation of negotiations between PES buyers we use a multiple round Stackelberg game and simplify by assuming equal delay costs for the buyers and setting the delay cost at exactly the cost of a single land parcel, c. Accordingly, with each round of negotiation each budget-constrained buyer incurs a cost which reduces the number of land parcels they can afford to fund by one. Those negotiation costs can be simply included in the budget constraint presented in 2.3.3 according to:

$$\sum_{j=1}^{N} cx_{A,j} \leq \bar{B} - c(d-1)$$

where the budget is reduced by the product of the cost of a single land parcel c and the round of negotiation d, such that in round one $d - 1$ equals zero and therefore the full budget is available.

In our analysis, we explore how different numbers of rounds of costly negotiation affect non-cooperative bargaining outcomes. The maximum number of rounds of negotiation is denoted $\bar{d}$, we calculate solution vectors x_A^* and x_B^* for a range of $\bar{d}$. In addition, the buyer moving first is varied for each of the $\bar{d}$, such that, a single $\bar{d}$ gives two solutions—one when buyer A is the leader in the first round and one when buyer B is the leader in the first round.

To solve the Stackelberg game over multiple rounds of negotiation we use backwards induction, starting by setting $d = \bar{d}$ we solve for a solution that optimises the benefits realised by the buyer with the advantage of being the proposer in the final round of negotiation. The method then moves back one round of negotiation so that $d = d - 1$ and the buyer that was previously the leader becomes the follower and the buyer that was previously the follower becomes the leader. In this problem we know that the new leader must make an offer that ensures that the follower receives at least as much benefit as realised in the solution to the problem when $d = \bar{d}$, minus the delay cost. Following that logic back up through the rounds of negotiation of the game until $d = 1$, we solve for the subgame-perfect equilibrium offer made by the leader in the first round of negotiation that will be accepted by the follower.

Since solving for optimal solutions to a multiple buyer negotiated decision problem is a difficult combinatorial problem, we employ a heuristic search method called the genetic algorithm[25] to solve the optimisation problem. An overview of the specific algorithm used is given in table 2.3. The basic premise of the genetic algorithm is to mimic the process of natural selection. To that end, many solution vectors $\boldsymbol{x_A^*}$ ($\boldsymbol{x_B^*}$) of the leader's problem are generated and those solution vectors together are called the

[25] The genetic algorithm is from the branch of computer science known as evolutionary computation used for solving combinatorial constrained optimisation problems.

population. The starting population is a number of randomly generated solution vectors across the solution space. The follower then performs a maximisation of their benefits over the $x_j's$ not selected in the leader's solution vector for each of the solution vectors in the population. Thus the follower solves the same number of problems as the number of vectors in the population. The follower's problems lead to different levels of benefits for the leader and subsequently those solution vectors in the population that result in a high payoff for the leader are selected for reproduction (or crossover) in the next generation. The crossover process combines $x_j's$ in one $\boldsymbol{x}_A^*$ with those of another $\boldsymbol{x}_A^*$ to create "offspring" that contain traits of both $\boldsymbol{x}_A^*$. We use a uniform crossover method (see Haupt and Haupt (2004)) that compares the $x_j{}'s$ of each parent one by one and flips the $x_j{}'s$ in the offspring according to a probability parameter (0.5 in our case) from 0 to 1 or 1 to 0. To protect the genetic algorithm method from fixing on a local optimum random mutation is added to the solution vectors, a mutation parameter of 0.001 is used such that approximately 1 x_j in every 1000 is flipped.

Parameters for multi round negotiation	Set: delay cost (cost of negotiation) Set: maximum number of rounds of negotiation
Start	Solved by backwards induction so start at the maximum number of rounds of negotiation and subtract the delay costs from the budgets of the buyers.
1. Initialisation	The 'population' of leader choices is randomised for the first iteration.
2. Fitness	The follower moves second optimising their objective subject to the leaders moves.
3. Selection	The solutions are ranked according to the benefit to the leader and the top 50% are kept with the bottom 50% discarded.
4. Crossover	A new generation of populations are created by recombining the x^*'s of two parent solution sets. In addition, the top performing population is carried over intact to the new generation.
5. Mutation	A mutation operator randomly flips $x_j's$ from ones to zeros and zeros to ones. A mutation parameter of 0.001 is used such that approximately 1 x_j in every 1000 is flipped.

6. Feasibility	If the populations are infeasible in that they do not meet the budget constraint of the problem then they are replaced by a randomly generated population that does meet the budget constraint.
7. Update population	The population of leader choices is updated.
8. Iterate until conversion	The process repeats until the genetic algorithm reaches convergence in that no improvements in the objective value is observed for a certain number of iterations (in our case 50 iterations).
Next round of negotiation	• Leader becomes the Follower and vice versa • Add the increment of the delay cost, c, to the budget of both buyers. • Include the constraint that the leader must offer the follower at least the amount they would have received as leader in the previous round.
Stop	Since backwards induction is employed the process stops when all rounds of negotiation have been solved.

Table 2.3. Genetic algorithm for non-cooperative multi-round bargaining.

Fully cooperative decision making

In the final multiple-purchaser decision making problem we imagine full cooperation between the buyers. In this problem the buyers give up control of their decisions to a trusted third party who jointly optimises the objectives of both buyers. Recall, that while our two buyers both receive benefits from the same change in land-management practices, they gain those benefits from the impact those changes have on two different ecosystems services and in different quantities. One way in which the trade-offs among ecosystem service benefits can be illustrated is using the production possibilities frontier (Nalle, et al. 2004; Kline and Mazzotta 2012; Lester et al. 2013). The production possibilities frontier shows the combinations of ecosystem services that can be produced on a landscape given the landscape's capacity for production. The capacity depends on physical features such as the size, also the existing land-use patterns and the spatial characteristics and interactions.

The production possibilities frontier combines the complex relationship between the production of one ecosystem service in terms of the production of another ecosystem service. In Figure 2.2, we show an example with two ecosystem services (I) and (II).

The production possibilities frontier traces out the limit of the joint production of both ecosystem services given a fixed set of inputs and technology. In Figure 2.2 any point inwards (to the left) of the production possibilities frontier is an achievable combination of the production of the two ecosystem services and any point outwards (to the right) of the production possibilities frontier is unattainable given a fixed set of inputs. In our case, the set of achievable combinations of production is defined by the set of all possible configurations of land parcels on which the two buyers could fund land-management changes within the limits set by their constraints.

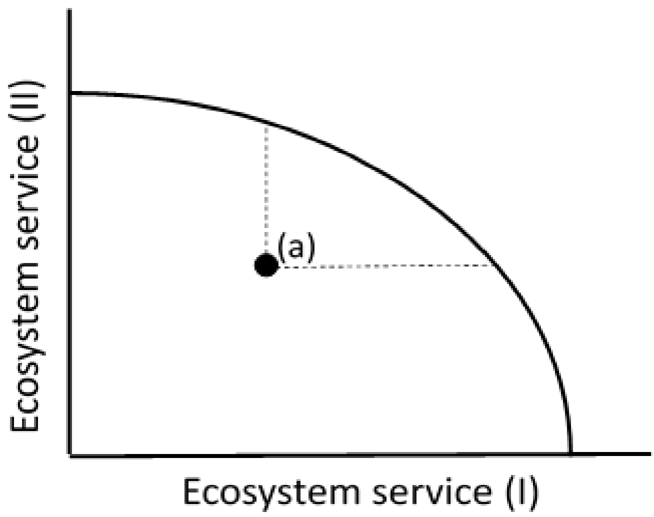

Figure 2.2. Example of joint production possibilities frontier for ecosystem services.

Point (a) in figure 2.2 represents the ecosystem service levels enjoyed by each buyer under some current funding allocation resulting in a particular landscape configuration. Pareto-improving outcomes would be ones that lead to landscapes that provide greater production of both ecosystem service (I) and ecosystem service (II). From point (a) anywhere within the dotted lines provides a Pareto-improvement and the ideal outcome would be to move onto the production possibility frontier itself.

The frontier itself is based on a fixed set of resource inputs and a fixed level of technology, if those inputs or technology are changed, for example, by a buyer increasing their budget so they can pay for more land parcels then the frontier can shift (outwards for increasing inputs and inwards for decreasing inputs).

The frontier can also be defined by its shape, which describes the relationship between the production of the two ecosystem services, for a review of the different possible shapes in the context of marine ecosystem services see Box 1 in Lester et al. (2013).[26] The three frontier shapes most widely discussed are (i) a direct trade-off, (ii) a concave trade-off and (iii) a convex trade-off. A direct trade-off between the two ecosystem services results in a linear production possibilities frontier, in that situation a land-use pattern that increases the provisioning of one service results in a proportional decrease of the other service, with no diminishing returns. A concave frontier, as in Figure 2.2, means that although there is a trade-off, there are scenarios that increase the delivery of one service substantially without a large cost to the other service. A convex frontier means that achieving even a small increase in the provisioning of one service comes at a large cost for the other service (Lester et al. 2013). Some have also suggested the complex production of ecosystem services may be subject to non-convexities (such as non-monotonic trade-offs or threshold trade-offs), with Brown et al. (2011) showing this to be the case when strong positive externalities are present.

To form the production possibilities frontier for two ecosystem services requires a method of joint optimisation in which consideration is given to the differing objective functions of the two buyers. In this chapter, we use the ε-constraints method. To implement the ε-constraints method one objective is maximised whilst introducing a constraint that the second objective must reach at least a certain level ε, by varying ε it is possible to trace out a whole production possibilities frontier. The ε-constraints is given as follows:

Buyer	Problem	Solution Vector
Buyer A and Buyer B	$\max_{x_j \in X} F_A(x)$ $s.t. F_B(x) \geq \varepsilon_i \ \forall i \in \{1,2,\dots,M\}$ $G_A(x) \leq 0$	x_C^*

where $F_A(x)$ is the objective function of buyer A and $F_B(x) \leq \varepsilon_i$ ensures that the buyer B's objective value is at least the value ε_i, with i representing the choice of one objective value out of M for the second buyer. Again $G_A(x)$ represents the constraint set of the first buyer and the joint solution vector is given by $\boldsymbol{x}_C^*$.

[26] The absolute slope of the production possibilities frontier is called the marginal rate of transformation, it is the rate at which one ecosystem service must be given up to produce more of the other ecosystem service.

Comparing the non-cooperative solutions with the production possibilities frontier allows us to examine the potential for gains from the different decision making to be assessed. For example, if the non-cooperative solutions sit to the left of the frontier, then we know that a cooperative solution has the potential to produce more ecosystem services for both buyers with the same fixed inputs.

2.5. Simulation 1

In this section we employ methods of simulation to explore outcomes under the different types of multiple-purchaser PES institutions outlined above: independent and simultaneous, sequential, negotiated, and cooperative. Our simulation environment maintains the assumption that the PES consists of just two buyers. Both buyers wish to fund the same land-management change on land parcels in a landscape but realise benefits from those changes through increases in two different ecosystem service flows. In the initial simulations reported in this section, we adopt the simplest assumption regarding the production process underpinning ecosystems services. In particular, we assume that the benefit realised by a buyer from a change in land-management practice on any one parcel is a constant that is independent of land-management practices on adjoining land parcels. Of course, those benefits may differ across the landscape such that there is spatial heterogeneity in production (see Section 2.3.2). Likewise the benefits realised from changes on a particular land parcel may differ across the two buyers. Indeed the key issue we explore in this simulation environment is how outcomes differ under different assumptions regarding the level of correlation in the two buyer's benefits across land parcels. Negative correlation (trade-offs) occur when an increase in one buyer's benefits comes at the expense of the other buyer. Positive correlation (synergies) occur when an increase in one buyer's benefits causes an increase in another buyer's benefits. If that correlation is positive, then both buyers will be motivated to invest in changes in the same land parcels through the PES, if it is negative then the two buyers will target land-management changes in different land parcels.

Setup

To explore those issues we created a simulation environment consisting of 225 square parcels of equal size arranged on a 15x15 square grid[27]. In each land parcel j some form of agricultural production is taking place. For the purposes of providing clear results, we assume a single land-management change on a land parcel produces two ecosystem services, one is beneficial for buyer A and the other is beneficial for buyer B. The key consideration of the buyers, results from the spatial heterogeneity in benefits and therefore the difference in benefits realised at different locations; for example, one particular land parcel may produce a lot of the ecosystem service which is beneficial to buyer A if switched to an alternative land-management practice but very little of the ecosystem service beneficial to buyer B.

Our simulation environment allows us to construct landscapes that offer different spatial patterns of benefit to the two buyers. The benefits to each buyer are simulated using two random draws from the standard normal distribution to create two vectors of benefits $\boldsymbol{b}_x$ and $\boldsymbol{b}_z$, following this we define a third vector, $\boldsymbol{b}_y = r\boldsymbol{b}_{\boldsymbol{x}} + \sqrt{1-\boldsymbol{b}_x^2}\boldsymbol{b}_z$, where $\boldsymbol{b}_y$ is also standard normal and r is the correlation coefficient between the vectors $\boldsymbol{b}_x$ and $\boldsymbol{b}_y$. The benefits across the landscape for buyer A can be represented by $\boldsymbol{b}_x$ and for buyer B, $\boldsymbol{b}_y$ and r can be varied from -1 to 1 to create specified correlation between the buyers' benefits. It should also be noted that the benefit vectors are shifted such that each element is strictly positive. To illustrate observe Figure 2.3 which depicts three different simulated landscapes that differ with regard to the levels of spatial correlation in the two buyer's benefits. A coefficient of 1 specifies that the benefits for buyer A are perfectly positively correlated with buyer B, a coefficient of 0 specifies no correlation, and a coefficient of -1 specifies perfect negative correlation between the benefits of buyer A and buyer B. In the figure darker cells represent higher levels of ecosystem service benefits to the buyer.

---

[27] We use a square land parcels on a square landscape of 15x15 land parcels for the simulation environment analysis although it should be noted that the same methodology can be applied to other geometric designs. As suggested by the agglomeration bonus literature using network games in laboratory experiments the shape of the network/landscape can affect the outcome (Cassar 2007; Banerjee et al. 2012; Banerjee et al. 2014). This is something we do not explore in this thesis but could be a topic worthy of future research.

	Buyer A	Buyer B
Perfect Negative Correlation r = -1	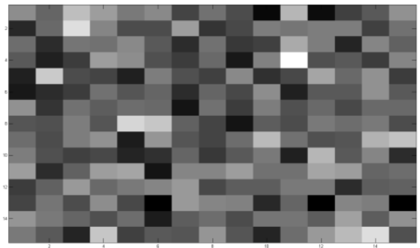	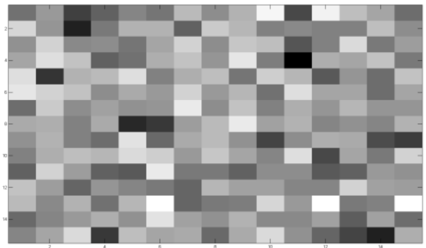
No Correlation r = 0	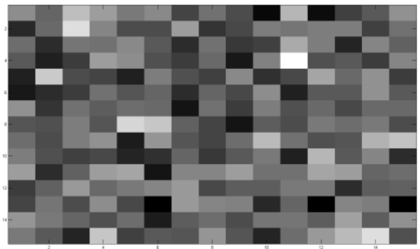	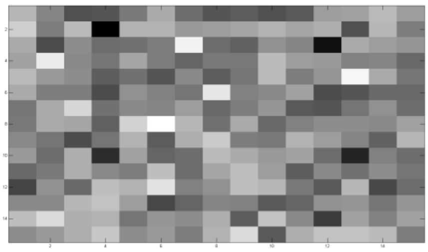
Perfect Positive Correlation r = 1	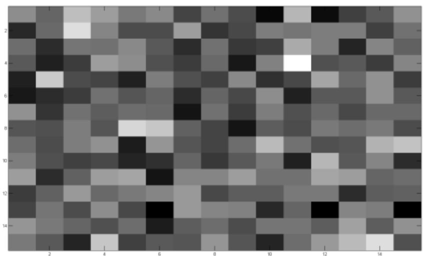	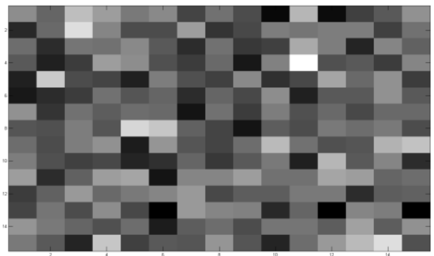

Darker cells represent higher ecosystem service benefits to the buyer.

Figure 2.3 Ecosystem service benefits to the buyers, showing three different correlation coefficients

Given these assumptions, we can produce a specific form for the mathematical programme describing each buyer's PES purchasing decision problem. In particular, a buyer whose benefits are dependent on the spatial heterogeneity of the landscape can be represented using an objective which maximises the sum of benefits, b_j realised by the buyer from funding changes on land parcels whilst meeting a budget constraint, $\bar{B}$ (equation 2.14)[28]:

$$\max_{x} \sum_{j} (b_j x_j)$$

$$s.t. \sum (c x_j) \leq \bar{B} \tag{2.14}$$

$$x_j \in \{0,1\} \quad j = 1,2, \ldots, N$$

[28] We have chosen to use a budget constrained optimisation problem although this could easily be substituted for a target constrained or profit maximising objective.

where we assume that the benefits to each buyer are strictly positive such that each land parcel when converted to an alternative land-management practice contributes at least some small positive benefit. Furthermore, we assume that the costs, c are uniform across the whole landscape.

Our simulation environment was developed in MATLAB R2013a and used the solver CPLEX 12.5.1 to identify optimal solutions to the integer-linear program described by the PES purchasing problem in (2.14). Modelling the buyer's choices as integer linear programmes creates complex combinatorial problems. In themselves they are difficult optimisation problems to solve, however, in our negotiated and cooperative multiple-purchaser decision we introduce substantial additional complexity. In our negotiated decision making model we nest the integer-linear programme within a heuristic search method, the genetic algorithm, which is solved over a number of rounds of negotiation. In our cooperative decision making model we jointly optimise the objective functions of two buyers' PES purchasing decisions and solve multiple times to trace out the production possibilities frontier. A sample of the code used to achieve the results is available in appendix B2.

Results

Figure 2.4 provides a graphical illustration of the PES-purchasing outcomes under the four different institutional setups. The Figure shows empirical outcomes from our simulation illustrating for each institution the benefits realised by Buyer A against those realised by Buyer B.

Observe that independent and simultaneous decision making gives a single solution in which both buyers act independently. In contrast, sequential decision-making could result in one of two solutions; one in which Buyer A moves first and one in which Buyer B moves first. Finally, negotiated decision-making results in a range of solutions depending on which buyer makes the first offer, the number of rounds of negotiation, and the costs to each buyer of a delay in reaching a bargaining solution. In figure 2.4 we visualise a selection of negotiated outcomes by varying both the number of rounds and negotiation and which buyer makes the first proposal.

Interestingly, with two buyers whose benefits depend only on the spatial heterogeneity across the landscape we find that one round negotiated decision making in which buyer A is the proposer provides the same solution as the sequential solution in which buyer

B is the first mover. Clearly, this is a result of the simplifications assumed in this simulation environment. In particular, the uniform costs across the landscape mean that each buyer can afford exactly p land parcels, the buyers will therefore choose the p land parcels that provide the most benefits to them. The best land parcels for one buyer will either be common to the other buyer or not, the interesting case is when common land parcels are in both buyers' best p land parcels. In that situation, one buyer gains an advantage from either being the leader in the one round negotiation, in which case they can leave the common land parcels for the other buyer and buy their second best land parcels, or as the second mover in the sequential decision making in which case the other buyer will have already purchased the common sites.

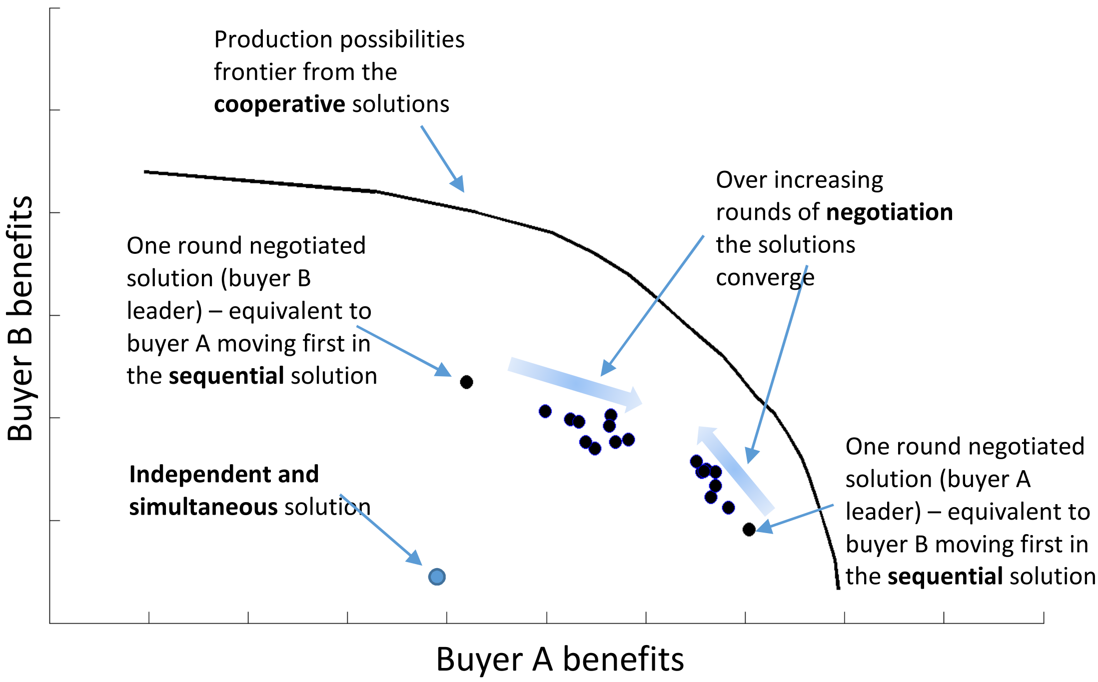

Figure 2.4. An example landscape with no correlation between the ecosystem services, showing solution outcomes for all four non-cooperative and cooperative decision making problems.

Figure 2.4: The benefits to each buyer are correlated random values simulated using the process outlined in the setup.

The equivalence between the sequential and negotiated outcomes only applies to the one round negotiated solution. As the number of rounds of negotiation increases, the negotiated solutions tend to converge. For the remainder of this simulation we focus on the negotiated solutions over the sequential solution, however, in the second simulation we investigate spatial interdependence in the ecosystem service benefits flows such that the equivalence between the two decision making problems is no longer present.

From the example results in Figure 2.4 it is clear that negotiated decision making provides Pareto-improving solutions when compared to independent and simultaneous decision making, moreover, cooperative decision making provides Pareto-improving solutions on negotiated decision making. For cooperative decision making, both buyers give up power over their decisions to a third party, by jointly optimising the objectives of both buyers we see a Pareto-efficient productions possibilities frontier.

Negotiated decision making

For each different negotiated solution that lies below the production possibilities frontier, cooperative decision making offers the opportunity for Pareto-improvements that would take the outcome to the production possibilities frontier. Moreover, if there exists a Pareto-improving zone, which we define as feasible solutions that provide Pareto-improvements for both buyers over all the negotiated outcomes, then there must exist solutions which Pareto-dominate any negotiated bargaining outcome. This distinction is important because if a Pareto-improving zone exists then, whatever the result of the negotiation, each of the buyers can realise a Pareto-improvement from participating in a cooperative mechanism. These Pareto-improving zones are akin to the self-enforcing properties of cooperative international environmental agreements discussed in Barrett (1994) such that neither buyer would be willing to sign up to a cooperative agreement (or indeed stay in a cooperative agreement) if they have a more attractive alternative option.

Our simulations reveal that Pareto-improving zones tend to exist when correlation is low to moderate. In other words, when the two purchasers prioritise different land parcels. If the two purchasers both prioritise changes in land-management practice on the same land parcel, then the opportunities for Pareto-improvements decline; above a certain threshold in the correlation coefficient (approx. $r = 0.65$) we see a switch from the Pareto-improving zone existing to it not existing. When both purchasers prioritise the same land parcels (very high positive correlation) the advantages from being the proposer in the negotiated decision making problem increase so that the negotiated solution moves closer to the frontier, as illustrated in Figure 2.5. The figure shows two scenarios, one with high positive correlation between the ecosystem service benefits and one with low positive correlation, for each scenarios two solutions are plotted using the results from one round of negotiation, one in which buyer A is the proposer and one in which buyer B is the proposer.

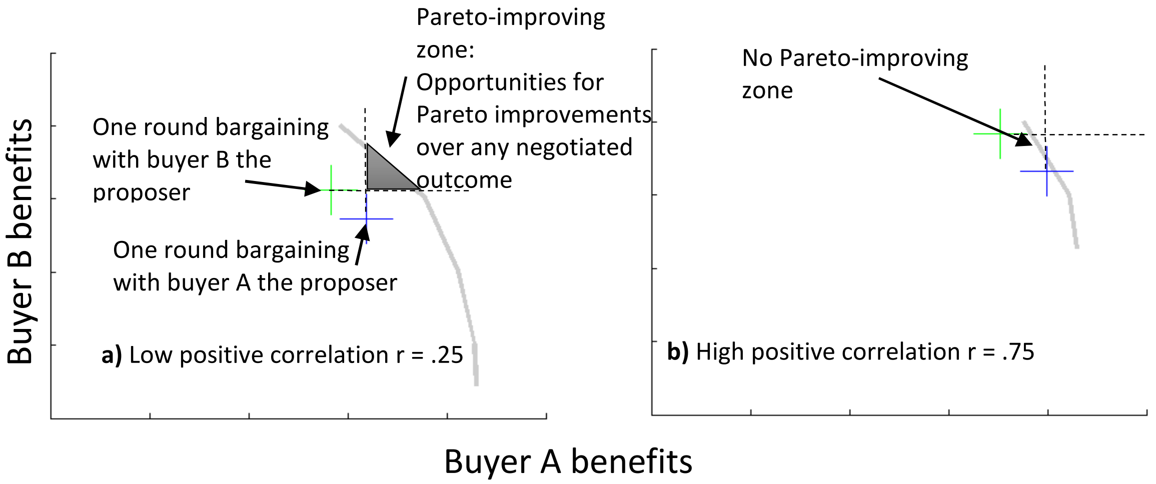

Figure 2.5. One round negotiated solutions showing Pareto-improving zone in a) with low positive correlation but not in b) with high positive correlation.

Figure 2.5: The benefits to each buyer are correlated random values simulated using the process outlined in the setup.

As well as comparing the negotiated solutions to the cooperative solutions defining the frontier, we can compare the extreme (one round of negotiation) outcomes with each other; in other words, compare the solution when buyer A makes the only proposal in the negotiation to the solution when buyer B makes that proposal.

Figure 2.6 plots, for five different randomly simulated environments, the Euclidean distance between the one round negotiated solutions (one in which buyer A is the proposer and one in which buyer B is the proposer) as the correlation coefficient is varied between -1 and 1. The results show that, for the negotiated solution, advantages from being the proposer can be observed above a correlation coefficient of approximately -0.4, peaking at moderately high levels of positive correlation. In other words, the first proposer advantage is largest when both the buyers prefer similar land parcels. To understand why this is so, imagine a situation in which both buyers get lots of benefit from a particular land parcel, 'land parcel *z*' Now, if buyer A is the proposer, they know that buyer B also values 'land parcel *z*' highly and will be prepared to pay for changes in land-management practices at that location. Accordingly, buyer A's best strategy is to allow buyer B to pay for 'land parcel *z*' whilst focusing their buying efforts on other valuable land parcels that buyer B will not want to fund. This gives an advantage to buyer A because buyer B will pay for the land-management change on 'land parcel *z*' that buyer A would have been prepared to fund. Conversely, if buyer B is the proposer, then they get an advantage.

Figure 2.6 also shows a decline in the first proposer advantage when the correlation coefficient is very highly positive, this decline occurs because the buyer proposing first is likely to replace 'land parcel *z*' with another land parcel that is beneficial to the second buyer because of the similarity in the buyers preferences.

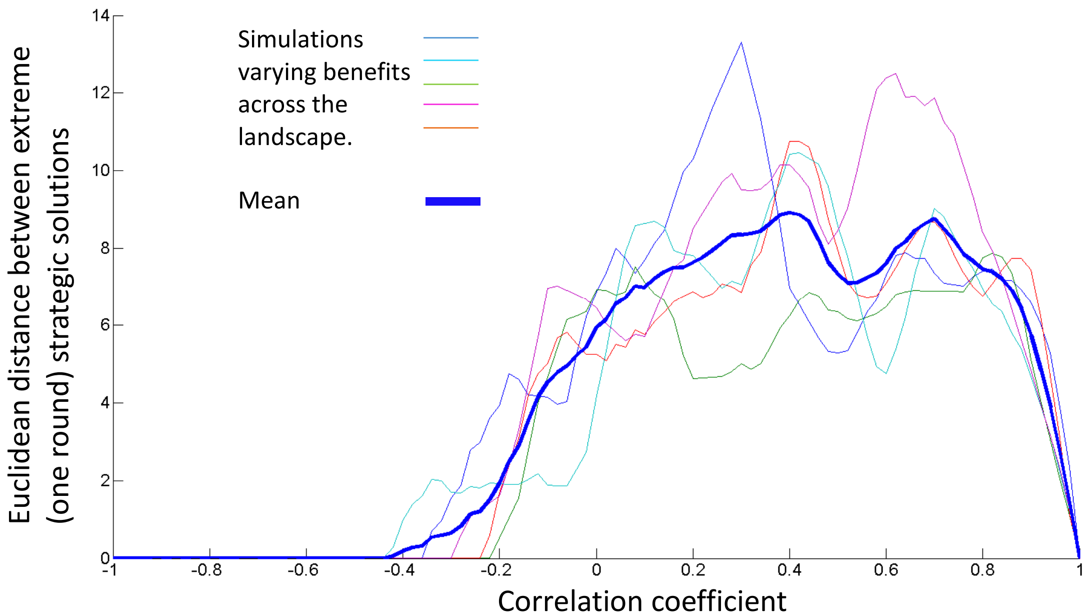

Figure 2.6. First mover advantage – measured as the Euclidean distance between the one round negotiated solutions (solution one is buyer A as the leader and solution two is buyer B as the leader).

Cooperative decision making

Figure 2.7 builds on information in Figure 2.4 on the production possibilities frontier, specifically it shows the production possibilities frontier for ecosystem service benefits as the correlation coefficient (*r*) is varied between 1 and -1.

The first thing to notice is that the frontiers do not reach the axes, this results from modelling all the benefits as strictly positive, so even using all the fixed resources to maximise buyer A's benefits leads to buyer B receiving some positive benefits. Figure 2.7 also shows differences in the shape of the frontier as the correlation is varied. When the benefits are perfectly negatively correlated (*r*= -1) the rate at which buyer A's benefits have to be given up for buyer B's benefits is constant and we see a straight line on Figure 2.7. The shape of that production possibilities frontier follows from the assumptions of our simulated environment. Specifically, as the cost of each land parcel is uniform across the landscape, the only important factor for the buyers in choosing where to pay for land-management changes is the benefits. For example, imagine the best we could do for buyer B in a cooperative institution is to use the combined budget to purchase the top *p* parcels on their ordered list, which as a result of perfect negative correlation will also constitute the worst *p* parcels on buyer A's list. To construct the production possibilities frontier, we would then wish to increment buyer A's benefit while maximising the benefits realised by buyer B. Given each land parcel can be funded at constant cost, that would mean choosing the land parcel one position up buyer A's ordered list, which (because of perfect negative correlation) is also one position down buyer B's ordered list. It follows that each time we perform this calculation we add a constant to buyer A's benefit and take a constant decrement from buyer B's benefit. The production possibilities frontier must, therefore, fall at a constant rate as we trade-off buyer B's benefit to get more for buyer A. At the other extreme with perfect positive correlation (*r*= 1) we only a single point. With perfect positive correlation the two buyers have exactly the same ranking of the land parcels. It follows that the same land parcels will be purchased regardless of how the combined budget is divided between the two buyers.

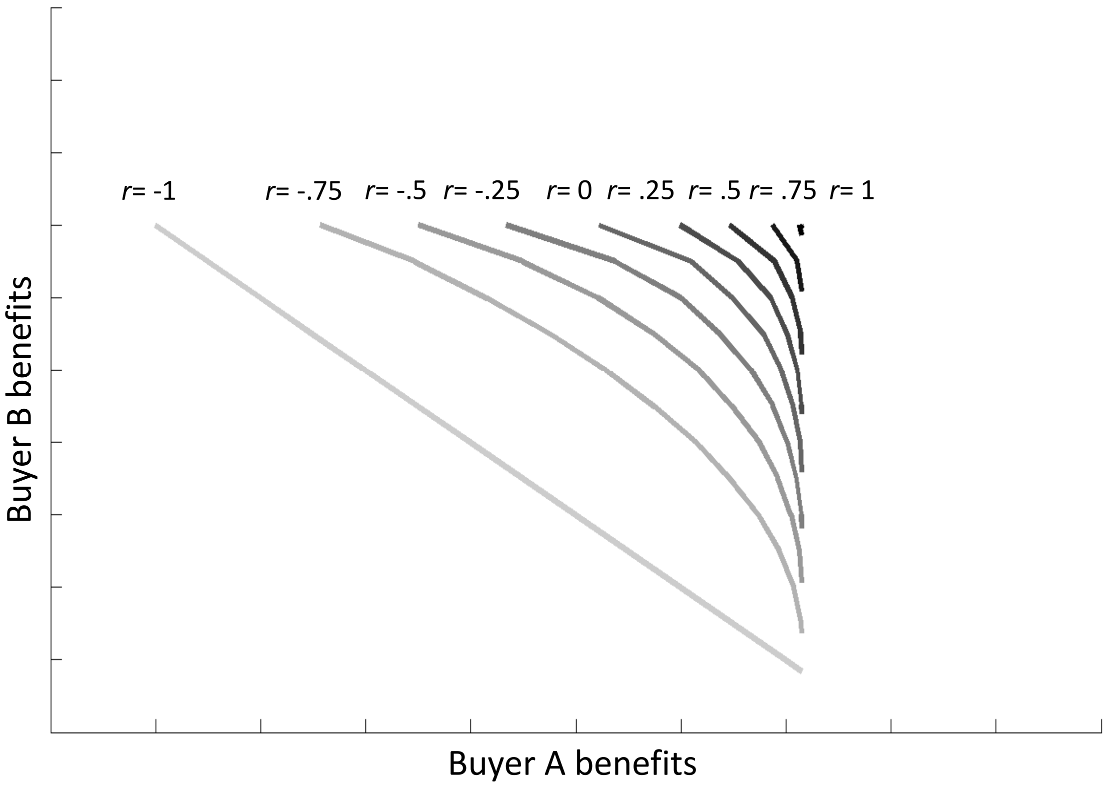

Figure 2.7. Joint production possibilities frontier with varying correlation in the ecosystem service benefits.

Figure 2.7: The benefits to each buyer are correlated random values simulated using the process outlined in the setup.

As the correlation coefficient increases from r = -1 all the way up to r = 1 we see a pattern of the frontier shifting outwards. Such a pattern exists because a higher correlation coefficient means that an increasing number of land parcels are highly beneficial to both buyers so that it is possible to achieve Pareto-improving joint production.

Comparisons of different PES purchasing institutions

For the final analysis in simulation 1 we look at the difference in the solutions across different PES-purchasing institutions; in particular, independent and simultaneous decision making, negotiated decision making and cooperative decision making. Due to the similarities in the sequential and negotiated (one round) solutions we do not include sequential decision making in this comparison. As a measure of difference between the solutions from the decision making problems we take the Euclidean distance in benefits. Due to cooperative and negotiated decision making having multiple solutions, the minimum Euclidean distance is used[29], as such the differences between the solution methods can be viewed as conservative. The results are presented in Figure 2.8 (a, b, c) which shows results from a series of five simulations in which the distribution of benefits across the landscape is randomly varied. The darker lines represent the mean value of the five simulations.

a. Cooperative versus negotiation decision making.

Figure 2.8a shows the minimum Euclidean distance between the cooperative solutions that define the production possibilities frontier and the negotiated solutions. At the two extremes in correlation, the negotiated solution coincides with the frontier such that there is no difference between the two. The solutions coincide because at perfect negative correlation the land parcels that offer the highest benefits for buyer A offer the lowest benefit for buyer B, as such the cooperative solution determines how much of the combined budget goes towards paying for land-management changes on land parcels offering high levels of benefits to buyer A and those parcels offering high levels of benefits to buyer B. The solutions coincide at perfect positive correlation because the same land parcels offer high levels of benefit to both buyers, as such the exact same land parcels would be paid for when the buyers act cooperatively as would be if the buyers were acting in their own self-interest.

If the negotiated solution is on the frontier, negotiated decision making is Pareto-efficient. At all other correlation coefficient values we see opportunities for some form of cooperative mechanism to do better than non-cooperative negotiated decision making. Observe that the largest Euclidean distance between the negotiated solution and the frontier tends to be larger when benefits are moderately negatively correlated and decreases as the correlation coefficient increases. At negative correlation coefficients, the land parcels offering high benefits to one buyer are likely to provide low benefits to the other buyer. As such, the buyers thinking in a non-cooperative

[29] For the strategic solution, we calculate only the two extreme one round negotiation solutions and calculate the minimum Euclidean distance from the two solutions. Although it is possible to calculate the whole range of strategic solutions for each correlation coefficient value the computational time would be substantial due to the nested integer programme and genetic algorithm running time.

strategic way optimise by choosing land parcels that are best for them, but these provide little benefits to the other buyer. In the cooperative solutions, the optimal land parcels are those that provide the highest total benefit, as such land parcels which provide a medium level of benefits to both buyers may be Pareto-improving. It follows that cooperative PES purchasing institutions are likely to offer greater efficiency gains when benefits are less positively correlated.

b. Negotiated versus independent and simultaneous decision making.

Figure 2.8b shows the Euclidean distance between the negotiated solutions and the independent and simultaneous solution. At highly negative correlation coefficients the two buyers have preferences for different land parcels, as the coefficient is increased the number of land parcels that the two buyers would both choose increases and so in the independent and simultaneous solution there is a high level of stacking or double payments. Subsequently, at higher levels of positive correlation there is an advantage in terms of efficiency for the purchasers, of institutions which allow buyers to think strategically in making purchasing decisions.

Cooperative versus independent and simultaneous decision making.

Figure 2.8c shows the difference between the cooperative solutions and the independent and simultaneous solution. The solutions are identical at perfect negative correlation because, as stated previously, there are no opportunities for gains from cooperation. As the correlation coefficient is increased the Euclidean distance between the solutions increases up to a peak at perfect positive correlation. The Euclidean distance increases at negative correlation coefficients for the same reasons as outlined in a.), furthermore at positive coefficients the Euclidean distance increases due to the double payments reasons outlined in b.).

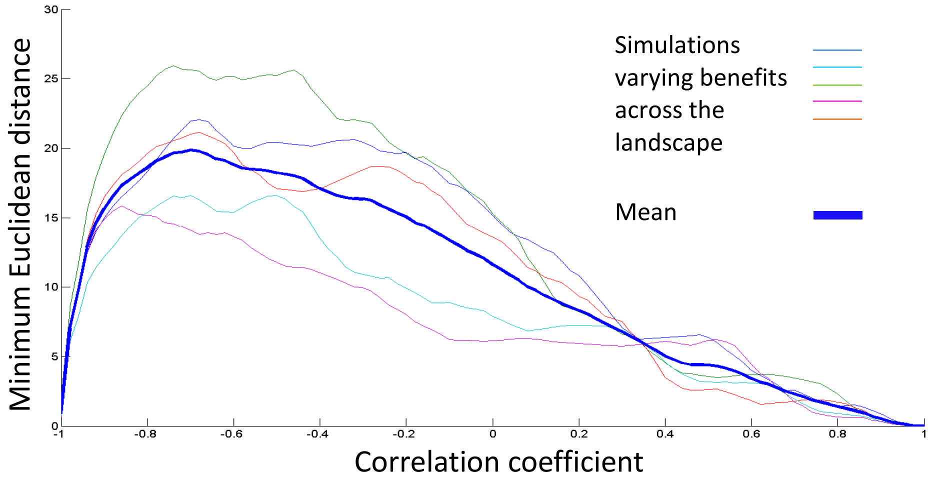

a) Differences between the **Cooperative** and the **Negotiated** solutions.

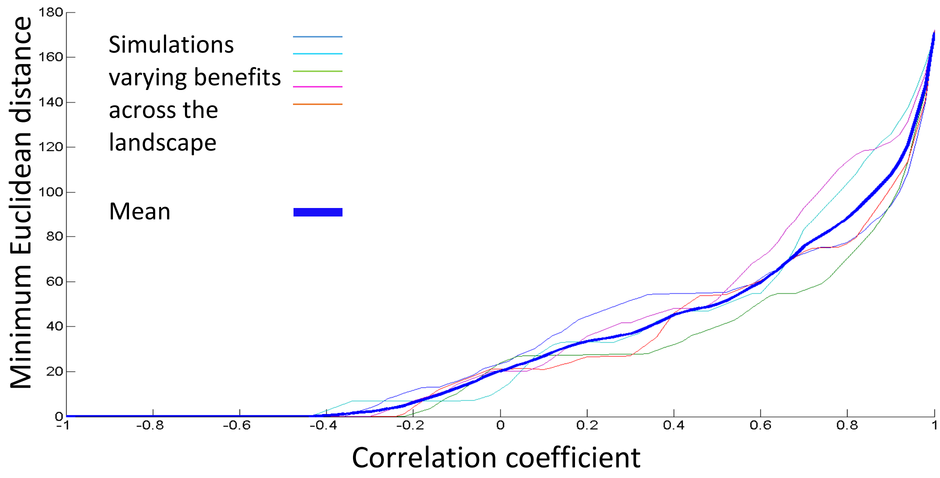

b) Differences between the **Negotiated** solutions and the **Independent and simultaneous** solution.

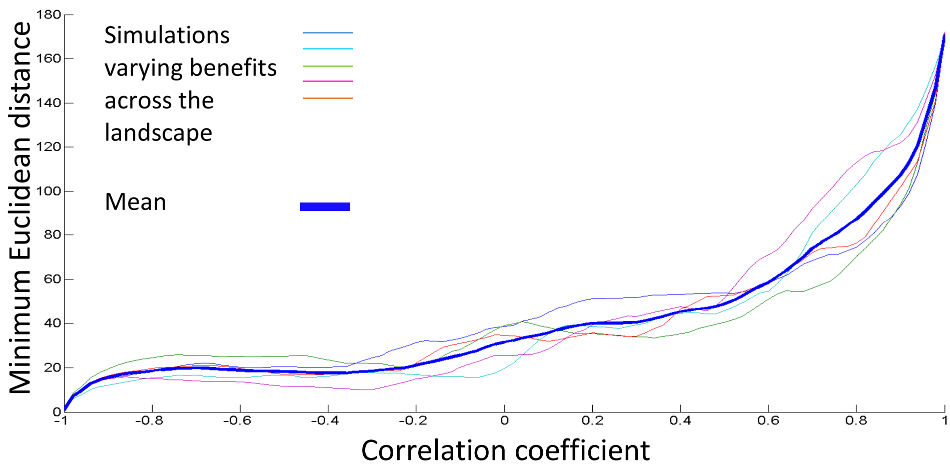

c.) Differences between the **Cooperative** and the **Independent and simultaneous** Solution.

Figure 2.8. Comparisons of different PES purchasing institutions using Euclidean distances.

2.6. Simulation 2

In this section, we introduce a more complex simulation environment which allows us to illustrate how the framework outlined in this chapter can be used by policy makers to assess the potential for multiple purchaser markets in a more realistic situation. The framework set out in this chapter is very general in that it can incorporate different buyers' objectives, for example objectives for different ecosystem service benefits, and include different constraints on those objectives; in addition, the framework solves a variety of PES purchasing institutions and does this over a variety of spatial landscapes. Here we present one specific simulation environment in which we found significant possibilities for Pareto-improvement through cooperative institutions.

In this simulation environment, we again limit ourselves to the two-buyer case. One of those buyers resembles the buyers described in simulation 1 insomuch as their benefits depend on the spatial heterogeneity of benefits across the landscape; each location offers a fixed benefit but those benefits differ across the landscape. In contrast, the second buyer's benefit flows depend on spatial interdependency, specifically the connectedness of the land parcels brought under alternative land-management practices.

Let us assume once again that the change in land-management practice involves taking a land parcel out of production (or at least moving to low intensity agriculture). Moreover, we can imagine that the buyer whose benefits depend just on the location of land parcels to be a water company. We assume that the value of taking a parcel out of production to the water company depends only on the proximity of that land parcel to water courses from which the company abstracts for water supply. Likewise, we could imagine the buyer whose benefits depend on the connectedness of land parcels to be the government interested in paying for land-management changes that would improve biodiversity. By introducing spatial interdependence into benefits we necessarily create a non-linear decision problem for the buyer. We show how our framework of methods is capable of creating solutions even for spatially interdependent benefits by forming a linearised version of the buyer's decision problem.

Setup

Consider the same size landscape as simulation 1, 225 square land parcels arranged on a 15x15 grid. This time however, that landscape constitutes a single catchment with a river system that flows through that catchment from north to south. Figure 2.9 shows the landscape with cell through which the river flows depicted in blue and other wholly agricultural land cells depicted in red.

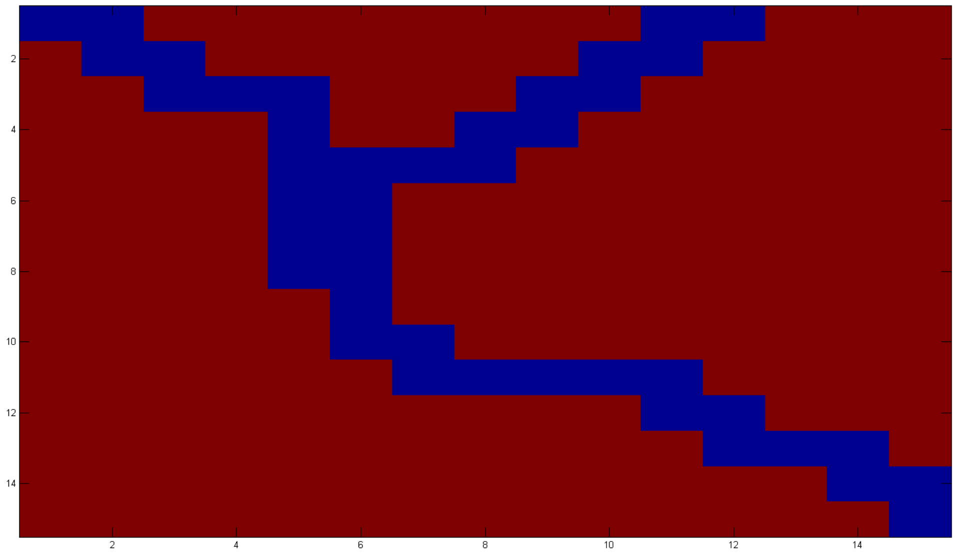

Figure 2.9. An example water catchment landscape and river, partitioned into square land parcels.

In each land parcel j, including those that contain the river, some form of agricultural production is taking place. The agricultural production leads to an initial level of pollution entering the water system, in addition the farmland supports an initial level of biodiversity. Both the water company and the government are interested in improving on these initial levels of ecosystem services (reducing water pollution for the water company and increasing biodiversity for the government). Again we assume that cost of taking a land parcel out of agricultural production is a constant, c.

Water quality buyer

As in simulation 1, equation 2.14 shows the optimisation problem for the water company with b_j again representing the benefit value at land parcel j; however this time b_j is formed by taking the Euclidean distances to the river and adding Gaussian noise, with mean 0 and standard deviation σ. The resulting b_j are shown in Figure 2.10

in which darker cells represent those cells with higher benefits to the water company. The more valuable cells are clustered around the river, with cells further from the river decreasing in value to the water company.

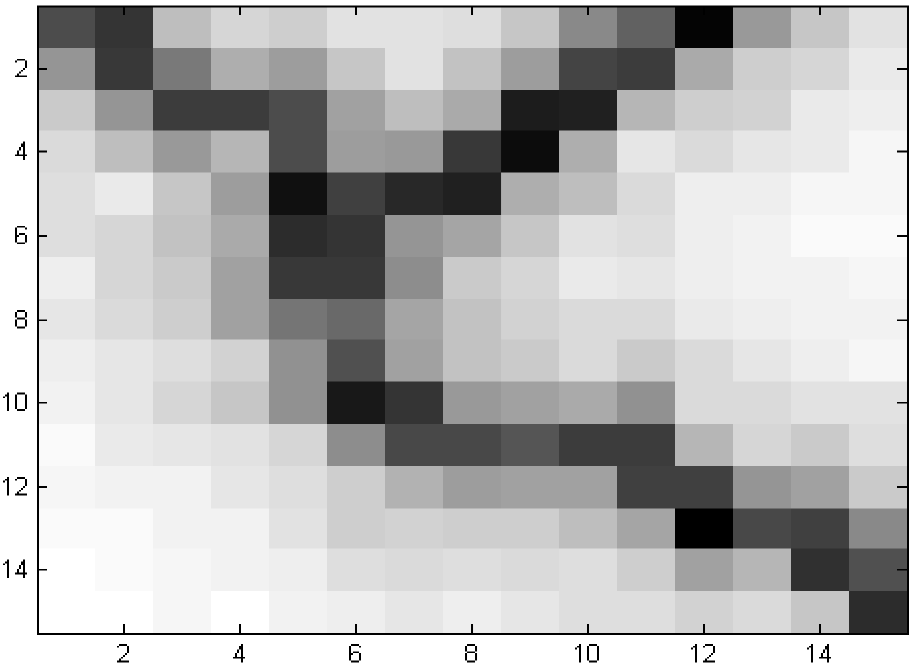

Figure 2.10. Ecosystem service benefits to the absolute spatial configuration buyer.

Biodiversity buyer

To represent the biodiversity buyer we imagine an objective of creating large / well-connected reserves, that is, contiguous land parcels taken out of agricultural production. A suitable method, taken straight from the reserve selection literature, is to choose so as to minimise the sum of the distances between neighbouring selected land parcels (Önal and Briers 2005; Williams, et al. 2005):

$$\min_{x} \quad \sum_{j\in J}\sum_{k>j} d_{jk}x_j x_k \tag{2.15}$$

$$\sum_{j=1}^{N} c x_j \quad \leq \bar{B} \tag{2.15a}$$

where d_{jk} is the pairwise Euclidean distance between land parcel j and land parcel k. The constraint ensures that the budget limit is enforced[30]. However, the objective function (2.15) is quadratic such that it cannot be simply included in the linear programming framework of methods outlined in this chapter. To enable this objective function to be included in our framework we convert the problem to a linear problem by introducing two new constraints as shown in the following:

$$\min_{x} \quad \sum_{j \in J} \sum_{k > j} d_{jk} u_{jk} \tag{2.16}$$

$$\begin{aligned} s.t.\, u_{jk} &\leq x_j \qquad \forall j \in J, \forall k \in D_j, k > j \\ u_{jk} &\leq x_k \end{aligned} \tag{2.16a}$$

where u_{jk} = 0 or 1; it is 1 if land parcels j and k are both selected and 0 otherwise. The two constraints (2.16a) ensure the definition of the binary u_{jk} variables, for example, if u_{jk} is 1 then both x_j and x_k must be greater than or equal to 1, since they are also binary variables they both have to equal 1. Finally, we take the inverse of the distances d_{jk}, and maximise, so we are maximising the sum of the inverse pairwise Euclidean distances in our integer linear programme for the biodiversity buyer:

$$Max \sum_{j \in J} \sum_{k > j} \left(\frac{1}{d_{jk}} \right) u_{jk} \tag{2.17}$$

$$\begin{aligned} s.t. \qquad u_{jk} &\leq x_j \qquad \forall j \in J, \forall k \in D_j, k > j \\ u_{jk} &\leq x_k \\ \sum_{j} c x_j &\leq \bar{B} \end{aligned} \tag{2.17a}$$

Now, since both the problem of the water quality buyer and the biodiversity buyer can be expressed in a linear programming model they can be incorporated into the framework of methods outlined earlier in the chapter.

[30] A species selection constraint can also be added in the form: $\sum_j \delta_{sj} x_j \geq 1$, where δ_{sj} is a parameter in which δ_{sj}= 1 if parcel j contains species s, and δ_{sj}= 0 otherwise. A species selection constraint ensures that at least one land parcel containing each species is selected and is a common method used in the conservation biology literature.

Results: Water quality buyer and a biodiversity buyer.

To visualise the results we present the simulation landscape in Figure 2.11, highlighting the specific land parcels purchased by the two independent and simultaneous buyers and one example from both the range of negotiated and cooperative solutions. In the figure the land parcels paid for by the water quality buyer are depicted in blue, those paid for by the government biodiversity buyer are depicted in green, those in which both buyers paid for the same land parcels are depicted in yellow (these are stacked payments and are only applicable to the independent and simultaneous PES purchasing institution) and those in which the cooperative buyer acting on behalf of both buyers paid are depicted in white (applicable only to the cooperative institution). For the negotiated solution, we show an example from an institution with one-round of negotiation in which the water quality buyer was the proposer. For the cooperative solution we show an example which lies in the Pareto-improving zone, as a reminder, that defines the feasible solutions that provide Pareto-improvements for both buyers over all the negotiated outcomes.

Figure 2.11a illustrates the independent and simultaneous solution landscape when each buyer acts independently. In this solution, because the buyers have not considered the existence of other buyers, it leads to a suboptimal configuration of land parcels being taken out of production. Observe in figure 2.11a the water quality buyer pays for land parcels along the river to be taken out of production as those parcels lead to the highest benefit, the biodiversity buyer however, pays for land parcels that creates the largest contiguous collection of sites. The biodiversity buyer does not take advantage of the land parcels paid for by water quality buyer; for example, the biodiversity buyer could have used the concentration of land parcels funded by the water quality buyer in the north west of the landscape to create a larger contiguous area of land taken out of the agricultural production than they could thinking independently about their PES purchasing choices. It is this lack of consideration for the other buyer that leads to the suboptimality of the independent simultaneous solution. Furthermore, the independent solution is suboptimal in this case despite not having any stacked payments for a single land parcel; in scenarios in which stacked payments exist, the independent and simultaneous solution would be even less desirable.

Figure 2.11b illustrates one particular solution arising from the negotiation between the buyers. The major difference between the independent and simultaneous solution and the negotiated solution can be seen in the location of land parcels paid for by the biodiversity buyer. Indeed, the figure shows the biodiversity buyer has changed the locations in which they purchase land parcels to take advantage of the configuration of the water company's purchases. Nevertheless, the water quality buyer also makes strategic decisions, and these decisions vary depending on whether they are the proposer or not in negotiated PES purchasing institution. When the water quality buyer is not the proposer it simply chooses the best land parcels in terms of water quality (from the remaining land parcels after the government has chosen). Alternatively, when the water company is the proposer it chooses in such a way as to influence the government into buying the land parcels that are beneficial to the water company[31]. The solution shown in the figure has more contiguous land parcels in the PES scheme and those land parcels are located close to the river system. The contiguity of the land parcels taken out of agricultural production increasing the benefits for the biodiversity buyer, the location of those parcels also has the effect of increasing benefits for the water quality buyer because more land parcels close to the river are converted to an alternative land-management practice. We see the benefits for the water quality buyer increase from 309.6 to 393.2, an increase of 27%, and the benefits for the biodiversity buyer increase from 54 to 61, an increase of 13%, compared to the independent and simultaneous institution. Hence, we conclude that, for our simulated environment, negotiated decision making Pareto-dominates independent and simultaneous decision making.

Figure 2.11c illustrates a cooperative solution sitting on the production possibilities frontier. The pattern of land parcels in which the land is taken out of agricultural production is different to the patterns seen in the negotiated and independent and simultaneous institutions. Since the land parcels are chosen to jointly maximise the benefits for the water quality buyer and the biodiversity buyer, those land parcels which simultaneously provide connected land parcels and land parcels close to the river system are favoured. The example solution presented in the figure shows the benefits for the water quality buyer increase from 393.2 to 420.0, an increase of 7%, and the benefits for the biodiversity buyer increase from 61 to 64, an increase of 5%

[31] Only if, that strategy provides more benefits overall to the water company

over the negotiated solution. As such, the solution is a Pareto-improvement on the non-cooperative solutions and is achieved by considering both buyers' objectives together. Hence, cooperative decision making, such as when the two buyers give up responsibility for their choices to a trusted broker working to their mutual best advantage, Pareto-dominates both independent and simultaneous and negotiated non-cooperative decision making. It is interesting to contrast these results with the results from the international environmental agreement literature, for example Barrett (1994) conclude that cooperative self-enforcing international environmental agreements may not be able to substantially improve on the non-cooperative outcome, particularly when the number of parties is large. All the results presented here are for just two ecosystem service buyers, it would be interesting in the future to expand our analysis to more than two buyers and see if the Pareto dominance of cooperative decision making is still present or if we see results similar to the IEA literature between the cooperative and non-cooperative outcomes.

a. Non-cooperative independent and simultaneous	
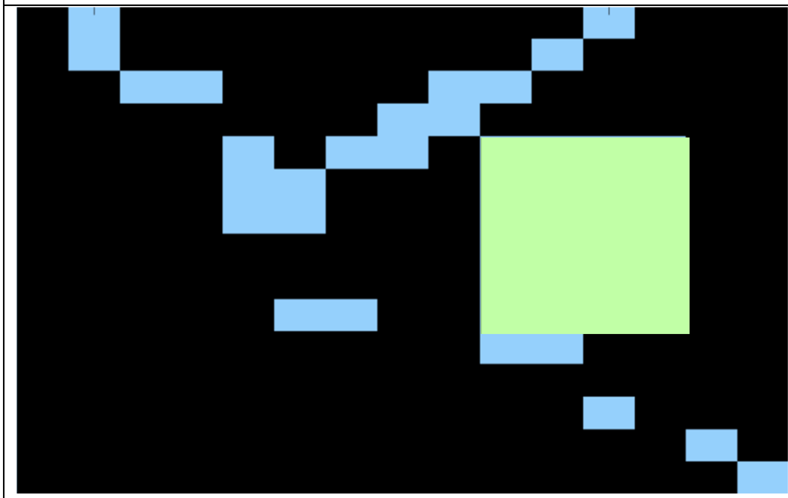	**Benefits** Water company = 309.6 Government = 54 (biodiversity) Water buyer Government buyer Both buyers
b. Non-cooperative negotiated solution	
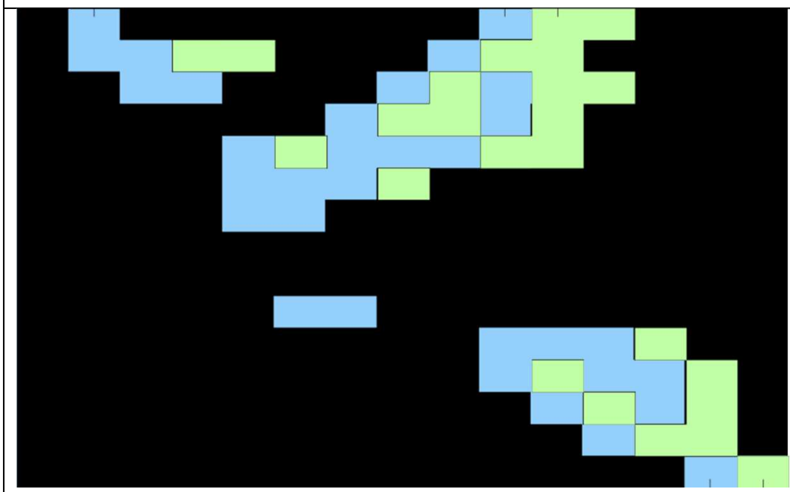	**Benefits** Water company = 393.2 Government = 61 (biodiversity) Water buyer Government buyer
c. Cooperative solution on the production possibilities frontier	
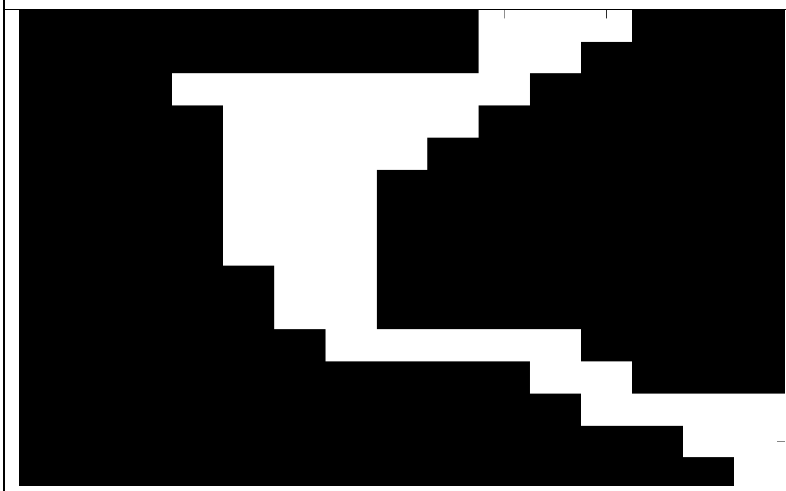	**Benefits** Water company = 420.0 Government = 64 (biodiversity) Cooperative buyer

Figure 2.11. Example naïve, strategic and cooperative simulation solutions to the problem of simultaneous purchasing from multiple buyers of ecosystem services.

2.7. Conclusion and discussion[32]

In this chapter, we have developed a powerful and flexible method for exploring the outcome of different purchasing institutions in a PES scheme with multiple independent buyers. Moreover, the analytical framework we describe can be used to identify optimal multi-objective patterns of land use. The general framework, which we have illustrated in two examples, allows important characteristics of PES to be incorporated, such as different costs and benefits from a land-management change, different buyer objectives, and importantly the spatial characteristics and dependency of the ecosystem services.

We have highlighted four multiple-purchaser decision making problems: independent simultaneous decision making, independent and sequential decision making, strategic negotiated decision making, in which buyers consider how the other buyer will react to their spatial choices of where to purchase ecosystem services, and fully cooperative decision making, in which, both buyers benefits from the ecosystem services are optimised jointly by a third party. These decision making arrangements are not exhaustive of all potential arrangements between multiple buyers but instead show a range of decision making problems which increase in strategic sophistication and the level of cooperation between the buyers. Our investigation shows that negotiated solutions (of which there are many) Pareto-dominate the independent and simultaneous solution, suggesting that, as a minimum, institutions should be created that coordinate and facilitate negotiation between ecosystem services purchasers in a particular landscape. Moreover, for many problems there exist a set of cooperative solutions that Pareto-dominate all negotiated solutions suggesting that coordinating action through empowering a trusted broker to make decisions on behalf of both buyers could potentially benefit both buyers.

We create two simulation environments in which to test our framework. In the first simulation we investigated the effect that correlation in the production of ecosystem services has on the efficiency for the multiple buyers using our four decision making problems. The results show that negotiated decision making and cooperative decision making provides Pareto-improvements over independent and simultaneous decision

[32] Further concluding remarks on all three chapters, in which we highlight potential future extensions, can be found at the end of this thesis.

making in the majority of scenarios. For each individual negotiated solution, Pareto-improvements are possible by employing a cooperative institution for all correlation values apart from perfect negative or perfect positive correlation, in which they do equally well to the cooperative solution. For negotiated decision making taken as a whole, we find that a Pareto-improving zone exists for all positive correlation values less than approximately r = 0.65 (high positive correlation) and for all negative correlation values apart from perfect negative correlation. The difference between the cooperative and negotiated solutions is highest when the ecosystem services are moderately negatively correlated, in other words when the buyers tend to favour different locations. It follows that cooperative PES purchasing institutions are likely to offer greater efficiency gains when benefits are less positively correlated. For the ecosystem services we consider in this chapter, such as biodiversity and water quality, the evidence suggests weak positive correlations are fairly likely. Maes et al. (2012) show a link between habitats in favourable conservation status and regulating ecosystem services such as water quality and Chan et al. (2006) provide evidence that such correlations between ecosystem services are likely to be weak, their results show correlation coefficient of less than ±0.3 for all 21 pairs they assess apart from carbon storage and water storage which has a correlation coefficient of 0.58. Therefore PES purchasing institutions for the type of ecosystem services we consider here are likely to offer moderate gains.

Our second simulation environment reflects a more complex and perhaps realistic situation. A landscape of agricultural land parcels and a river was created and two buyers imagined, one (a water quality buyer) whose benefits relied on the spatial heterogeneity of the landscape—sites closer to the river were more beneficial to a water quality buyer—and another (a biodiversity buyer) whose benefits relied on the spatial interdependency and configuration of the landscape—connected habitats provide more benefits to the biodiversity buyer. By introducing spatial interdependence into either the costs or benefits we necessarily create a non-linear decision problem for the buyer. We show how a non-linear spatially interdependent problem can be linearised and solved within our framework. Indeed, any buyer's decision problem can be included as long as it can be represented in a linear way. Policy makers are thus able to use the method outlined in this chapter to study specific landscape configurations solutions. Moreover, the framework of methods can be used by a third party broker to generate real solutions—exactly which sites to purchase and who should purchase them—that Pareto-dominate any solutions that could be negotiated by multiple purchasers thinking in their own self-interest.

In this chapter we assume a number of simplifications to provide clear results but real world complications and heterogeneity can straightforwardly be included in the framework of methods. The framework we propose is capable of incorporating real world data where such data exists, for example, on the potential costs and benefits of providing ecosystem services. In addition, specifics of the spatial landscape can also be included. Throughout this chapter we have assumed that the land parcels are exactly the same size and shape and that they can be purchased as individual parcels, in reality there may be large land parcels, small land parcels and also farmers not interested in participating in the scheme. Although these issues add complexity, there are no theoretical barriers that mean these nuances could not be incorporated into expanded versions of the methods outlined in this chapter. We have also arranged all land parcels in square grids, this is clearly unrealistic, and different arrangements of the landscape could affect the outcomes. The use of network games in the laboratory experiments have shown that the shape of the network matters (Cassar 2007; Banerjee et al. 2012; Banerjee et al. 2014), this result is relevant for spatially interdependent benefits where the shape of the landscape determines which land parcels share borders with others and could be further explored in future research. It is also important to remember our analysis is based on there being a single land-management change that the landowner can undertake which leads to multiple ecosystem benefits. Although it may be common for an action taken by land owners to lead to multiple ecosystem services we acknowledge that our results do not necessarily apply to schemes in which a particular action leads to a specific ecosystem benefits or where farmers can claim property rights over the multiple benefit flows.

This chapter shows that Pareto-improving cooperative solutions exist but whether the multiple purchasers will in reality agree to such a solution depends on the mechanism put in place, the negotiation process between the two purchasers, and the negotiation between the purchasers and the landowners. Insights into these matters are not possible within the methods and simulation environment used in this chapter but are explored in the complementary experiment on multiple purchasers in the next chapter.

CHAPTER 3

AN EXPERIMENTAL EXPLORATION OF MULTIPLE-PURCHASER PAYMENT FOR ECOSYSTEM SERVICES SCHEMES

3.1. Introduction

In this chapter, we build on the work of the previous chapter on multiple ecosystem services buyers. In the previous chapter, we modelled the buyers' procurement decisions in a range of non-cooperative and cooperative decision settings. We concluded that when buyers were modelled as strategic negotiators, solutions Pareto-dominate outcomes where buyers were modelled as independent and uninformed decision makers. While we were able to demonstrate Pareto-improving solutions arising from negotiations in the simplified setting of a simulation environment, it does not necessarily follow that the same outcomes will be achieved in the far more complex context of real negotiations. To gather insights as to whether such solutions might be achievable in a real PES scheme and to explore the factors shaping the division of gains from negotiations between multiple purchasers and sellers in such a scheme, in this chapter, we turn to experimental economic techniques.

Negotiation is not the only exchange mechanism that might be used to facilitate transactions and establish prices in a PES scheme. Alternatives include fixed prices and competitive bidding. The choice of mechanism for PES schemes will depend on the specific circumstances of any particular scheme. The procurement literature, for example, has shown that negotiation might be preferred to competitive bidding when the good(s) being procured are complex and there are few available bidders (Bajari and Tadelis 2001; Bajari, et al. 2009). In addition, both fixed price mechanisms and auctions may stifle communication between buyers and sellers, preventing the valuable exchange of knowledge and expertise (Goldberg 1977; Bajari, et al. 2009). Since the exchanges transacted in PES schemes are often complex (for example, stipulating the timing, locations and types of undertaking expected of a farmer) and sellers of ecosystem services are generally assumed to have superior information (for example, concerning the complexities surrounding the translation of land-management changes into ecosystem service benefits) negotiation between buyers and sellers may also play an important role in certain PES mechanisms.

Accordingly, our experimental investigation focuses on an exchange process facilitated through the *multilateral bargaining* of buyers and sellers. The laboratory environment allows for careful control over the information given to the participants but also allows for financial decisions to be made by human decision makers rather than simple programming of a computer to simulate rational decision making. For the purposes of clarity, our experiments involve just three parties to those negotiations: two buyers and one seller. By including a seller in the negotiation process buyers not only have to negotiate between themselves regarding how much each might contribute to the payment offered to the seller but they must also ensure that that offer is satisfactory for the seller. The experiment is conducted over a maximum of 15 rounds of negotiation, although each time a participant rejects an offer there is an increasing risk (presented clearly to the participants) that negotiations will fail and therefore no deal will be agreed.

We use this experimental framework to investigate a number of complexities of the negotiating environment that might typically arise in a PES scheme. First, the degree to which the buyers offer (and the seller accepts) an amount over and above the sellers costs. Second, the degree to which *asymmetry in the gains* enjoyed by the two buyers from a successful transaction affects the outcome of negotiations. Here we imagine that one buyer would benefit more from a PES scheme being implemented. Third, the degree to which *asymmetry in the income* of the two buyers (irrespective of their gains from the transaction) affects the outcome of negotiations. Here we imagine that one buyer might be a large, wealthy organisation and that the relatively less wealthy buyer might be more inclined to free ride on the wealthy buyer's contribution to the PES scheme. Fourth, how negotiations differ under conditions of *incomplete information*. Here we imagine that differences in knowledge exist between the buyers and sellers, for example the seller may know the costs for supplying the environmental output but the buyers might not. Finally, how negotiations evolve when the benefits enjoyed by the buyers from the transaction are not known for sure but are *stochastic* in nature. Here we imagine that the buyers are paying for the seller to undertake an action and therefore they are not entirely certain of the actual environmental output that will be produced, this could be due to unpredictable phenomenon such as weather patterns. This is a common situation for PES schemes but has received relatively little attention in the experimental economics literature. This chapter explores these five issues within our experimental framework to not only establish whether participants can successfully negotiate multilateral agreements in such a purchasing setting but also to explore how the gains from successfully-negotiated exchanges are partitioned both between the purchasers and between the purchasers and sellers.

Previous experimental studies assessing bargaining in the purchase of ecosystem services are limited in number and have typically involved just a single buyer and seller (Bruce and Clark 2010a,b) or a single buyer and multiple sellers as in the auction literature (Reeson et al. 2011) and agglomeration bonus literature (Banerjee et al. 2014). In this chapter, we move beyond bilateral negotiated ecosystem service procurement schemes which have been successfully implemented both in laboratory experiments (Bruce and Clark 2010b,2012) and in real world schemes, Perrier-Vittel (Perrot-Maître 2006) and United Utilities UK (Smith, et al. 2013) to study multi-lateral ecosystem service bargaining—with two buyers and one seller of ecosystem services. To do that, we use the non-cooperative, alternating-bargaining setup discussed in Section 3.2. A general discussion on non-cooperative bargaining is followed by specific literature reviews of the four complexities of multilateral bargaining identified above. In Section 3.3, we describe our experimental design and in Section 3.4 report the results of their implementation. Finally in Section 3.5 we offer concluding remarks.

3.2. Bargaining: an ecosystem services procurement mechanism

Studying negotiation necessarily involves studying the dynamics between the two classes of agent, buyers of ecosystem services and sellers of ecosystem services. Bargaining is characterised by agents with common interests of cooperation, but with conflicting interests about how it is achieved and about the resulting payoffs. For the majority of people, businesses and organisations, bargaining is a commonplace activity. Numerous day to day tasks either at home or in the workplace involve bargaining, companies are frequently bargaining to get the best prices or wages, and the policies that emerge from political parties or governments are often the result of long and repeated bargaining processes.

Negotiation as an ecosystem service procurement mechanism has received little attention despite being utilised in a variety of real world PES schemes, such as the Perrier-Vittel scheme in France (Perrot-Maître 2006) and the United Utilities scheme (Smith, et al. 2013). Perrier-Vittel purchased large swathes of farm land around water springs and then offered the farmers the rights to farm on the land provided they followed management practices that caused minimum water pollution. The International Institute for Environment and Development concluded that the success of the scheme was the result of the extensive trust building with farmers and a set of mutually agreeable negotiated incentives for the farmers (Perrot-Maître 2006). United Utilities implemented The Sustainable Catchment Management Programme (SCaMP) programme between 2005 and 2010 with the aim of improving raw water quality and conditions for sites of special scientific interest. Through directly negotiating contracts with farmers a wide variety of capital items were installed to help farmers deliver additional ecosystem services beyond the minimum legal standards[33].

In the wider economics literature research on bargaining and its outcomes has been pursued in two parallel fields; non-cooperative game theory and cooperative game theory. Non-cooperative game theory focuses on the particular equilibrium outcomes that arise from some defined procedure for bargaining between multiple strategic agents (Sutton 1986). Cooperative game theory puts no structure on the bargaining procedure and instead focuses on the benefits different agents might enjoy when they act together in particular combinations (Osborne and Rubinstein 1994). Due to our interest in studying the negotiation outcomes for ecosystem services based on the structural rules imposed by a PES scheme we proceed to explore bargaining through the lens of non-cooperative game theory.

Non-cooperative game theory considers bargaining to be fully specified by the procedural rules of the negotiation process. A bargaining strategy specifies the action of an agent at each stage of the negotiation process. The outcome of a negotiation is identified as a Nash equilibrium, that is to say a set of bargaining strategies at which no agent could benefit from unilaterally changing their strategy (Osborne and Rubinstein 1994). If bargaining proceeds through a pre-defined sequence of

[33] An additional study conducted in the Fowey river catchment, UK, recommended pursuing different procurement mechanisms in different situations (Day and Couldrick 2013). They recommend an ‘Advisor-led mechanism’ (where farm advisors go out to visit and negotiate directly with farmers) for small scale schemes, where farm advisors have good local knowledge, and known farms are likely to yield positive outcomes. In contrast, auction mechanisms are recommended for large scale schemes where the buyers have little local knowledge.

negotiations then equilibria can be defined by the more specific concept of a subgame perfect Nash equilibrium.[34]

One of the most studied non-cooperative bargaining games is the ultimatum game. In that game, one of the players propose a split of some surplus and the other player has the option to either accept or reject that proposal; if the responder rejects then both players get nothing. In the ultimatum game the subgame perfect equilibrium stipulates that the proposer asks for the entire surplus and the responder accepts, however in laboratory experiments the common result is that the responders are able to get more than the game theoretical prediction of close to zero but generally less than an equal split (Roth 1995). An extension of the ultimatum game is the alternating bargaining game which adds multiple stages of offer and counter-offer. In a two round version, player 1 makes a proposal and player 2 can accept or reject. If player 2 accepts, the game ends, if they reject, then player 2 gets to make a counter-offer. With any finite number of rounds the unique subgame perfect equilibrium is one in which the player who is the proposer in the last round has all the bargaining power and can claim all the surplus. One way in which the alternating bargaining game differs from reality, however, is that it fails to recognise that bargaining is itself costly. Stahl's bargaining model (Stahl 1972) captures that complication through adding a cost to each stage of a negotiation, a structure that Rubinstein (1982) extended to a game with infinite possible rounds of negotiation. These models show that when negotiation is costly a subgame perfect equilibrium of the game is achieved in which a more equitable split of the surplus is agreed upon in the first round of negotiation. The intuition here is that Player 1 is better off proposing a fairer division of the full surplus in the first round than holding out for the possibility of a larger share of a diminished surplus later in the negotiations. Rubinstein's model was later generalised to $n \geq 3$ players by Herrero (1985) but the uniqueness of the subgame perfect equilibrium disappears, a result which is relevant here as more than two participants are likely for all PES.

The alternating bargaining game is employed in our experimental design and in the following sections we discuss existing experimental literature particularly relevant to

[34] The strategies chosen by all players are said to be a Nash equilibrium if no one could benefit from unilaterally changing their strategy, the subgame perfect equilibrium expands this concept to all subgames of an extended form game.

this experimental investigation, namely, heterogeneous purchasers, in terms of asymmetries in benefits inside and outside of a deal, incomplete information, stochastic payoffs and multilateral bargaining.

3.2.1. Heterogeneous purchasers

In a multi-purchaser PES scheme, buyers are unlikely to be homogeneous. For example, organisations may differ in terms of the size of the benefit they stand to gain from participating in the scheme. Likewise, organisations may differ in terms of their relative size as measured, say, by their wealth or the income they gain from activities outside the PES scheme. As a motivating example, imagine the difference between a small environmental NGO interested in river ecology and a large water company concerned with the quality of raw water abstracted for water supply. Both organisations would benefit from a PES reducing diffuse agricultural pollution, but the absolute size of those benefits may differ across the two organisations, likewise the size of those benefits relative to each organisations income may differ. Since such disparities may be common in multi-purchaser PES schemes, we are interested to explore how heterogeneity in purchasers impacts on the outcome of bargaining.

Such disparities have been extensively examined, under a bilateral negotiation framework in the experimental bargaining literature since (Rubinstein 1982) proposed his model, summaries can be found in (Roth 1995; Camerer 2003; Zwick and Mak 2012). This literature has been defined by experimental behavioural observations that often deviate from game theoretical predictions but sometimes come close to those predictions. Zwick and Mak (2012) categorise this debate in terms of “gaming” and “fairness”. As an example, imagine both parties know that one of the negotiators has an advantage, the negotiator with an advantage will try to exploit the advantage (gaming tendencies); conversely, the disadvantaged negotiator will exhibit fairness tendencies. They propose that the failure of theoretic predictions can be explained by three principles: firstly, the same negotiator can be both self-centred (“gamesman”), or inequity averse (“fairman”) depending on the context; secondly, the source of the bargaining advantage matters; and thirdly, when fairness has a price, the higher the price, the lower the demand for fairness.

Zwick and Mak (2012) identify characteristics of the negotiation environment to explain why game theoretic predictions work well in some cases but not in others. In particular, they propose that if advantages are gained through intrinsic characteristics (the characteristic of the negotiator) rather than extrinsic (the characteristics of the negotiation procedure) the advantages are more readily exploitable and the results are more likely to approximate the game theoretic predictions. Intrinsic characteristics are those that the negotiator brings to the table, for example their gains from a deal, or their time preferences (patience). On the other hand, extrinsic characteristics are the characteristics of the negotiation procedure, such as being the first or the last to make a proposal. Zwick and Mak (2012) also propose that demand for fairness is subject to some form of cost-benefit evaluation undertaken by the actors involved. With heterogeneous negotiators, the degree to which advantaged negotiators attempt to exploit their advantage depends on their cost and the degree to which disadvantaged negotiators attempt to achieve fairness depends on their cost, in other words, the higher their costs the less extreme their demands (Zwick and Chen 1999).

Although negotiations with heterogeneous subject characteristics have been heavily studied this has not been extended to multilateral negotiation environments relevant to multiple-purchaser PES schemes. Nevertheless, the vast experimental literature suggests that the level of gaming or fairness in the results could be influenced by the heterogeneous characteristic of negotiation, such as the size of the gains they stand to receive if a deal is agreed or the size of the payoff relating to the ‘outside option’ (the amount a player could receive if negotiations break down or they elect to exit negotiations).

3.2.2. Incomplete information

It seems likely that a real world multi-purchaser PES scheme will be characterised by information asymmetries. Compared to the purchasers, the seller of an ecosystem service will have a greater certainty of the opportunity cost they incur in providing that service and hence the minimum price they are willing to accept to participate in the scheme (Salzman 2005). In a similar manner, the exact amounts that different purchasers are willing to pay for ecosystem services and the value of their default payment is information known only to them (Ferraro 2008).

Asymmetries in information fundamentally change the negotiation dynamic. With incomplete information players can strategically manipulate their offers and responses to offers in an attempt to impart themselves with an advantage. For example a player may reject a favourable offer in order to signal high patience, a strong bargaining position or a large outside option. Ultimately, such strategizing may lead to less efficiency in bargaining outcomes (Kennan and Wilson 1993).

Bargaining games with incomplete information can be one-sided or two-sided. In one-sided incomplete information one of the players has private information but the other player's information is public. In addition to theoretical work (a summary can be found in Ausubel, Cramton and Deneckere (2002)) there is a large literature on experimental studies of incomplete information in bargaining protocols (summaries can be found in Roth (1995) and Camerer (2003). Once again the majority of studies use the ultimatum game format. For example, Croson (1996) conducted a one-sided incomplete information ultimatum game to contrast with a full information treatment, showing that varying the information given to the responder affected both the offers made and the demands made by the responder. In the treatment conducted in dollars, under incomplete information offers made were smaller, in the treatment conducted in percentage, under incomplete information demands were higher. The general conclusion from incomplete information ultimatum games is that offers are significantly lower and the responders accept lower offers (Guth, et al. 1996; Rapoport, et al. 1996; Croson, et al. 2003; Schmitt 2004). Shupp et al. (2013) explore the effect of incomplete information on multilateral bargaining using alternating bargaining. They find that incomplete information increases bargaining delay and the likelihood of failed agreements.

Experimental evidence focused on ecosystem services includes Bruce and Clark (2010b,2010a) who study bargaining between two stakeholders using an axiomatic, cooperative bargaining approach, where bargaining happens in continuous time. The experiment found that subjects were able to reach an agreement that was Pareto-efficient compared to a disagreement payoff. Bruce and Clark (2012) expand on that work to include incomplete information treatments, in which the players only know their own payoffs and not the other player's payoffs. They found that with incomplete information, Pareto-efficient deals were almost as likely as with full information. However, for agglomeration bonus schemes, where the aim is to induce adjacent landowners to coordinate for the production of ecosystem services, Banerjee et al. (2014) found that if subjects were informed about their neighbours actions then socially optimal outcomes were more likely.

3.2.3. Stochastic payoffs

In reality, the level of ecosystem service benefits enjoyed by the buyers when a seller takes a particular action may not be certain. For example, the size of benefits delivered by a scheme paying farmers to change land-management practices may be determined by unpredictable phenomenon such as weather patterns (Latacz-Lohmann and Schilizzi 2005). Indeed, natural processes may mediate the relationship between land-management change and service flow in such a way that the benefits realised by paying for the action are essentially stochastic.

While some experimental bargaining studies (Roth and Malouf 1979; Roth and Murnighan 1982) introduce uncertainty into the players' payoffs these are designed to induce risk neutrality in participants. There have been relatively few experimental studies explicitly examining the effect of stochastic payoffs on the negotiation procedure, one such example is (Pillutla and Murnighan 1996) who examine responder behaviour for an ultimatum game with an unknown pie size but a known small outside option ($0, $1, or $2), they find that offers are typically low and are frequently rejected. Furthermore, there is, some evidence that stochasticity affects cooperation in public goods games. Berger and Hershey (1994) show lower contributions from participants when payoffs are stochastic compared to when payoffs are deterministic. Dickinson (1998) introduced an element of exogenous risk where the public good may not be produced even with positive contributions and a situation where the risk decreases with the level of contribution. He shows some evidence of a reduction in contributions associated with the uncertain production of the public good. Keser and Montmarquette (2008) switch the situation around to assess where contributions reduce the risk of loss rather than the chance of gain and show evidence of a decrease in contributions with uncertain production.

3.2.4. Multilateral bargaining

The majority of theoretic and experimental evidence on bargaining concerns bilateral bargaining. When negotiation involves agreement by more than two independent agents the term multilateral bargaining is used. The literature studying multilateral bargaining is not as well developed; however, PES schemes with multiple buyers and at least one seller are examples of a multilateral bargaining situation.

Multilateral bargaining has been studied in the experimental literature using the ultimatum game (Güth and van Damme 1998; Kagel and Wolfe 2001; Schmitt, et al. 2008). In this literature there remains one proposer, one responder and an additional player who is a non-responder. Using a multilateral ultimatum games with incomplete information, Güth and van Damme (1998) show that proposers are not intrinsically motivated by fairness but instead want to appear fair to the responder.

An important element that arises with three or more players is the possibility of coalitions. Although coalitions between ecosystem service sellers and buyers are possible, we are specifically interested in studying the potential opportunities and barriers for multiple-purchaser PES schemes, we therefore look at a subset of multilateral problems which require universal consent from all parties. The problem of multilateral bargaining with universal consent has been studied in a range of situations such as biodiversity conservation, land development, international trade agreements and international environmental policy agreements. Lennox et al. (2012) study landowner's ability to holdout for higher payments for voluntary conservation agreements. They show that the holdout potential for landowners is significant and that this could have implications for which land parcels should be conservation priorities. Furthermore, a contemporary series of experimental papers focus on land acquisition and development (Cadigan, et al. 2009; Shupp, et al. 2013). Cadigan et al. (2009) study the issue of "holdout" in which one player can delay a project by rejecting agreement in the hope of receiving higher compensation, they show holdout to be a common problem in a range of bargaining institutions.

In our experiments each party has the power to delay and attempt to holdout for a higher payoff. Overall, the experimental design we use to study multiple purchaser negotiations for ecosystem services is closest to the experimental design in Shupp et al. (2013) who examine multilateral bargaining with incomplete information. We build on their work by including a number of characteristics potentially relevant to PES schemes; such as, multiple buyers, differing outside options, and stochastic payoffs to represent the uncertain realisation of ecosystem services flows from land-management changes. To the best of the author's knowledge, no other study has investigated bargaining for ecosystem services with more than a single buyer or with stochastic payoffs.

3.3. Experimental design and Hypotheses

Multilateral bargaining experiment

In our experiment, three players form a group. Two of the players assume the role of purchasers of ecosystem services, with the third player as the seller of the ecosystem services.[35,36] Neutral labels were used in the experiment, so that subjects were labelled player's 1, 2 and 3; nevertheless, for the remainder of this chapter we shall refer to the two purchasers as *Purchaser 1* and *Purchaser 2* and the seller as the *Provider* to denote we are referring to the participants roles in the experimental setup.

The vast majority of experimental investigations of bargaining involve just two players negotiating over how to split a 'pie'. Our experiments differ from that bilateral design in two important ways. First, through the inclusion of a third party in the negotiations in the form of the Provider; and, second, the bargaining problem faced by the Provider differs from that faced by the two purchasers. This particular setup was chosen to more realistically simulate potential multiple-purchaser PES negotiations as it seems unlikely that the Provider would be able to specify offer amounts for each purchaser. As such, the Provider was offered one amount—the sum of the purchasers' offers—and could either accept or reject that offer but did not have the opportunity to propose alternative offers.

In naturally occurring bargaining, players often alternate between offers and counter-offers, a structure mimicked by our experiment. Accordingly, bargaining began with Purchaser 1 submitting a proposal to Purchaser 2. Purchaser 1's proposal suggests an amount that the two purchasers should offer to the Provider, detailing the contribution to that amount that they themselves are willing to make and the contribution expected from Purchaser 2. If Purchaser 2 agreed to that proposal, then the offer was sent on to the Provider for their consideration. If the Provider also agreed to the payment then the negotiations are over and a deal is done. Alternatively, if the proposal or offer was rejected then a second round of bargaining was initiated. In this second round,

[35] Although we only have one seller of the ecosystem services it is perhaps more insightful to think of the seller as a single representative of multiple farmers charged with the task of negotiating on their behalf.

[36] Experiment participants did not know the identities of the other members of their group and the members of each group changed from task to task.

Purchaser 2 had the opportunity to make a counterproposal. If this counterproposal was agreed by both Purchaser 1 and the Provider then a deal was done; otherwise the process of proposal and counterproposal continues until all players agree or until negotiations fail. Figure 3.1 details the negotiation procedure. In order to keep the decision process moving and so that all groups progressed at the same speed, all decisions, proposals and acceptances, were subject to time limits. In the event that a Purchaser timed out when making a proposal default contributions of £0 for each Purchaser were submitted; when a responder (Purchaser or Provider) timed out the default submission was rejection.

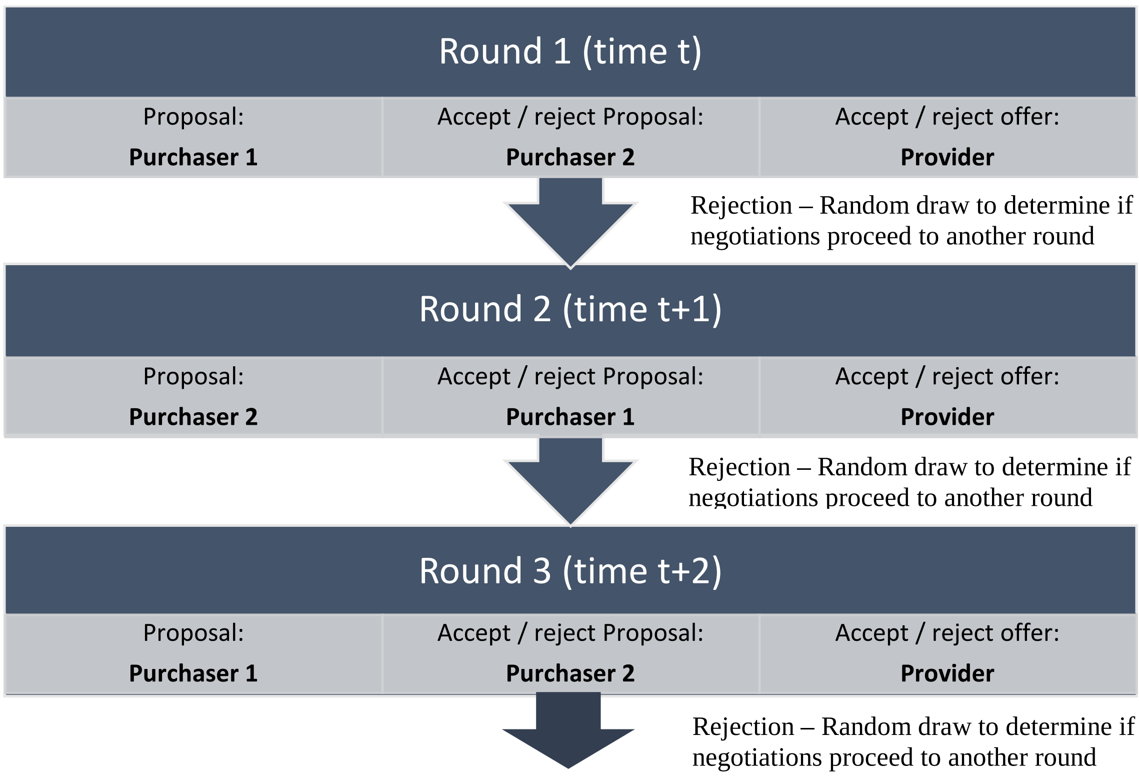

Figure 3.1. The negotiation procedure used in all experiment treatments.

Accordingly, negotiations proceed through a maximum of 15 rounds in which objections could be raised either by a purchaser who did not like the levels of contribution proposed by the other purchaser or by the Provider who felt that an offer made by the purchasers was not sufficiently generous. To represent the cost of negotiating rejecting a proposal and therefore blocking a deal came with an exogenous risk of the negotiations failing[37]. In the experiment that risk is randomly realised by the computer each time a rejection is made, with odds of failure being clearly displayed to the subjects. Those odds of failure increased as the rounds of negotiation progressed putting increasing pressure on the subjects to reach a deal.

Payments

In a task, both purchasers and the Provider began the experiment entitled to some *default payment*. Each participant was informed that they could forego their default payment in favour of a *deal payment*, but only if all three members of the group agreed to that arrangement. For the two purchasers, their deal payment exceeded their default payment such that they were always incentivised to reach a deal. In contrast, for the Provider their deal payment was always zero.

To persuade the Provider to incur the cost of forgoing their default payment, the two purchasers were given the opportunity to team-up and offer the Provider a payment. The essential challenge of the experiment was for the purchasers to agree how much each should contribute to an offer that was sufficiently generous to persuade the Provider to agree to the deal. Obviously, the most it would ever make sense for one purchaser to contribute to such a payment is the difference between their deal payment and their default payment. Though for each purchaser, the incentive is to contribute as little as possible hoping that the other purchaser will commit sufficient funds to the offer to ensure that the Provider to agrees to the deal.

Experimental treatments

Five key measures are used to assess differences in the experimental treatments results. The first two measures, success and ease of a deal, measure if a deal was reached and the ease of the negotiation process. The first measure, success of the negotiation, measures the proportion of deals in which agreement was reached between all parties. The second measure, ease of reaching a deal, measures the rounds of negotiation taken to reach a successful deal. Three further measures assess the composition of payoffs and contributions making up a deal. The third measure, contribution, measures the

[37] An example of an exogenous risk to negotiations can be seen in (Muthoo 2002) who propose that while two firms negotiate over how to divide the profit from the exploitation of a particularly technology a third firm invents a new technology that makes their existing technology obsolete.

payment made by an individual Purchaser's to the Provider, in addition, we measure the division of contribution as the Purchasers' share of the payment made to the Provider. The fourth measure uses the concept of a surplus, here we define the surplus from a deal as the sum of the differences between the deal payments and the default payments, in other words, the surplus is the aggregate gains from a deal.[38] In addition, the division of surplus measures the share of surplus each party gained in a deal. The fifth measure, the level of payoffs, measures the final payoffs received by each party.

Using the range of measurements outlined above, these experiments explore bargaining as a mechanism for effecting transactions in a multi-purchaser PES scheme. In particular, to characterise how bargaining outcomes might differ under a range of circumstances that might well arise in a multiple-purchaser PES scheme.

Our experimental design features a total of 10 treatments designed to explore five different issues: (i) the degree to which the buyers offer (and the seller accepts) an amount over and above the sellers costs; (ii) the degree to which *asymmetry in the gains* enjoyed by the two buyers from a successful transaction affects the outcome of negotiations; (iii) the degree to which *asymmetry in the income* of the two buyers (irrespective of their gains from the transaction) affects the outcome of negotiations; (iv) how negotiations differ under conditions of *incomplete information*; particularly regarding the payoffs both inside and outside of a deal; (v) how negotiations evolve when the benefits enjoyed by the buyers from the transaction are not known for sure but are *stochastic* in nature.

Offers above seller's costs with symmetric purchasers (Treatments 1A and 1B).

Two treatments are included with symmetric default and deal payments for the purchasers. In Treatment 1A, the purchasers have deal payments of £15 and all parties have default payments of £7.50. Under these circumstances, it was just cost effective to make a bilateral deal between one purchaser and the Provider. For example, if one Purchaser chose to compensate the Provider alone then it would require all their gains from a deal (£15 - £7.50 = £7.50) to make the Provider indifferent between agreeing to the deal or not. Under these circumstances we are interested in finding out if the

[38] For example, consider a situation in which all parties have a default payment of £7.50, the sum of default payments is therefore £22.50 (3 x £7.50), if the sum of deal payments is equal to £25 we have a surplus of £2.50, if the sum of deal payments is equal to £30 the surplus is equal to £7.50.

purchasers split the contributions equally. Furthermore, differences in payoffs between the purchasers and Provider would suggest differences in negotiating power of purchasers who make offers and Provider who can only accept. Under these circumstances, do the purchasers just offer enough to keep the Provider interested or are they forced to give up some of the surplus to the Provider? Results from ultimatum experiments show that responders are able to get more than the game theoretical prediction of close to zero but generally less than an equal split. In addition, evidence from (Schmitt 2004) on an alternating bargaining suggests that players adjust their offers until they find a minimum acceptable offer.

In Treatment 1B we increased the size of the surplus in the deal by having deal payments of £20 for both of the purchasers. By increasing the size of the deal payments, we were able to explore how negotiated agreements changed as the benefits the purchasers realised from making a deal increased. We were interested to see whether deals were reached more easily or more quickly in these circumstances and the extent to which the Provider was able to claim some of that increased surplus by holding out for a higher payment.

Asymmetry in gains from a deal (Treatment 2A).

Treatment 2 introduces asymmetry in deal payments. In Treatment 2A, the deal payments are £18 for Purchaser 1 and £12 for Purchaser 2 with default payments of £7.50 for all three players. Accordingly, it is possible for Purchaser 1 to form a cost-effective bilateral deal with the Provider but not Purchaser 2. Observe that the deal payments in this Treatment give the same overall surplus from a deal as Treatment 1A (sum of deal £30 - sum of default £22.5 gives a surplus of £7.50).

In many potential multiple-purchaser ecosystem service schemes we might expect one of the purchasers to do relatively better out of a deal than others. By including asymmetry in the benefits of a deal we wished to explore whether such asymmetry would result in a difference in contributions towards the negotiated payments. A number of experiments—notably, Roth, et al. 1981; Hoffman and Spitzer 1985; Bruce and Clark 2010a, and 2010b—found that their subjects were drawn towards Pareto-efficient outcomes that equalized payoffs. In order to achieve equal payoffs within our multilateral multiple-purchaser setup the Purchasers would need to agree to an unequal division of contributions. This will provide insights into whether the subjects are motivated by equality of contribution, over fairness in division of the gains from a deal or equal payoffs.

Asymmetry in the income of the purchasers (Treatments 3A).

While purchasers may not differ in terms of the additional benefits they stand to realise from a deal, they may differ in terms of the overall level of benefits they enjoy. In an ecosystem service market, that might correspond to having one purchaser that is a large, high-profit company and another that is a small, low-profit company.

In the experiment, Treatment 3A deal payments are £18 for Purchaser 1 and £12 for Purchaser 2, while their corresponding default payments were £10.50 and £4.50. As such, both purchasers stood to make the same gain of £7.50 from a deal. Observe that the total surplus is comparable to previous Treatments 1A and 2A.

We were particularly interested in whether the difference in default payments of the purchasers would affect the nature of the negotiations. Two differences are important in Treatment 3A compared to 2A. Firstly, equalising all three players payoffs cannot be achieved without Purchaser 1 accepting a deal that gave her less than her default payment. Secondly, the setup means the Purchasers cannot simultaneously reach a deal with an equal division of surplus and equal payoffs. This will provide insights to answer whether Purchasers are motivated by fairness in division of the gains from a deal over equality in payoffs.

Incomplete information (Treatments 2B, 2C, 3B and 3C).

In both Treatments 2B and 2C the deal and default payments are identical to Treatment 2A; therefore the deal payments are £18 for Purchaser 1 and £12 for Purchaser 2 with default payments of £7.50 for all three players. Two additional treatments, 3B and 3C are analogous to Treatments 2B and 2C but with the deal and default payments of Treatment 3A. The purchasers have different default payments (£10.50 and £4.50) but stand to gain the same amount from a deal.

In the real world, it is unlikely that purchasers, negotiating over contributions, will reveal the level of benefits they stand to enjoy from a deal being done. To reflect that reality, the deal payments in Treatments 2B and 3B were private information; only a Purchaser knew what they stood to gain from a deal and Treatments 2C and 3C introduce completely private information, with each player only knowing their own deal and default payments. Shupp et al. (2013) explore incomplete information in an alternating bargaining environment and find that negotiation takes longer and negotiations are more likely to result in failure. Contrary evidence shows that acceptance rates are higher with incomplete information perhaps because envy of the other player's payoff is removed with players instead focusing on their monetary payoff rather than relative payoff (Bolton and Ockenfels 2000; Schmitt 2004). The experimental design employed here allows the measurement of the number of rounds until negotiation leads to a deal or to failure allowing a measurement of the ease of reaching a deal.

Bruce and Clark (2010a) investigate unstructured bargaining under private information. When comparing their results to full information, they found that subjects were less drawn to outcomes that equalised payoffs. By including incomplete information in our experimental treatments the subjects are no longer able to equalise payoffs or to equalise the division of surplus because they no longer have all the required information to reach such outcomes. Under incomplete information we are therefore interested in any changes to the division of surplus between the parties, for example do the Purchasers move to equalise contributions to the payment under incomplete information treatments and is the Provider able to extract any of the surplus?

Furthermore, we explore if subjects are undertaking some form of cost benefit analysis as suggested by Zwick and Chen (1999) and Zwick and Mak (2012). Under this hypothesis, we would expect that in Treatment 2, as Purchaser 1 has more to gain from a deal, and therefore also more to lose from not reaching a deal, that Purchaser 1's contributes more to the payment compared to Treatment 3.

Stochastic benefits (Treatments 4+ and 4-).

Treatments 4+ and 4- explored stochastic outcomes. In those treatments, participants are faced by an uncertain future characterised by two possible states of the world, *a* and *b*, each of which has a 0.5 probability of becoming reality. The deal payments enjoyed by purchasers in these two states of the world are different. Accordingly, our experiment allows purchasers to make conditional offers; that is to say, purchasers decide how much they are going to contribute to a payment in each 'state of the world'. In the stochastic treatments, therefore, the Provider receives two offers; one detailing how much they will be paid by the purchasers in the event of ‘state of the world *a’* and another stating how much they will be paid in ‘state of the world *b’*. In these treatments during negotiations purchasers and the Provider could signal to one another regarding which, if either, of the payments in the two ‘states of the world’ had caused them to reject a proposal or offer.

Stochastic benefits are common for PES scheme in that buyers often have to enter PES schemes without full knowledge of the benefits they will receive due to the random qualities of natural processes. Of primary interest was to see if stochastic benefits alter the success and ease of reaching deal as well as any differences in the composition of payments and contributions in different states of the world. Do the Purchasers share the risk of the stochastic benefit between each other and is any of that risk transferred onto the Provider.

The experimental design is summarised in Table 3.1.

Implementation

The experiment was conducted in the Centre for Behavioural and Experimental Social Sciences at the University of East Anglia (UEA). In total 204 subjects participated in 14 sessions with between 12 and 18 subjects in each session. The subjects were recruited from the UEA undergraduate and graduate populations from a variety of disciplines using the online recruitment system (ORSEE) (Greiner 2004). No subject participated in more than one session, but each subject participated in multiple treatments within a single session, before each treatment the subjects were assigned a new role (Purchaser 1, 2 or Provider) and matched with a new group. The experiment was programmed and conducted in the z-Tree software (Fischbacher 2007).

The same procedure was followed in each session. All students were seated at a private computer terminal with no communication allowed. Instructions (see Appendix C) were read aloud and included a detailed walkthrough in which subjects could see how the tasks would progress from the perspective of all the different roles. Any questions were answered in private. At the end of the experiment subjects completed a short questionnaire on the computer. The final payments were from one randomly selected round and an additional £2.50 participation fee. The sessions lasted approximately 1 hour and 30 minutes including payments.

3.4. Results

Irrational responses

In total 412 deals were completed. If we take the standpoint that it would be irrational for the participants to accept a deal which would make them worse off than their default payment; then, of the 412 deals, 47 deals can be classified as irrational for at least one of the participants, leaving 365 rational deals. One can hypothesise about the reason for the high number of irrational deals, for example, although great care was taken to explain the procedure, the tasks were fairly complex and required the participants to be engaged. In addition, the role of the subject, deal and default payments changed from task to task. This complexity combined with the time pressure which the participants were under could go some way to explaining irrational bargaining behaviour. A small number of participants in the questionnaire admitted to making a mistake due to the time pressure. It should also be noted that other reasons such as altruistic behaviour could potentially explain some of the deals we have classified as irrational. Unless otherwise stated we use only those deals which can be classified as rational for all players in the following analysis.

Table 3.1 summarises the results for all treatments of the experiment and Table 3.2 summarises the division of contribution, division of surplus and division of payoffs for all deterministic treatment of the experiment.

Treatment		Information	Default payment			Deal payment			Deal Success	Average Rounds to completion	Average Contribution		Average Payoff		
			P1	P2	Provider	P1	P2	Provider			P1	P2	P1	P2	Provider
1	**A**	**Full Public**	£7.50	£7.50	£7.50	£15	£15	£0	39/40	3.3 (2.8)	£4.78 (0.83)	£4.82 (0.74)	£10.22 (0.83)	£10.18 (0.74)	£9.60 (1.49)
	B	**Full Public**	£7.50	£7.50	£7.50	£20	£20	£0	39/40	3.0 (2.3)	£6.20 (0.77)	£6.20 (0.87)	£13.80 (0.77)	£13.80 (0.87)	£12.40 (1.61)
2	**A**	**Full Public**	£7.50	£7.50	£7.50	£18	£12	£0	63/73	4.6 (3.5)	£7.55 (0.91)	£2.26 (0.77)	£10.45 (0.91)	£9.74 (0.77)	£9.81 (0.78)
	B	**Deal Private**	£7.50	£7.50	£7.50	£18	£12	£0	32/38	7.6 (3.2)	£5.89 (1.21)	£3.46 (0.91)	£12.11 (1.21)	£8.54 (0.91)	£9.35 (0.70)
	C	**Deal & Default Private**	£7.50	£7.50	£7.50	£18	£12	£0	24/30	8.2 (3.4)	£5.70 (1.42)	£3.53 (1.04)	£12.30 (1.42)	£8.47 (1.04)	£9.23 (0.92)
3	**A**	**Full Public**	£10.50	£4.50	£7.50	£18	£12	£0	36/40	6.5 (3.7)	£6.08 (0.94)	£3.36 (1.13)	£11.92 (0.94)	£8.64 (1.13)	£9.44 (1.07)
	B	**Deal Private**	£10.50	£4.50	£7.50	£18	£12	£0	32/41	7.4 (3.1)	£5.50 (1.05)	£3.81 (0.81)	£12.50 (1.05)	£8.19 (0.81)	£9.32 (0.84)
	C	**Deal & Default Private**	£10.50	£4.50	£7.50	£18	£12	£0	35/40	6.2 (3.7)	£4.84 (0.81)	£4.40 (0.83)	£13.16 (0.81)	£7.60 (0.83)	£9.24 (1.34)
4 **+**	**a** **b**	**Full Public**	£7.50	£7.50	£7.50	£21 £15	£15 £9	£0	34/38	5.3 (3.7)	£8.47 (0.97) £6.42 (0.73)	£2.59 (1.28) £1.75 (1.36)	£12.53 (0.97) £8.58 (0.73)	£12.41 (1.28) £7.25 (1.36)	£11.06 (1.13) £8.16 (0.82)
4 **-**	**a** **b**	**Full Public**	£7.50	£7.50	£7.50	£21 £15	£9 £15	£0	31/33	3.9 (2.7)	£9.13 (1.29) £4.91 (0.45)	£0.57 (0.72) £4.81 (0.52)	£11.87 (1.29) £10.09 (0.45)	£8.43 (0.72) £10.19 (0.52)	£9.70 (0.80) £9.73 (0.88)

Table 3.1. Summary of results across all treatments.

For treatment 4, + represents treatment with positive correlation between the purchasers deal payments and – represents the treatment with negative correlation. a and b represent the responses to ‘state of the world a’ and ‘state of the world b’ respectively. Averages are means with standard deviations in brackets.

Treatment		Information	Division of contribution %			Division of surplus %				Division of Total Payoffs %			
			Total Contribution to Provider	P1	P2	Total surplus	P1	P2	Provider	Total payoff	P1	P2	Provider
1	A	Full Public	£9.60	49.8%	50.2%	£7.50	36.2%	35.8%	28.0%	£30	34.1%	33.9%	32.0%
	B	Full Public	£12.40	50.0%	50.0%	£17.50	36.0%	36.0%	28.0%	£40	34.5%	34.5%	31.0%
2	A	Full Public	£9.81	77.0%	23.1%	£7.50	39.4%	29.8%	30.8%	£30	34.8%	32.5%	32.7%
	B	Deal Private	£9.35	63.0%	37.0%	£7.50	61.4%	13.9%	24.7%	£30	40.4%	28.5%	31.2%
	C	Deal & Default Private	£9.23	61.7%	38.3%	£7.50	64.0%	12.9%	23.1%	£30	41.0%	28.2%	30.8%
3	A	Full Public	£9.44	64.4%	35.6%	£7.50	19.0%	55.2%	25.9%	£30	39.8%	28.8%	31.5%
	B	Deal Private	£9.32	59.1%	40.9%	£7.50	26.6%	49.2%	24.2%	£30	41.7%	27.3%	31.1%
	C	Deal & Default Private	£9.24	52.4%	47.6%	£7.50	35.5%	41.3%	23.2%	£30	43.9%	25.3%	30.8%

Table 3.2. Summary of the division of contribution, division of surplus, and division of payoffs across all deterministic treatments.

Result 1. Offers above seller's costs with symmetric purchasers.

In Treatment 1 the Purchasers' deal and default payments are symmetric and all information is public. Treatment 1A and 1B differ only in the deal payments for the two purchasers; £15 in Treatment 1A compared to £20 in Treatment 1B.

Observe the deal success and average rounds to completion data in Table 3.1. The majority of the negotiations resulted in deals being completed, 39 from 40 for both Treatment 1A and 1B, and those deals were completed relatively quickly, with the mean number of rounds to competition of 3.3 for Treatment 1A and 3.0 for Treatment 1B.

In our experimental design the Provider was not able to specify offer amounts, this was chosen to more realistically simulate potential multiple-purchaser PES negotiations as it seems unlikely that the Provider would be able suggest amounts that each Purchaser should contribute. A question that arises from such a design is whether the relative bargaining strength of the purchasers relative to the Provider results in differing payoffs. In Table 3.1 the average payoffs for Treatment 1A for Purchaser 1 and Purchaser 2 were £10.22 and £10.18 respectively, while for the Provider the average payoff was £9.60. In Treatment 1B both purchasers received average payoffs of £13.80, and the Provider £12.40. Both treatments show a statistically significant difference between the Provider and the Purchasers payoffs (Treatment 1, Wilcoxon-Rank-Sum, *p*-value = 0.000 and Treatment 2, Wilcoxon-Rank-Sum, *p*-value = 0.000). For the Provider, rather than being able to bargain for equal payoffs, the negotiations move to a minimum acceptable amount above the Providers default payment. This is consistent with experimental evidence from the ultimatum and alternating bargaining literature that negotiations seek a minimal acceptable amount for the responder (Roth 1995; Schmitt 2004).

Result 2. Asymmetry in gains from a deal

Treatment 2 introduces asymmetry into the deal payments; specifically, Purchaser 1 has a deal payment of £18 and Purchaser 2 has a deal payment of £12.

In Treatment 2A[39] our first observation is that participants found it harder to complete a deal; the number of negotiation rounds averaged 4.6 as compared to 3.3 in the Treatment 1A and 3.0 in Treatment 1B. Figure 3.2 shows the distribution of the number of negotiation rounds until deals were completed for symmetric deal treatments (1A and 1B) and asymmetric deal Treatments (2A) a Wilcoxon-Rank-Sum test shows deals were completed quicker in the symmetric treatments (*p*-value = 0.010).

[39] An additional treatment was included in the study design which tested for difference in which Purchaser opened the negotiations. In the experiment, it was always Purchaser 1 who began the negotiations by making a first proposal of payments to be made to the Provider. Accordingly, it was always the purchaser with the higher (£18.00) deal payment that began the negotiations. We included a similar treatment to Treatment 2A except for the fact that the deal payments of Purchaser 1 and Purchaser 2 were swapped. Accordingly it is the subject with the lower, £12.00, deal payment who opened the negotiations. The results were qualitatively identical and are therefore reported grouped together with treatment 2A.

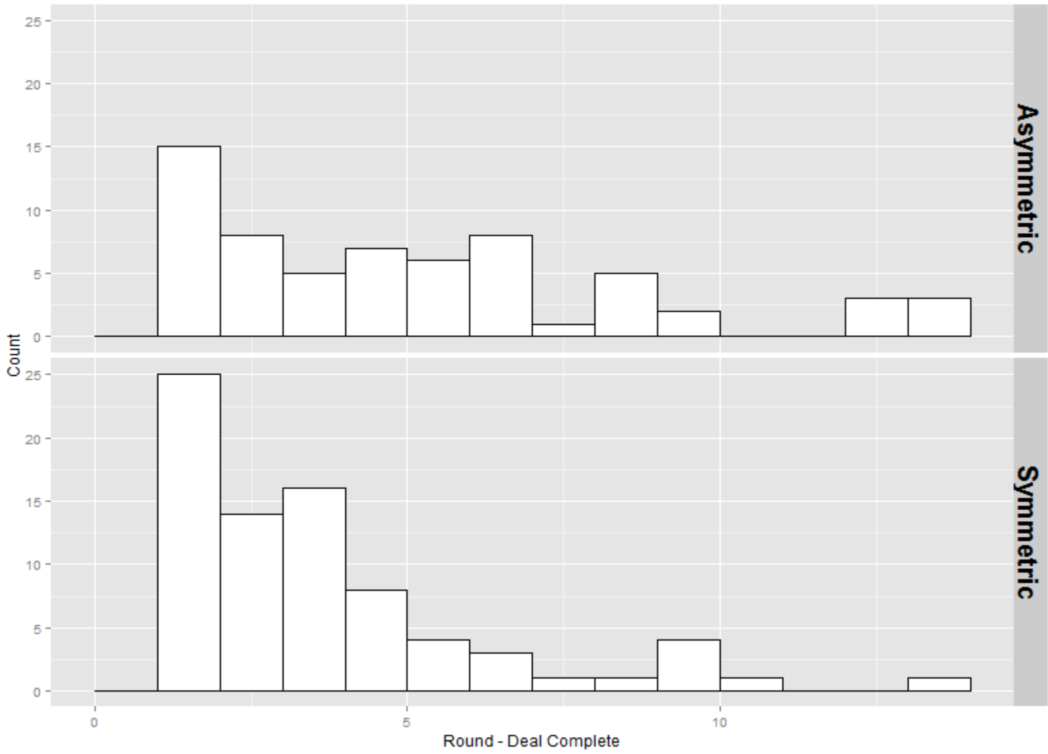

Figure 3.2. Number of rounds until deal completed for symmetric (1A 1B) and asymmetric (2A) experimental treatments.

With asymmetry in the gains from a deal the subjects cannot simultaneously achieve equal contributions and equal payoffs or equal contributions and equal division of surplus. In Treatment 2A statistically significant asymmetry in the contributions of the Purchasers are observable; the average contribution of Purchaser 1 is £7.55 and for Purchaser 2 it is £2.26 (Wilcoxon-Rank-Sum, *p*-value = 0.000). Purchaser 1 therefore contributed more to the deal but on average still ended up with a significantly higher average payoff £10.45 than both Purchaser 2, £9.74 (Wilcoxon-Rank-Sum, *p*-value = 0.000) or the Provider, £9.81 (Wilcoxon-Rank-Sum, *p*-value = 0.000). Since the default payments for the Purchasers are equal at £7.50 equivalent statistically significance differences are also observable for the division of surplus from a deal.

We have strong evidence that equal contributions are not the primary motivating factor in negotiations since Purchaser 1 is contributing significantly more in Treatment 2A than Purchaser 2, as a Consequence, Purchaser 1's share of the surplus is a lot smaller than their share in terms of gains from a deal. Figure 3.3 plots the share of the surplus for deals agreed in the symmetric (1A) and asymmetric treatments (2A). It is constructed as a ternary plot with Purchaser 1's share of the surplus on the left edge, Purchaser 2's share on the bottom edge and the Provider's share on the right. Each edge goes from 0-1 which represents the division of surplus as a proportion. Each dot represents a completed deal and the mean division of surplus amounts (reported in Table 3.2) are plotted in Figure 3.3 as the solid lines. In the symmetric treatment both Purchaser 1, 36.2%, and Purchaser 2, 35.8%, claim a significantly higher division of the surplus than the Provider, 28.0% (Wilcoxon-Rank-Sum, *p*-value = 0.001 and *p*-value = 0.001). However, in the asymmetric treatment Purchaser 1, 39.4%, is able to claim a statistically significant higher division of the surplus than either Purchaser 2, 29.8%, or the Provider, 30.8% (Wilcoxon-Rank-Sum, *p*-value = 0.000 and *p*-value = 0.000).

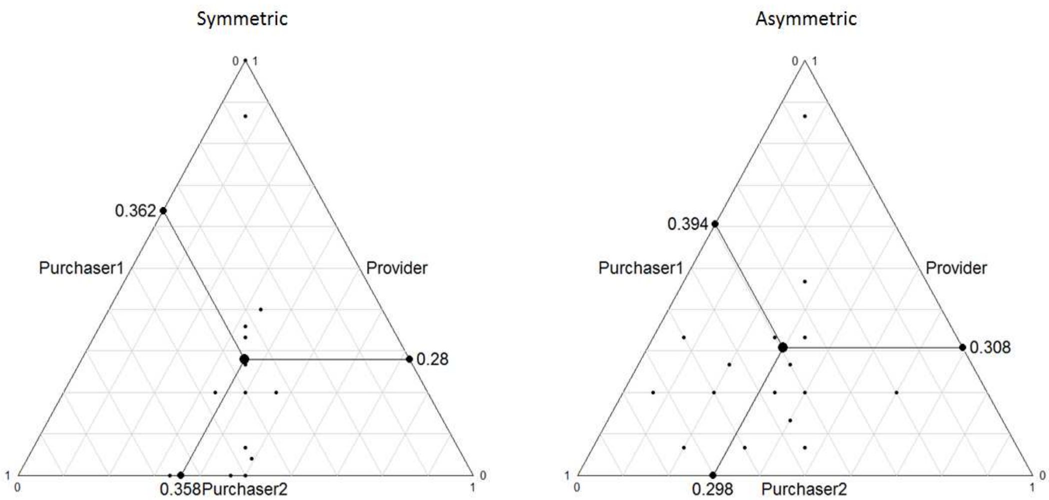

Figure 3.3. Ternary plot of the proportional split in the total surplus from the symmetric (1A) and asymmetric (2A) experimental treatments, with the lines showing the mean amount for each agent.

If the Purchasers were attempting to share the surplus equally we would expect no differences between the division of surplus in Treatment 1A and Treatment 2A. Comparing Treatment 2A to 1A using average division of surplus provides conflicting evidence on whether the subjects are attempting to share the surplus equally. One the one hand, statistically significant difference is observable between the division of surplus for Purchaser 2 (35.8%) in the symmetric treatment and Purchaser 2 in the asymmetric treatment (29.8%) (Wilcoxon-Rank-Sum, *p*-value = 0.000); on the other hand, there is no statistically significant difference for Purchaser 1's surplus

(Wilcoxon-Rank-Sum, *p*-value = 0.403) despite an increase from 36.2% to 39.8% of the division of the surplus. The number of deals in which the Purchasers negotiated to exactly equal payoffs provides additional evidence on this. In two-thirds of the deals in Treatments 2A (42 from 63) the Purchasers agreed to an equal division of the surplus and therefore equal payoffs too. It therefore looks like they are either (i) sharing the surplus roughly equally or (ii) agreeing a deal in which they their payoffs are roughly equal. Since the default payments are the same for the two Purchasers in Treatment 2A we cannot tell between those two hypotheses in this treatment, to provide further clarity we turn to Treatment 3A.

Result 3. Asymmetry in the income of the purchasers.

Treatment 3A provides the same deal payments as Treatment 2A, the difference is the introduction of asymmetry into the default payment. It is Purchaser 1 who gets the high deal payment of £18 and also the high default payment of £10.50; Purchaser 2, has a low deal payment of £12 and a low default payment of £4.50. Accordingly, both purchasers stand to make the same gain from agreeing a deal, £7.50; as in previous treatments all information is public.

With asymmetry in the default payments the subjects cannot simultaneously achieve equal division of surplus and equal payoffs. If Purchasers are motivated by equal divisions of surplus then that would lead to unequal payoffs as the default payments are different, conversely if Purchasers are motivated by equal payoffs then that would lead to unequal divisions of surplus. Our results show very unequal divisions of surplus and relatively equal payoffs.

In Treatment 3A, Purchaser 2 claimed a significantly larger share of the surplus, 55%, compared to 19% for Purchaser 1 (Wilcoxon-Rank-Sum, *p*-value = 0.000). Figure 3.4 plots the division of surplus for Treatment 3A.

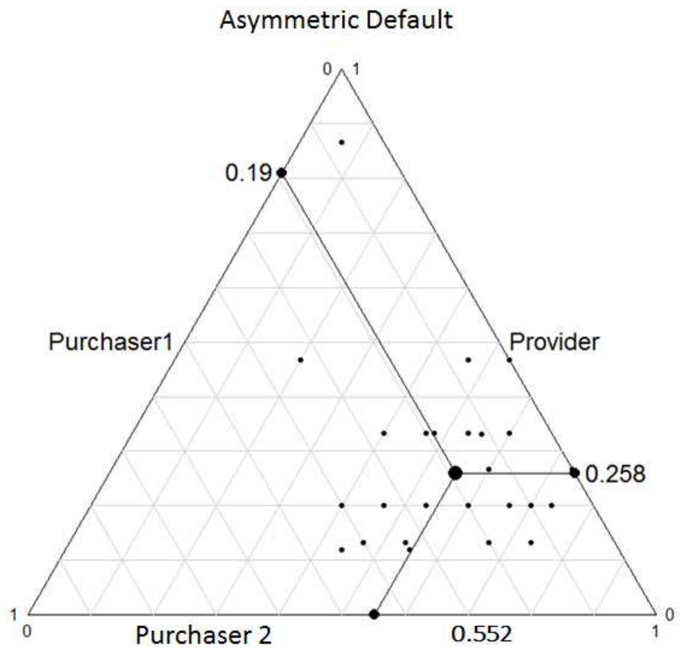

Figure 3.4. Ternary plot of the proportional split in the total surplus from the asymmetric default (3A) experimental treatments, with the lines showing the mean amount for each agent.

In Treatment 3A, it was not possible for the players to achieve equal payoffs unless Purchaser 1 decided to accept a deal in which she received less than her default payment of £10.50; however, in our data we observe 3 deals in which Purchaser 1 accepted an amount less than their default payment and everyone received a payoff of £10. It seems that, for some subjects, equal payoffs is still desirable despite one Purchaser having to accept less than their default amount. All other deals in Treatment 3A resulted in unequal payoffs, on average, Purchaser 1 contributed significantly more to the payment (Wilcoxon-Rank-Sum, *p*-value = 0.000) but still ended up with a higher overall payoff. Statistically significant differences are observable between the payoffs of Purchaser 1 and Purchaser 2 in Treatment 3A; the average payoff of Purchaser 1 was £11.92 and for Purchaser 2 was £8.64 (Wilcoxon-Rank-Sum, *p*-value = 0.000). In addition, payoffs in Treatment 3A were significantly different to payoffs in Treatment 2A (Purchaser 1, Wilcoxon-Rank-Sum, *p*-value = 0.000) (Purchaser 2, Wilcoxon-Rank-Sum, *p*-value = 0.000).

Given full information, it seems that equal payoffs is a strong point of attraction for bargaining outcomes. In Treatment 3A, even though both Purchasers have the same gains from a deal, and therefore bring equal amounts of surplus to the negotiating table, Purchaser 1 contributed significantly more to the deal and thus claimed a smaller share of the surplus, 19%, with Purchaser 2 claiming 55%. This evidence agrees with previous experimental studies that found that even with asymmetric gains from a deal their subjects were drawn towards outcomes that equalized payoffs (Roth, et al. 1981; Hoffman and Spitzer 1985; Bruce and Clark 2010a,2010b).

Result 4. Incomplete information.

Treatments 2B, 2C, 3B and 3C include incomplete information. The deal payments are private information in 2B and 3B, and the deal and default payments are private in 2C and 3C. Treatments 2B and 2C have the same deal and default payments as Treatment 2A, whereas, Treatments 3B and 3C have the same deal and default payments as Treatment 3A.

Our results provide no support for the hypothesis that incomplete information implies negotiations are more likely to result in failure. Comparing the proportion of successful deals in Treatment 2A to the proportion of successful deals in Treatment 2B and 2C, and similarly 3A to 3B and 3C, using Fisher's test of equality of proportions, reveals no statistically significant differences between the proportions (Treatment 2: Fisher exact test, p-value = 1, Treatment 3: Fisher exact test, p-value = 1).

Our results provide some support for the hypothesis that incomplete information implies more prolonged negotiation. For Treatment 2 an observable difference between the incomplete information treatments and full information treatment shows bargaining to be relatively harder; the average number of rounds of negotiation is 7.6 in Treatment 2B and 8.2 in Treatment 2C compared to 4.6 in Treatment 2A. Figure 3.5 shows the number of negotiation rounds taken to complete a deal for the full information treatment (2A) compared to the incomplete information treatments (2B and 2C). The results show that in the full information treatment deals were completed significantly quicker than the incomplete treatment deals (Wilcoxon-Rank-Sum, p-value = 0.000). However, these differences are not repeated when comparing Treatment 3A with 3B and 3C (Wilcoxon-Rank-Sum, p-value = 0.516).

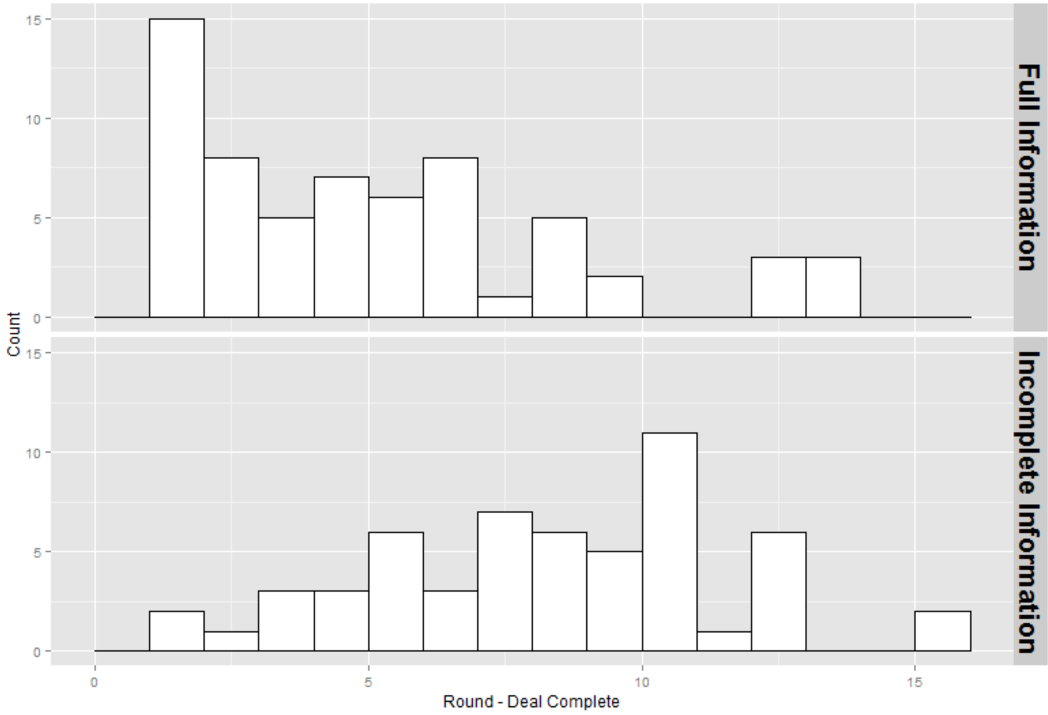

Figure 3.5. Number of rounds until deal completed for full information (2A) and incomplete information (2B 2C) experimental treatments.

We therefore have some evidence to support the conclusion in Shupp et al. (2013) who found that bargaining was more prolonged in incomplete information treatments but no support for their conclusion that incomplete information led to more failures to reach a deal.

Figure 3.6 illustrates the division of contributions between the two purchasers and the division of surplus between the purchasers and the Provider for all asymmetric treatments with complete and incomplete information.

For Treatment 2 a statistically significant reduction is observable for Purchaser 1 in the division of the contribution from 76.9% to 63% between Treatment 2A and 2B (Wilcoxon-Rank-Sum, *p*-value = 0.000), but only a small non-statistically significant drop between Treatment 2B and 2C (Wilcoxon-Rank-Sum, *p*-value = 0.390). This translates to significant changes in the division of surplus between Treatment 2A and 2B but not between Treatment 2B and 2C. Purchaser 1 was only able to capture 39.4% in Treatment 2A but in Treatment 2B they could capture 61.4% of the surplus (Wilcoxon-Rank-Sum, *p*-value = 0.000). At the same time, Purchaser 2 was able to capture 29.8% of the surplus in Treatment 2A but only 13.9% in Treatment 2B

(Wilcoxon-Rank-Sum, p-value = 0.000). With Private unequal deal payments the purchaser with the high deal payment (Purchaser 1) was able to capture a much higher share of the surplus.

For Treatment 3 statistically significant reductions are observable for Purchaser 1 in the division of contribution from 64.4% to 59.1% between Treatment 3A and 3B (Wilcoxon-Rank-Sum, p-value = 0.022) and from 59.1% to 52.4% between Treatment 3B and 3C (Wilcoxon-Rank-Sum, p-value = 0.009). These translate to statistically significant changes in the division of surplus between Treatment 3A and 3B, and 3B and 3C with Purchaser 1 moving from 19% in Treatment 3A to 26.6% in Treatment 3B (Wilcoxon-Rank-Sum, p-value = 0.022) and 35.5% in Treatment 3C (Wilcoxon-Rank-Sum, p-value = 0.009). Statistically significant differences are also observed for Purchaser 2's share of the surplus between Treatment 3A and 3B (Wilcoxon-Rank-Sum, p-value = 0.018) and 3C (Wilcoxon-Rank-Sum, p-value = 0.011).

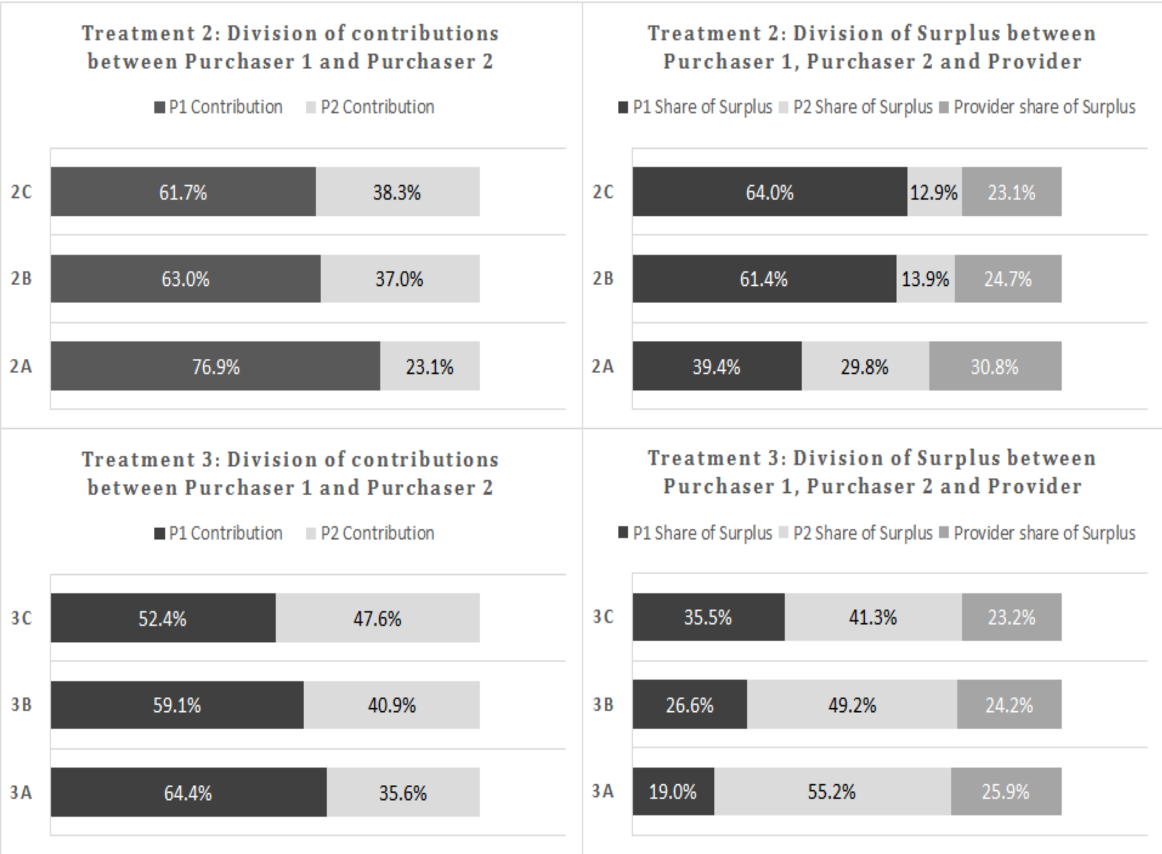

Figure 3.6. Division of contributions and division of surplus for asymmetric deal and asymmetric default treatments (Treatments: 2A, 2B, 2C, 3A, 3B, 3C).

Since the purchasers did not know the other purchasers default payment in either Treatment 2C or 3C we claim that a reasonable assumption for a purchaser would be to assume that the other purchaser had the same default payment as them. Under that assumption, we would expect no difference in the division of surplus between Treatment 2B and 2C because the purchasers in treatment 2C would assume, in this case rightly, that they have equal default payments. At the same time, we would expect to see a difference between treatment 3B and 3C because the assumption of equal default payments in Treatment 3C would be false. The division of surplus results support this hypothesis. A statistically significant difference was observable between the division of surplus in Treatment 3B and 3C (Wilcoxon-Rank-Sum, Purchaser 1: p-value = 0.009, Purchaser 2: p-value = 0.011) but not Treatment 2B and 2C (Wilcoxon-Rank-Sum, Purchaser 1: p-value = 0.387, Purchaser 2: p-value = 0.755).

Finally, we explore the difference in the contributions in Treatment 2C compared to Treatment 3C. In Treatment 3C the purchasers' contributions to the payment are close to equal, Purchaser 1 contributes on average 52.4% whereas Purchaser 2 contributes 47.6%. In the absence of knowledge about the other purchaser's deal and default payment we might expect the Purchasers to decide that equal contributions is fair. Although in Treatment 3 there is a small difference in contributions there is a much larger difference observable in Treatment 2C, Purchaser 1, 61.7% and Purchaser 2 38.3%. This suggests a fundamental difference between the deals completed in the two fully private information treatment. One possible explanation comes from Zwick and Mak's (2012) cost-benefit proposal that because Purchaser 1 has more to gain from a deal in Treatment 2C than Purchaser 2, they will be more willing to compromise and offer to contribute a higher share of the payment.

Result 5. Benefits of the deal are stochastic.

The Treatments 4+ and 4- are full information treatments, such that the deal and default payments are public information, but the actual deal benefits are stochastic. The key design feature that distinguishes Treatment 4+ and 4- is that in Treatment 4+ deal payments are positively correlated; both Purchaser 1 and Purchaser 2 get their higher payment in 'state of the world a' and their lower payment in 'state of the world b'. In contrast, in Treatment 4- deal payments are negatively correlated; Purchaser 1 realises their high payment when Purchaser 2 realises their low payment and vice versa; note that even purchasers lower payment are higher than their default payment. In addition, the payments are designed to maintain an asymmetry between the purchasers; Purchaser 1 has opportunities for deal payments that are as large if not larger than those of Purchaser 2. For both purchasers the deal payment in one 'state of the world' is larger than that in the other.

Despite the added complexity in the bargaining procedure, participants were able to reach a deal in a similar number of rounds to asymmetric full information treatments. Figure 3.7 shows the number of negotiation rounds until deals were completed for the deterministic treatments (Treatments: 2A, 3A) compared to the stochastic payments (Treatments: 4+, 4-). The results show no statistically significant differences in the number of rounds to complete a deal (Wilcoxon-Rank-Sum, *p*-value = 0.247).

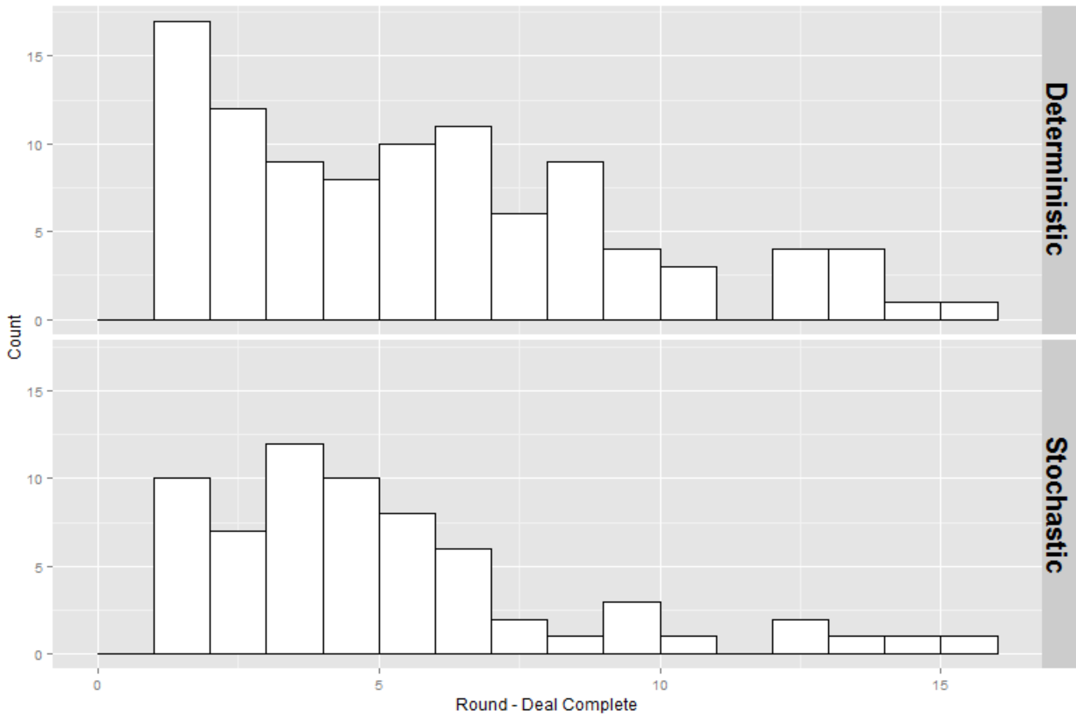

Figure 3.7. Number of rounds until deal completed for deterministic (2A, 3A) and stochastic (4+ 4-) experimental treatments.

Consider now the outcome of Treatment 4+ in Figure 3.8 where deal payments are positively correlated across states of the world. Rather than offering the Provider the same payment no matter what the 'state of the world', the deal that the purchasers agree to is one where they pay the Provider more, £11.06, if things turn out well and they both get their high payments ('state of the world a') and less, £8.16, if things turn out badly and they both get their low payments ('state of the world b'). In effect, in this treatment the purchasers push some of the risks of a bad outcome onto the Provider.

In contrast, observe the outcome of Treatment 4- in Figure 3.8 where deal payments are negatively correlated. In this case, the Provider gets very similar payments in both 'states of the world', £9.70 and £9.73. The purchasers, on the other hand, arrange their payments quite differently. Compared to positively correlated case, they pay relatively more when they are the purchaser to enjoy the 'good state of the world' (such that the other purchaser experiences their 'bad state') and relatively less when they are the purchaser to enjoy the 'bad state of the world' (such that the other purchaser experiences their 'good state'). Essentially, in Treatment 4- the purchasers share the risk of different possible outcomes between themselves rather than with the Provider.

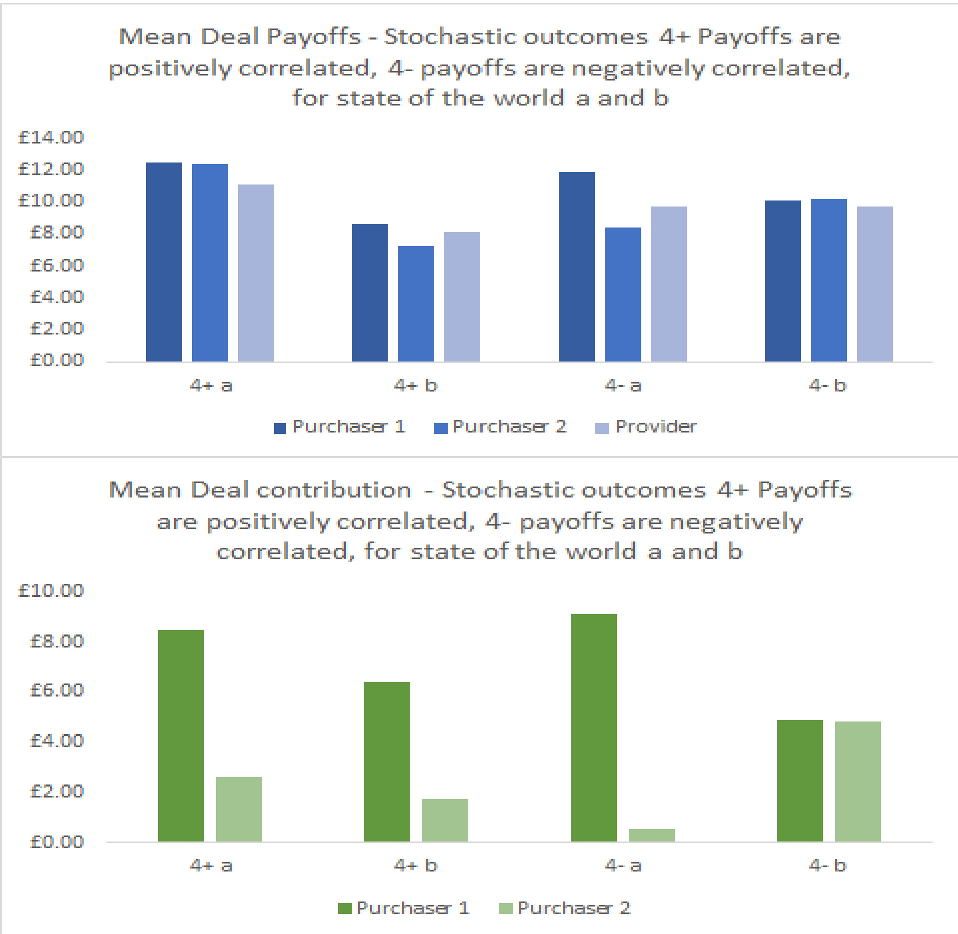

Figure 3.8. Deal payoffs and contributions from the stochastic benefit treatments (Treatments: 4+ a, 4+ b, 4- a, 4- b).

Furthermore, the stochastic treatments provide additional evidence that the participants are drawn to equal payoffs when information is complete. The top panel of Figure 3.8 shows the mean deal payoffs for Treatments 4+ and 4- are very similar for each participant in all of the stochastic treatment scenarios, this is achieved despite the large differences in the potential deal benefits (£21 Purchaser 1, £9 Purchaser 2 in 4- a). The bottom panel of Figure 3.8 shows that Purchaser 1, the purchaser who stands to gain the most from a deal, has contributed much larger amounts on average than Purchaser 2 in order to achieve equal or near equal payoffs.

3.5. Conclusion and discussion[40]

In this chapter, we analyse dual-purchaser multilateral bargaining as a mechanism for procuring ecosystem services. In the real world, it is likely that multiple-purchaser PES schemes are going to include organisations which vary with respect to the amount of benefits they will receive from a successfully negotiated deal, those benefits may not be known exactly by the purchasers or be difficult to quantify due to the range of stochastic natural processes involved, in addition, such organisation may also vary with respect to the benefits outside of a PES scheme, such as the costs and benefits of an alternative. As such, we implemented a broad series of experimental treatments to understand the nature of bargaining outcomes under a range of circumstances that might characterise a PES mechanism.

In all treatments, the Provider, unlike the purchasers, was not able to specify offer amounts, this difference resulted in the Provider, rather than being able to bargain for equal payoffs, instead negotiating towards a minimum acceptable amount above their default payment. This is consistent with experimental evidence from the ultimatum and alternating bargaining literature that negotiations seek a minimal acceptable amount for the responder (Roth 1995; Schmitt 2004).

A previous multilateral bargaining study found that with incomplete information negotiations were (i) more likely to result in failure and (ii) take longer to reach a successful deal (Shupp, et al. 2013). We find no evidence to support the first claim, the proportion of successful negotiations was similar in all the treatment scenarios

[40] Further concluding remarks on all three chapters, in which we highlight potential future extensions, can be found at the end of this thesis.

presented—symmetric and asymmetric benefits inside and outside a deal, complete and incomplete information and stochastic payoffs—the vast majority of groups were able to reach agreements. We find some evidence to support their second claim, the number of rounds of negotiation was significantly higher in the incomplete asymmetric deal payment treatments (2B and 2C) compared to Treatment 2A but this pattern was not repeated in the incomplete asymmetric default payment treatments (3B and 3C) compared to Treatment 3A. Overall, as we added realism into the experiment through inequality in benefits and incomplete information, the number of rounds needed to reach a deal increased. While in our experimental framework most of the deals were still completed it is important to note that the increase in the length of negotiations could cause real world negotiations to be a drawn-out and potentially costly process.

Existing studies have found that subjects were drawn towards outcomes that equalized payoffs (Roth, et al. 1981; Hoffman and Spitzer 1985; Bruce and Clark 2010a,2010b). The results of this chapter support this conclusion, given full information, it seems that equal payoffs is a strong point of attraction for bargaining outcomes.

Furthermore, through varying the level of information, treatments were included in which the subjects could not identify the contributions that would lead to equal payoffs. By varying the level of information two hypotheses were tested. Firstly, one might speculate that with private information, since the purchasers can no longer negotiate to a fair distribution of payoffs, subjects might instead be drawn towards a deal in which the purchasers make equal contributions towards the payment to the Provider. Alternatively, one might speculate that even with private information if subjects are undertaking some form of cost benefit analysis as suggested by Zwick and Chen (1999) and Zwick and Mak (2012) we would expect that in Treatment 2, as Purchaser 1 has more to gain from a deal, and therefore also more to lose from not reaching a deal, that Purchaser 1 contributes more to the payment compared to Treatment 3. Our results support the cost-benefit analysis hypothesis. When purchasers have different default payments, but equal gains from a deal, the contributions were fairly equal, however, when one purchaser has more to gain from a deal, and therefore more to lose if a deal does not go ahead, that purchaser contributes more. We see this result even when all information is private and the purchasers do not know their relative advantages or disadvantages.

Finally, the stochastic benefits treatments provide further evidence of a preference for equality of payoffs when information is public. Ecosystem services differ from most traditional goods in that there exists an inherent uncertainty in their production. When paying for ecosystem services the actual outcome and payoffs are unknown to the purchasers. We achieved this by having a 50% chance of one outcome and a 50% chance of another; although we acknowledge that a continuum of potential outcomes would be more realistic, this was not feasible in a laboratory experiment. In the stochastic treatments with positive correlation between the purchasers' benefits, the purchasers shared the risk with the Provider, such that in the 'good state of the world' the Provider would receive a higher payment compared to the 'bad state of the world'; in contrast, in the stochastic treatment with negative correlation the purchasers shared the risk between each other. This pattern of payments led to fairly equal final payoffs between the subjects across the different treatments and different states of the world. In addition, the proportion of successful deals and the number of rounds of negotiation was comparable in the stochastic treatments to other asymmetric full information treatments.

The policy recommendations from these experiments are clear. Participants in negotiated dual-purchaser PES schemes can reach deals that are agreeable to all parties in a variety of scenarios relevant to real world negotiations, including asymmetric benefits and asymmetric incomes, incomplete information and stochastic benefits. However, the nature of both the benefits from negotiation and the structure of negotiation leads to different patterns of response. For example, the negotiations seek a minimal acceptable amount for the responder (the ecosystem service seller); with full information, the negotiations seek equal payoffs, this includes asymmetric benefit, asymmetric income and stochastic benefits (when the benefits were not known with certainty); with incomplete information, the evidence suggests that the purchasers implement some sort of cost-benefit thinking such that when one purchaser has more to gain from a deal, and therefore more to lose if a deal does not go ahead, that purchaser contributes more.

CONCLUDING REMARKS

In Part I of this thesis we used an existing CV dataset to explore respondents' preference uncertainty. Chapter 1 set out two clear aims: the first, to explore the hypothesis that common elicitation anomalies observed in CV studies may arise because respondents with uncertain preferences are required to answer as if those preferences were precisely-defined even when valuing complex and unfamiliar non-market goods; the second, to develop and implement a novel econometric method of analysing uncertain WTP data.

We found evidence that uncertainty alone cannot explain common elicitation anomalies such as starting point bias and higher WTP estimates in DC questions compared to OE questions. Uncertain preferences, like certain preferences, can be shifted up or down by the elicitation method used. The exact mechanism through which elicitation anomalies manifest is not yet apparent. Our study looked purely at uncertainty as an explanation; in reality there may be a combination of a number of factors which are leading to elicitation anomalies including other aspects regarding the form of the preferences and the idea that certain formats of CV elicitation encourage strategic (Carson et al., (2001)) or ill-considered responses (Poe and Vossler (2009), Hutchinson et al. (2007)).

To analyse our data we made use of a multi-state semi-parametric estimator, adapted from the duration modelling literature of the medical sciences. Our model assumed that individuals transition to different states of certainty as the amount on the WTP scale is increased to higher amounts. For example, an individual starts out, at the lower end of the WTP scale, certain that they would pay and as the amount increases they transition into a state of uncertainty and finally into a state of certainty about not paying. We consider the duration model to be a step forward compared to other models used to analyse uncertain CV data such as Wang (1997) and Evans et al. (2003) and is comparable to Kobayashi et al. (2012) in its ability to analyse thresholds in which respondents switch their certainty about paying for a good. Moreover, our estimator allows statistical analysis over the full range of the WTP value distribution without requiring restrictive parametric assumptions. A straight-forward expansion of the estimator would be to include more than three states, for example, by dividing the state of uncertainty into "probably would pay", "not sure" and "probably would not pay"; however including the polychotomous choice responses relies on the respondents interpretation of the term "probably", as noted by Hanley et al. (2009) this is unlikely to be identical for every respondent. A further extension to the model would be the inclusion of covariates such as income or experience/knowledge of the good, this would add to the richness of the model and allow further exploration into the uncertainty range. For example, respondents with larger incomes may be willing to pay more, shifting their uncertainty range higher up the willingness to pay scale.

Directly asking respondents in CV surveys their WTP is one of the few quantitative methods available to assess full economic value (including both use and non-use values) of non-market goods. It therefore remains vitally important to better understand the reasons we consistently observe elicitation anomalies in such surveys. Asking people to answer such questions as if they had precisely-defined preferences when they are uncertain of their preferences may be one element of this, however, our study showed that such uncertainty could not fully explain elicitation anomalies. It would be interesting to explore if more experienced or knowledgeable respondents had narrower uncertainty ranges, perhaps because their preferences were more well-defined prior to the survey, and if those respondents provided values that were procedurally invariant. A starting point might be to examine familiar goods, possibly in a laboratory environment, to understand if the uncertainty range is still malleable, or to examine experience goods to understand the influence of learning and experience on the uncertainty range.

Part II of this thesis considered multiple buyers PES schemes. In Chapter 2 we focused on the issue of PES mechanism design when the activity incentivised through the scheme benefits multiple groups each of whom might be prepared to contribute to payments made through the scheme. In particular, we focused on the issue of spatial coordination on the demand side of the market; that is to say, the question of which beneficiary of the PES scheme buys land-management changes on which land parcels. To study multiple-purchaser PES schemes, Chapter 2 developed a framework of methods. The framework can incorporate different buyers' objectives, for example objectives for different ecosystem service benefits, and include different constraints on those objectives; in addition, the framework solves a variety of PES purchasing institutions and does this over a variety of spatial landscapes.

Two simulation modelling environments were created to highlight the flexibility and power of the framework of methods. From these environments we are able to draw conclusions about the situations in which we expect a multiple-purchaser PES scheme to be practical. We conclude that negotiated solutions Pareto-dominate the independent and simultaneous solution suggesting that, as a minimum, institutions should be created that coordinate and facilitate negotiation between ecosystem services purchasers in a particular landscape. Moreover, for many problems there exist cooperative solutions that Pareto-dominate all strategic solutions suggesting that coordinating action through empowering a trusted broker to make decisions on behalf of both buyers could potentially benefit both buyers.

It would be of great benefit to investigate these findings under a wide range of different environments and a wide range of multiple-purchaser decision making problems. The example we presented in Section 2.5 showed two budget constrained buyers, although this is a common way of modelling ecosystem services buyers it is just one of a number of possible combinations of potential buyers. In Appendix B we show how the framework can easily incorporate not just budget constrained buyers but also target constrained, or profit maximising decision problems.

In Section 2.6 we presented a more complex and realistic example, in that, one buyer's benefits rely on the spatial heterogeneity of benefits from different land uses in the landscape and another buyer's benefits rely on the spatial interdependence and configuration of land uses in the landscape. By introducing spatial interdependence into either the costs or benefits we created a non-linear decision problem for the buyer. We showed how a non-linear spatially interdependent problem can be linearised and solved by our framework. Indeed, any buyer's decision problem can be included as long as it can be represented in a linear way, however, we acknowledge that the practicalities of solving some decision problems may not, in practice, be a simple task, particularly if exact optimal solutions are required. In future work, it would be interesting to investigate the technical challenges of combining non-linear decision problems; in addition, there are likely to be further technical challenges in expanding the multiple-purchaser framework to include more than two buyers' decision problems.

Furthermore, in Section 2.6, we presented an example of how our framework can also be used to generate optimal patterns of land use across a more realistic spatial landscape, this type of exercise could potentially be of interest to both buyers of ecosystem services and policy makers. Policy makers could, for example, study specific land-use configurations, producing solutions—exactly which sites to purchase and who should purchase them—that Pareto-dominate any solutions that could be negotiated by multiple-purchasers thinking of their own self-interest. The method presented therefore provides the groundwork for a potential policy-relevant practical tool for facilitating multiple-purchaser PES schemes.

One potential avenue for future research would be to apply the framework of methods developed in Chapter 2 to real world data. Modelling an actual catchment with real supply prices and real buyers of ecosystem services. By proving the method outside of the test environment it would provide increased policy relevance. A more ambitious expansion could be to develop the framework into an optimal spatial ecosystem service decision making tool for direct use by policy makers.

Inside the simulation modelling environment there are a number of potential expansions/improvements that could be pursued in future work. The negotiated decision problem is currently not solved to a point of convergence as would be expected by Rubinstein's (1982) alternate bargaining theory. This is due to the computationally intensive nested optimisations and genetic algorithm that is built into the negotiated decision problem, and would therefore require a smaller and simpler test environment to prove the concept. In addition, the decision making problems modelled in Chapter 2 do not allow the buyers to offer or receive any form of side payment, this is an area which has been shown to be important in the outcome of international agreements (Barrett 1994; Barrett 2001; Barrett and Stavins 2003). Finally, the landowner's costs used in the simulation environment are modelled as constant, this is clearly unrealistic. Heterogeneous costs can easily be included in the current framework and this includes real world data, where such costs exist. However, in PES schemes these costs are not just the cost to the farmer of providing environmental output, as shown in the auctions literature they may include an element of bid shading where the landowner would like to make a profit, alternatively, the landowner may be willing to accept less than their cost if they get some benefit from engaging in pro-environmental land-management. One way to model the incentives of multiple sellers within framework of methods would be to utilise agent-based modelling techniques.

Chapter 3 of this thesis also explored the issue of multiple purchasers for ecosystem services but focused on negotiation as a procurement mechanism. We designed and conducted novel three-person bargaining experiments in which two potential buyers negotiated not only between each other but also with a seller of ecosystem services to reach a mutually beneficial outcome. The experiments were structured as non-cooperative alternating bargaining experiments, where two buyers alternate in proposing how much each buyer should pay and therefore also how much the seller receives.

The experiments extend the literature on negotiation as an ecosystem service procurement mechanism by moving beyond bilateral negotiation to consider multilateral negotiation with multiple purchasers of ecosystem services. The results showed that the negotiation outcome is pulled around by the nature of the bargaining setup, future experimental work could explore the bargaining setup further. For example, previous experimental work from Binmore et al. (1991) highlights the difference between an exogenous random termination and a choice to exit negotiations. Additional experimental work could be undertaken using our multilateral ecosystem procurement setup but with participants able to opt out of the bargaining process and take their outside option. Giving participants the option to opt out of negotiation could lead to more breakdowns in the negotiation process, particularly if some participants view the benefits of the negotiation as unfavourable when compared to the time costs involved in the negotiation. Alternatively, it may also have the effect of increasing pressure on the proposer to provide payoffs considered fair to all parties earlier in the negotiation process to avoid negotiation breakdown.

Furthermore, Chapter 3 showed that asymmetric benefits both inside and outside of a deal affect negotiation between multiple purchasers, this could be a key factor in determining the outcomes of negotiation between multiple ecosystem service purchasers. The potential exists for future work inside an experimental framework, for example, taking the differences between the purchasers' deal payments to more extreme levels so that one purchaser has considerably higher benefits. We could then test for higher contributions from the purchaser with the most to gain from a deal; this would add evidence to answer the hypothesis that subjects are undertaking some form of cost-benefit analysis as suggested by Zwick and Chen (1999) and Zwick and Mak (2012).

Inside the experimental environment there are a number of potential expansions/improvements that could be pursued in future work. Currently the series of experiments lack a comparison with a treatment without a participant playing the role of the seller. Instead of a human seller, a control experiment could be conducted in which the buyers are negotiating towards a set price. Although one would assume that the buyers would negotiate so that they only paid that set price exactly, there may be interesting outcomes when the set price cannot be split evenly between the buyers. In the current design, the series of experiments are all set up such that the seller cannot propose a price that they are willing to accept. At the time of designing the experiment it was decided that it would be unrealistic to allow the seller to specify the amounts that they wanted the buyers to pay individually, however, one way around this would be to have the seller request a total amount and see if the buyers offer such an amount. Finally, as with Chapter 2, Chapter 3 could benefit from future work which included multiple sellers of ecosystem services. Additional experiments could be run as an auction in which the sellers of ecosystem services compete on the price they are willing to accept.

Extra experimental treatments would provide additional realism to the multiple-purchaser ecosystem service procurement experimental framework set out in this thesis, providing further evidence for the opportunities and barriers of multiple-purchaser ecosystem service schemes. Ultimately though, expansion beyond the experimental techniques to other quantitative techniques may be necessary, for example, moving beyond three way negotiation in a laboratory environment may prove too computationally challenging for the participants. An alternative would be to move to simulation modelling; for example, simulation through agent based modelling would allow for numerous providers of ecosystem services and enable the simultaneous study of both the demand and supply side of PES schemes under the same framework. Ideally, investigations into negotiated multiple-purchaser PES schemes would then move to small scale field trials to gain further insights of the kind that can only be attained through practical implementation.

Overall, this thesis applies quantitative techniques of environmental economics, contributing methodological advances in the econometric analysis of uncertain WTP data from CV surveys, the modelling of multiple purchaser PES schemes and the study of negotiated multiple purchaser PES schemes in laboratory experiments. Being able to draw on a variety of quantitative techniques provides the variety of evidence needed for policy makers to make better informed decisions on future resource allocation such as the topics discussed in this thesis—the decision to invest in new coastal defences or the decision to pay landowners to produce environmental output.

APPENDIX A1

The following appendix refers to Chapter 1 and derives the estimating equations for (p, λ, β) used in the semi-parametric estimator.

We first differentiate (1.5) with respect to p_s:

$$G(P_s) \;=\; \sum_{i=1}^{N}\left[\sum_{j=1}^{M+1} {d_{ij}}^{t}\;\ln(p_j)\right] + \;\mu_0\left(1 - \sum_{j=1}^{M+1} p_j\right) - \sum_{j=1}^{M+1} \mu_j\, p_j \qquad \text{(A.1)}$$

Taking the derivative of the first element:

$$= \;\sum_{i=1}^{N}\left[\sum_{j=1}^{M+1} {d_{is}}^{t}\;\left(\frac{1}{p_s}\right)\right]$$

d_{is} represents one specific d_{ij} for one p_s. The sum of d_{is} across all intervals and all individuals equates to the number of respondents who survived which we will denote n_s therefore:

$$= \;\frac{n_s}{p_s}$$

p_s is also present in two of the primal constraints which yield:

$$\frac{\partial}{\partial \mathrm{p}_s} = \;-\,\mu_0 \quad \& \quad \frac{\partial}{\partial \mathrm{p}_s} = \;-\,\mu_j$$

After combining the equations and setting equal to zero we rearrange to yield:

$$\frac{n_s}{p_s} - \;\mu_j = \;\;\mu_0$$

Multiplying through by p_s provides the following equation due to the complementary slackness constraint, $\mu_j p_j = 0$.

$$n_s = \;p_s\mu_0 \qquad \text{(A.2)}$$

Summing over all j where $j = 1,2, \dots, M + 1$

$$\sum_{j=1}^{M+1} n_j = \mu_0 \sum_{j=1}^{M+1} p_j$$

Hence after substituting $\sum_{j=1}^{M+1} n_j$ = N and $\sum_{j=1}^{M+1} p_j$ = 1 we obtain:

$$N = \mu_0$$

Finally substituting back into to (A.2) and rearranging gives the empirical estimator for p_s (1.6).

$$p_s = \frac{n_s}{N} \tag{A.3}$$

Next we differentiate (1.5) with respect to λ_r:

$$\begin{aligned} G(\lambda_r) = \sum_{i=1}^{N}\Bigg[\sum_{k=1}^{K} d_{ir}{}^{w}\ \ln\left(\frac{\lambda_r\, e^{\beta ti}}{1+\ \lambda_r\, e^{\beta ti}}\right) \\ +\ \delta_{ir}{}^{w}\ln\left(1-\ \frac{\lambda_r\, e^{\beta ti}}{1+\ \lambda_r\, e^{\beta ti}}\right)\Bigg] - \sum_{k=1}^{K}\gamma_k\lambda_r \end{aligned} \tag{A.4}$$

$$\sum_{i=1}^{N}\left[\sum_{k=1}^{K} d_{ir}{}^{w}\ \ln\left(\frac{\lambda_r\, e^{\beta ti}}{1+\ \lambda_r\, e^{\beta ti}}\right)\right] \tag{A.5}$$

$$\sum_{i=1}^{N}\left[\sum_{k=1}^{K} \delta_{ir}{}^{w}\ \ln\left(1-\ \frac{\lambda_r\, e^{\beta ti}}{1+\ \lambda_r\, e^{\beta ti}}\right)\right] \tag{A.6}$$

$$-\ \sum_{k=1}^{K}\gamma_k\lambda_r \tag{A.7}$$

Differentiating (A.5) we obtain the following equation due to $d_{ir}{}^{w}$ $(r\ =\ 1,2,\ldots,K)$ being one for respondents who failed in interval r and zero otherwise:

$$\sum_{i=1}^{N}\left[d_{ir}{}^{w}\ \left(\frac{\lambda_r\, e^{\beta ti}}{1+\ \lambda_r\, e^{\beta ti}}\right)^{-1}\ \left(\frac{e^{\beta ti}}{1+\ \lambda_r\, e^{\beta ti}} -\ \frac{\lambda_r\, e^{2\beta ti}}{(1+\ \lambda_r\, e^{\beta ti})^2}\right)\right]$$

Rearranging and cancelling we obtain:

$$\sum_{i=1}^{N}\left[d_{ir}{}^{w}\ \left(\frac{1}{\lambda_r} - \frac{e^{\beta ti}}{1+\ \lambda_r\, e^{\beta ti}}\ \right)\right]$$

Differentiating (A.6) we obtain the following equation due to $\delta_{ir}{}^{w}$ $(r\ =\ 1,2,\dots,K)$ being one for respondents who survived up to interval r and zero otherwise:

$$-\sum_{i=1}^{N}\left[\delta_{ir}{}^{w}\ \left(1 - \frac{\lambda_r\, e^{\beta ti}}{1+\ \lambda_r\, e^{\beta ti}}\right)^{-1}\left(\frac{e^{\beta ti}}{1+\ \lambda_r\, e^{\beta ti}} - \frac{\lambda_r\, e^{2\beta ti}}{(1+\ \lambda_r\, e^{\beta ti})^2}\right)\right]$$

Rearranging and cancelling we obtain:

$$\sum_{i=1}^{N}\left[\delta_{ir}{}^{w}\ \frac{e^{\beta ti}}{1+\ \lambda_r\, e^{\beta ti}}\ \right]$$

Differentiating (A.7) we obtain:

$$-\gamma_r$$

Therefore,

$$\frac{\partial \mathrm{G}}{\partial \lambda_r} = \sum_{i=1}^{N}\left[d_{ir}{}^{w}\ \left(\frac{1}{\lambda_r} - \frac{e^{\beta ti}}{1+\ \lambda_r\, e^{\beta ti}}\ \right) - \ \delta_{ir}{}^{w}\ \frac{e^{\beta ti}}{1+\ \lambda_r\, e^{\beta ti}} \right. \\ \left. -\ \gamma_r\right] \tag{A.8}$$

Setting the expression in (A.8) to zero and multiplying through by λ_r gives the following equation (due to complementary slackness condition of $\gamma_k \lambda_k = 0$ the final term in (A.8) drops out):

$$\frac{\partial \mathrm{G}}{\partial \lambda_r} = \sum_{i=1}^{N}\left[d_{ir}{}^{w}\ \left(1\ -\ \frac{\lambda_r\, e^{\beta ti}}{1+\ \lambda_r\, e^{\beta ti}}\ \right) - \ \delta_{ir}{}^{w}\ \frac{\lambda_r e^{\beta ti}}{1+\ \lambda_r\, e^{\beta ti}}\ \right] = 0$$

Rearranging we obtain:

$$\sum_{i=1}^{N} d_{ir}^{w} = \sum_{i=1}^{N}\frac{\lambda_r\, e^{\beta ti}}{1+\ \lambda_r\, e^{\beta ti}}(d_{ir}^{w} +\ \delta_{ir}^{w})\quad r = 1,2,\dots,K \tag{A.9}$$

By taking the λ_r outside the summation and replacing the notation for $\sum_{i=1}^{N} d_{ir}^{w}$ with n_r^{w} we obtain:

$$n_r^w = \lambda_r \sum_{i=1}^{N} \frac{e^{\beta ti}}{1 + \lambda_r e^{\beta ti}} (d_{ir}^w + \delta_{ir}^w) \quad r = 1,2, \dots, K$$

Rearranging to leave λ_r on the right side of the equation gives:

$$\frac{n_r^w}{\sum_{i=1}^{N} \frac{e^{\beta ti}}{1 + \lambda_r e^{\beta ti}} (d_{ir}^w + \delta_{ir}^w)} = \lambda_r$$

And finally multiplying the top and bottom of the left hand side of the equation by λ_r gives the estimating equation (1.7):

$$\frac{\lambda_r \; n_r^w}{\sum_{i=1}^{N} \frac{\lambda_r e^{\beta ti}}{1 + \lambda_r e^{\beta ti}} (d_{ir}^w + \delta_{ir}^w)} = \lambda_r \tag{A.10}$$

Next we differentiate (1.5) with respect to $\boldsymbol{\beta}$.

$$G(\beta) = \sum_{i=1}^{N} \left[\sum_{k=1}^{K} d_{ik}{}^{w} \ln\left(\frac{\lambda_k e^{\beta ti}}{1 + \lambda_k e^{\beta ti}} \right) + \delta_{ik}{}^{w} \ln\left(1 - \frac{\lambda_k e^{\beta ti}}{1 + \lambda_k e^{\beta ti}} \right) \right] \tag{A.11}$$

$$\sum_{i=1}^{N} \left[\sum_{k=1}^{K} d_{ik}{}^{w} \ln\left(\frac{\lambda_k e^{\beta ti}}{1 + \lambda_k e^{\beta ti}} \right) \right] \tag{A.12}$$

$$\sum_{i=1}^{N} \left[\sum_{k=1}^{K} \delta_{ik}{}^{w} \ln\left(1 - \frac{\lambda_k e^{\beta ti}}{1 + \lambda_k e^{\beta ti}} \right) \right] \tag{A.13}$$

Differentiating (A.12) we obtain the following equation:

$$\sum_{i=1}^{N} \sum_{k=1}^{K} \left[d_{ik}{}^{w} \left(\frac{\lambda_k e^{\beta ti}}{1 + \lambda_k e^{\beta ti}} \right)^{-1} \left(\frac{\lambda_k e^{\beta ti} \, ti}{1 + \lambda_k e^{\beta ti}} - \frac{(\lambda_k e^{\beta ti})^2 \, ti}{(1 + \lambda_k e^{\beta ti})^2} \right) \right]$$

Rearranging and cancelling we obtain:

$$\sum_{i=1}^{N} \sum_{k=1}^{K} \left[d_{ik}{}^{w} \, ti \left(1 - \frac{\lambda_k e^{\beta ti}}{1 + \lambda_k e^{\beta ti}} \right) \right]$$

Differentiating (A.13) we obtain the following equation:

$$\sum_{i=1}^{N}\sum_{k=1}^{K}\left[\delta_{ik}{}^{w}\left(1-\frac{\lambda_k\, e^{\beta ti}}{1+\ \lambda_k\, e^{\beta ti}}\right)^{-1}\ \left(\frac{\lambda_k\, e^{\beta ti}\ ti}{1+\ \lambda_k\, e^{\beta ti}}-\ \frac{(\lambda_k\, e^{\beta ti})^2\ ti}{(1+\ \lambda_k\, e^{\beta ti})^2}\ \right)\right]$$

Rearranging and cancelling we obtain:

$$-\sum_{i=1}^{N}\sum_{k=1}^{K}\left[\ \delta_{ik}{}^{w}\ \ ti\ \left(\frac{\lambda_k\, e^{\beta ti}}{1+\ \lambda_k\, e^{\beta ti}}\right)\right]$$

Therefore:

$$\frac{\partial \mathrm{G}}{\partial \beta}\ =\ \sum_{i=1}^{N} ti \sum_{k=1}^{K}\left[\ d_{ik}\ \left(1-\frac{\lambda_k\, e^{\beta ti}}{1+\ \lambda_k\, e^{\beta ti}}\right)-\ \ \delta_{ik}{}^{w}\ \left(\frac{\lambda_k\, e^{\beta ti}}{1+\ \lambda_k\, e^{\beta ti}}\right)\right]=0$$

After rearranging, we obtain the estimating equation in (1.8):

$$\begin{aligned}\frac{\partial \mathrm{G}}{\partial \beta}\ &=\ \sum_{i=1}^{N} ti \sum_{k=1}^{K}\left[\ d_{ik}\ \left(1-\frac{\lambda_k\, e^{\beta ti}}{1+\ \lambda_k\, e^{\beta ti}}\right)\right]\\ &=\ \sum_{i=1}^{N} ti \sum_{k=1}^{K}\left[\delta_{ik}{}^{w}\ \left(\frac{\lambda_k\, e^{\beta ti}}{1+\ \lambda_k\, e^{\beta ti}}\right)\right]\end{aligned} \qquad (A.14)$$

APPENDIX A2

The following Appendix builds on Appendix A1 to incorporate the initial bid level (the dichotomous choice or open ended contingent valuation task, task 1 in Chapter 1) dummy variables for the transition from uncertainty to certainly would not pay. In particular we show how differentiating our models differs when we add in the initial bid level dummy variables.

$$G_2 = \ ln\, L \ + \ \mu_0 \left(1 - \sum_j p_j\right) - \ \sum_j \mu_j p_j \ - \ \sum_{k=1}^{K} \gamma_k \lambda_k$$

The only changes are contained inside the hazard function (h_{ik}) which makes differentiating very simple.

$$\begin{aligned} ln\, G_2(\lambda_r) = \ & \sum_{i=1}^{N} \left[\sum_{k=1}^{K} d_{ir}{}^{w} \ \ln(h_{ik}) + \ \delta_{ir}{}^{w} \ln(1 - \ h_{ik}) \ \right] \\ & - \ \sum_{k=1}^{K} \gamma_k \lambda_r \end{aligned} \qquad \text{(A.15)}$$

Differentiating (A.15) w.r.t λ_r we obtain:

$$\frac{\partial \mathrm{G}_2}{\partial \lambda_r} = \sum_{i=1}^{N} \left[\left(\frac{d_{ir}{}^{w}}{\lambda_r} - \ h_{ik} \right) (d_{ir}{}^{w} + \ \delta_{ir}{}^{w}) \ - \gamma_r \right] \qquad \text{(A.16)}$$

Differentiating (A.15) w.r.t $\beta_{\boldsymbol{q}*}$ we obtain:

$$\frac{\partial \mathrm{G}}{\partial \beta_{\boldsymbol{q}*}} = \sum_{i=1}^{N} \begin{bmatrix} t_i \\ \vdots \\ q_7 \end{bmatrix} \sum_{k=1}^{K} [\, d_{ik}(1 - h_{ik}\,) - \delta_{ik}{}^{w} h_{ik}] \qquad \text{(A.17)}$$

Include initial bid level dummy variables for the transition from certainly would pay to uncertainty.

$$G_3 = \ ln\, L \ - \ \sum_{j=1}^{M+1} \gamma_j \phi_j$$

Taking logs gives:

$$ln\ \mathrm{G}_3(\phi_j) = \sum_{i=1}^{N}\left[\sum_{j=1}^{M+1} d_{ij}{}^{t}\ \ln(h^{c}{}_{ij}) + \delta_{ij}{}^{t}\ln(1 - h^{c}{}_{ij})\right] - \sum_{j=1}^{M+1}\gamma_j\phi_j. \tag{A.18}$$

Differentiating (A.18) w.r.t ϕ_j we obtain:

$$\frac{\partial \mathrm{G}_3}{\partial \phi_j} = \sum_{i=1}^{N}\left[\left(\frac{d_{ij}{}^{t}}{\phi_{j*}} - h^{c}{}_{ij}\right)(d_{ij}{}^{t} + \delta_{ij}{}^{t}) - \gamma_r\right] \tag{A.19}$$

Differentiating (A.18) w.r.t $\alpha_{\boldsymbol{q}}$ we obtain:

$$\frac{\partial \mathrm{G}}{\alpha_{\boldsymbol{q}}} = \sum_{i=1}^{N}\begin{bmatrix} q_0 \\ \vdots \\ q_7 \end{bmatrix}\sum_{j=1}^{M+1}[\ d_{ij}{}^{t}(1 - h^{c}{}_{ij}) - \delta_{ij}{}^{t}h^{c}{}_{ij}] \tag{A.20}$$

APPENDIX A3

We include here the contingent valuation survey used to collect the data for Chapter 1. Note that question 14 varies depending on whether the respondent was answering a dichotomous choice or open ended question.

Location of interview..

CONFIDENTIAL

-Date ☐☐ ☐☐ ☐☐

-Day (circle correct day)

1= MON 2= TUE 3=WED 4=THU 5=FRI

6= SAT 7=SUN

- Time interview started ☐☐ ☐☐

  (24 hour clock)

- Time interview ended ☐☐ ☐☐

  (24 hour clock)

- Weather conditions (circle the correct response)

(a)	Sunny	=	1	(c)	Dry	=	1
	Broken Cloud	=	2		Drizzle/Showers	=	2
	Overcast	=	3		Persistent rain	=	3
(b)	Hot (>20)	=	1	(d)	Calm	=	1
	Warm (15-20)	=	2		Breezy	=	2
	Cool (10-15)	=	3		Windy	=	3
	Cold (<10)	=	4				

- Tide Level (circle the correct response)

Low tide	=	1
Mid tide	=	2
High tide	=	3
Not known	=	4

- Is the sea (circle the correct response)

Rough = 1

Moderately Rough =2

Calm = 3

INTERVIEWER INSTRUCTIONS

1. Statements and questions to be read out are shown in bold type;

2. When recording answers circle the number of the appropriate response or fill in boxes as indicated;

3. If interviewing a family group you should aim to interview the head of household.

Hello, I am (GIVE NAME and show ID) **from the University of East Anglia we are conducting a survey regarding the beaches at Southwold. Would you mind answering a few questions, it will take about 10 minutes and any information you provide will be kept strictly confidential.**

If willing then proceed. If not then withdraw politely and make a note of the refusal on the tally sheet.

1. **Before I start can I just check if you live in the UK or not ?**

If answer = Yes, then proceed

If answer = No, then explain that the questionnaire is only applicable to UK citizens, ask for country of residence then withdraw politely and record this on the refusal tally sheet making a note of the country of residence.

I want to show you this map of Southwold (show card 1**). Now as you may know, here is the pier** (INDICATE)**, here is the North beach which you may know as Easton Bavents** (INDICATE) **and South of the pier is the Town Front beach** (INDICATE)

2. **Can you tell me how often you visit each of these areas** (show card 2) **First the North Beach at Easton Bavents (**get response**). The Town Front Beach south of the pier (**get response) **and finally the pier itself (**get response). (NOTE: for holidaymakers ensure they do not answer 1, 2, 3 etc - all answers refer to per year visits, therefore holidaymakers will usually be response 4 or 5)

	North Beach (Easton Bavents)	Town Front Beach (South)	The Pier
1 = I visit at least daily			

2 = I visit at least 3 times a week			
3 = I visit at least once a week			
4 = I visit about 10 days a year			
5 = I visit about 5 days a year			
6 = I visit about once a year			
7 = I visit less than once a year			
8 = I have never visited here before			

3. **On a typical trip would you visit more than one of these areas** (If yes circle all that apply for a typical visit)

0 = No, I would only visit one area on a typical visit (go to Qu.5)

1 = Yes, I typically visit more than one of these areas (go to Qu.4)

4. **And which of the areas are they?** (circle all that apply for a typical visit)

1 = North Beach (Easton Bavents)

2 = Town Front Beach

3 = The pier

0 = Not asked (said no to Qu.3)

5. **We want to find out where visitors come from, I am not going to ask for your full address but can you tell me your home postcode?** (note that we are not asking for their house number so they will not be receiving any mail)

(Get **full** postcode)

Alternatively, if you do not know your full post code could I have your approximate address ignoring the house number and street name (typical examples are area in a city and that city name [not just city name], or village or nearest town. In all cases also elicit the county name).

Village or Area within city

Nearest town or City

County

6. **Roughly how far away is it from your <u>home</u> address to here?**

One way distance in miles
Or distance in metres

7. **Did you set out from that <u>home</u> address today?**

Yes = 1 (go to Q.10)

No = 0 (go to Q.8)

8. **Are you staying with family or friends or in rented holiday accommodation here?**

1 = Staying with family/friends

2 = Staying in rented holiday accommodation

3 = Other (please specify) ....

9. **How far did you travel <u>today</u> to get here, just the one way distance?**

Distance in miles: <u>OR</u> distance in metres:

1 metre is approximately 1 yard

10. **How long did it take you to get here today?**

Hours Minutes

11. **When you come to Southwold beach which of these do you do often, sometimes or never** (SHOW CARD 3)

ACTIVITY	OFTEN	SOMETIMES	NEVER
1. Relaxing on the beach	1	2	3

2. Walking the dog	1	2	3
3. Other Walking	1	2	3
4. Picnicking	1	2	3
5. Fishing	1	2	3
6. Boating/sailing	1	2	3
7. Swimming/paddling/surfing	1	2	3
8. Going to restaurants / pubs /cafes	1	2	3
9. Visit local shops or arcades	1	2	3
10. Bird / wildlife watching	1	2	3

12. **Which of the above, or any other activity, do you feel is your main reason for visiting Southwold sea front today?**

Write activity code number (FROM CARD 3/Q.10) in the following box:

OR write in other main activity:

....

13. **How long in total will you spend on the seafront today?**

[ ][ ] Hours [ ][ ] Minutes

Don't know = 0

I now want to show you some information (show info card 1**). Here again is the map of Southwold** (indicate). **At present the sea defences along the coast here are old and in a poor state of repair. This photo shows the North Beach** (indicate upper left photo) **where you can see the stumps of the old wooden defences called groynes** (indicate). **At the Town Front Beach the wooden groynes are also in a poor state of repair (**indicate lower left photo**).**

This has resulted in the erosion of the beach. As you can see in these pictures the beach is very narrow at high tide (indicate left hand side pictures) **it's actually considerably narrower than it used to be.**

Government funding will ensure the sea wall is maintained to defend the properties in the town. However, additional defences could be put in place to extend the size of the beach. (show info card 1) **It is proposed that new sand will be brought in from the sea and added to the beaches. To stop future erosion new rock defences would be built at the North Beach** (indicate upper right photo) **and new timber groynes built at the Town Front Beach** (indicate lower right photo). **This will substantially increase the size of both beaches as shown in these photos** (indicate right hand side photos).

The scheme to enlarge the beach would result in additional costs. These costs would have to be met from extra general taxes as paid by your household on the everyday things you purchase.

In a moment I am going to ask you what is the <u>most</u> your household would be prepared to pay per year in extra general taxes to fund the beach enlargement. However, before you answer I want you to think about all of the following (SHOW CARD 4)**:**

1. **Irrespective of this scheme, the sea wall at Southwold will be maintained and the properties will be protected from flooding**

2. **There are alternative beaches which you could travel to;**

3. **And any money you would pay towards this scheme would not be available to you for other purchases.**

<u>Dichotomous Choice</u>

14. **So please tell me whether your household would be willing to pay £ ……….. per year for the scheme to enlarge the beaches at Southwold?**

1 = yes

2 = no

<u>Open Ended</u>

14. **So please tell me what is the most that your household would be willing to pay per year for the scheme to enlarge the beaches at Southwold?**

Answer £...per annum

15. **Now I realise that was a difficult question and that you may not be very certain of your answer. Take a look at this card. Can you tell me which of these statements best describes your feelings about paying £ ……….. for the enlarged beaches?** (circle appropriate number)

Definitely would pay this amount	Probably would pay this amount	Not sure if I would pay this amount	Probably would not pay this amount	Definitely would not pay this amount
1	2	3	4	5

16. **We are interested in finding out the amounts of money that you definitely would pay and those that you would definitely *not* pay for the enlarged beaches.**

Take a look at this list of money amounts (flick through valuation sheets to show the range of values**).**

IF *YES* (*NO)* TO QUESTION 14 (*DEFINITELY YES (NO))* TO QUESTION 15):

Just now you said that you would/(WOULD NOT) pay £ ……….. for the enlarged beaches. I'll indicate your answer by placing a tick in the "Definitely Yes"/("DEFINITELY NO") box next to that amount.

Now consider the higher/(LOWER) amounts on the list. (Pass list and clipboard to respondent**). Starting with £ ……….. (**next highest/(LOWEST) amount**), work down/(UP) the list considering each of these amounts in turn until you reach an amount that there's a possibility you would not/(WOULD) pay, however small. (**Allow respondent time to determine this amount**). Again, looking at the card decide which category best describes your response to that amount and tick the corresponding box on the list.**

Continue working down/(UP) the amounts on the list, ticking one box for each amount. Stop once you reach an amount that you definitely would not/(WOULD) pay.

IF UNSURE TO QUESTION 14 (PROBABLY/NOT SURE TO QUESTION 15):

Just now you said that you "probably would"/"uncertain whether would"/"would not" (answer to Question 15**) pay £ ……….. for the enlarged beaches. I'll indicate your answer by placing a tick in the "probably yes"/"not sure"/"probably no" box next to that amount.**

Now consider the next amount down on the list. Still looking at the card, if the amount was (£next highest amount) **which of the categories on the card best describes your response to that amount (**tick in appropriate box next to amount on valuation sheet**).**

Now it's your turn (Pass list and clipboard to respondent**). Work down the amounts on the list, ticking one box for each amount. Stop once you reach an amount that you DEFINITELY WOULD NOT PAY.**

Now I'd like you work up the amounts on the list. Starting at (£next lowest amount) **tick one box for each amount and stop once you reach an amount that you DEFINITELY WOULD PAY.**

17. (FOR RESPONDENTS WILLING TO STATE AN AMOUNT) **Why were you prepared to pay towards this scheme?**

..

NOW GO TO Q.19

18. (FOR RESPONDENTS NOT WILLING TO PAY ANYTHING) **Why were you not willing to pay for this scheme?**

....

19. **There is an alternative scheme** (show info card 2) **which is the same in all respects except that instead of timber groynes, rock defences would be used on the Town Front Beach** (indicate lower right hand side photo)

Thinking back to the previous scheme using wooden groynes. You said that the most you would <u>definitely</u> pay for that scheme was £X. Would you also definitely pay £X for the alternative scheme using rock groynes rather than timber groynes.

IF YES THEN ASK HIGHER AMOUNTS, STOP WHEN RESPONDENT IS NOT DEFINETELY SURE THEY WILL PAY THE AMOUNT

IF NO THEN ASK LOWER AMOUNTS STOP WHEN THE RESPONDENT IS NOT DEFFINETELY SURE THEY WILL PAY THE AMOUNT

(only tick the highest amount they would definitely pay)

For zero payers in first scheme: **Thinking back you said you would not pay for the timber groynes, would you change your answer for a scheme with rock defences?**

0 = No

1 = Yes (if so ask higher amounts and tick as appropriate)

20 . (FOR RESPONDENTS WILLING TO STATE AN AMOUNT) **What is the main reason for your answer?**

..

21. **If Southwold had a bigger beach as described would you visit more often, less often or about the same?**

0 = less often

1 = about the same (go to Qu.23)

2 = more often

22. **So roughly how many more / less times would you visit each year?**

... more / less times per year

23. **Now from this card** (SHOW CARD 6) **could you tell me which letter corresponds to your age group?**

LETTER	AGE IN YEARS
A	0-4
B	5-9
C	10-15
D	16-19
E	20-29
F	30-39
G	40-49
H	50-59
I	60-69
J	70 +

Age group (Letter please) ☐

24. **Again using the same card** (SHOW CARD 6) **could you tell me how many people in your household fall into each age category.**

LETTER	AGE IN YEARS	NUMBER IN HOUSEHOLD
A	0-4	
B	5-9	

C	10-15	
D	16-19	
E	20-29	
F	30-39	
G	40-49	
H	50-59	
I	60-69	
J	70 +	

25. **How would you classify your employment status:** (SHOW CARD 7)

1 = Full time employed

2 = Full time self employed

3 = employed part time

4 = unemployed

5 = on a government training scheme

6 = retired

7 = homemaker

8 = student

9 = other (please specify)

26. **Could you please tell me which of these letters, A to I** (SHOW CARD 8)**, best describes your total household income (pre-tax including state benefits, pensions, interest on investments, etc.).** If necessary please stress:

a. All answers are completely anonymous and confidential;

b. The importance of getting an accurate reply to this question - we need to account for the fact that ability to pay clearly influences responses to tax and entrance fee questions.

Letter	**Total Household Income (£)**	
	Yearly (£)	**Weekly (£)**
A	**0-4,999**	**0-96**
B	**5,000-7,499**	**96-144**
C	**7,500-9,999**	**144-192**

D	**10,000-14,999**	**192-288**
E	**15,000-19,999**	**288-385**
F	**20,000-29,999**	**385-577**
G	**30,000-39,999**	**577-769**
H	**40,000-49,999**	**769-962**
I	**50,000+**	**962+**

INCOME CATEGORY LETTER: ☐

27. **Is anyone in your household a member of any of the following groups?**

(SHOW CARD 9 Circle all that apply)

1 = Any sports club

2 = Any church/religious/charity group

3 = Lions/Rotary etc.

4 = Women's Institute

5 = Sailing/ Boating Club

6 = Angling Club

7 = Beach/coastline campaign group

8 = National Trust

9 = RSPB

10 = Greenpeace/Friends of the Earth

11 = World Wide Fund for Nature

12 = Other local or County nature trust, society or volunteers

13 = Other social group (please specify)____________________________

14 = Other not covered above (please specify)____________________________

THANK YOU VERY MUCH FOR YOUR HELP

RESPONDENT SEX (circle number)

Female = 0

Male = 1

APPENDIX A4

We include here the contingent valuation payment card used to collect the data for Chapter 1.

Value	Definitely Yes	Probably Yes	Not Sure	Probably No	Definitely No	Value	Definitely Yes	Probably Yes	Not Sure	Probably No	Definitely No
£1	☐	☐	☐	☐	☐	£51	☐	☐	☐	☐	☐
£2	☐	☐	☐	☐	☐	£52	☐	☐	☐	☐	☐
£3	☐	☐	☐	☐	☐	£53	☐	☐	☐	☐	☐
£4	☐	☐	☐	☐	☐	£54	☐	☐	☐	☐	☐
£5	☐	☐	☐	☐	☐	£55	☐	☐	☐	☐	☐
£6	☐	☐	☐	☐	☐	£56	☐	☐	☐	☐	☐
£7	☐	☐	☐	☐	☐	£57	☐	☐	☐	☐	☐
£8	☐	☐	☐	☐	☐	£58	☐	☐	☐	☐	☐
£9	☐	☐	☐	☐	☐	£59	☐	☐	☐	☐	☐
£10	☐	☐	☐	☐	☐	£60	☐	☐	☐	☐	☐
£11	☐	☐	☐	☐	☐	£61	☐	☐	☐	☐	☐
£12	☐	☐	☐	☐	☐	£62	☐	☐	☐	☐	☐
£13	☐	☐	☐	☐	☐	£63	☐	☐	☐	☐	☐
£14	☐	☐	☐	☐	☐	£64	☐	☐	☐	☐	☐
£15	☐	☐	☐	☐	☐	£65	☐	☐	☐	☐	☐
£16	☐	☐	☐	☐	☐	£66	☐	☐	☐	☐	☐
£17	☐	☐	☐	☐	☐	£67	☐	☐	☐	☐	☐
£18	☐	☐	☐	☐	☐	£68	☐	☐	☐	☐	☐
£19	☐	☐	☐	☐	☐	£69	☐	☐	☐	☐	☐
£20	☐	☐	☐	☐	☐	£70	☐	☐	☐	☐	☐
£21	☐	☐	☐	☐	☐	£71	☐	☐	☐	☐	☐
£22	☐	☐	☐	☐	☐	£72	☐	☐	☐	☐	☐
£23	☐	☐	☐	☐	☐	£73	☐	☐	☐	☐	☐
£24	☐	☐	☐	☐	☐	£74	☐	☐	☐	☐	☐
£25	☐	☐	☐	☐	☐	£75	☐	☐	☐	☐	☐
£26	☐	☐	☐	☐	☐	£76	☐	☐	☐	☐	☐
£27	☐	☐	☐	☐	☐	£77	☐	☐	☐	☐	☐
£28	☐	☐	☐	☐	☐	£78	☐	☐	☐	☐	☐
£29	☐	☐	☐	☐	☐	£79	☐	☐	☐	☐	☐
£30	☐	☐	☐	☐	☐	£80	☐	☐	☐	☐	☐
£31	☐	☐	☐	☐	☐	£81	☐	☐	☐	☐	☐
£32	☐	☐	☐	☐	☐	£82	☐	☐	☐	☐	☐
£33	☐	☐	☐	☐	☐	£83	☐	☐	☐	☐	☐
£34	☐	☐	☐	☐	☐	£84	☐	☐	☐	☐	☐
£35	☐	☐	☐	☐	☐	£85	☐	☐	☐	☐	☐
£36	☐	☐	☐	☐	☐	£86	☐	☐	☐	☐	☐
£37	☐	☐	☐	☐	☐	£87	☐	☐	☐	☐	☐
£38	☐	☐	☐	☐	☐	£88	☐	☐	☐	☐	☐
£39	☐	☐	☐	☐	☐	£89	☐	☐	☐	☐	☐
£40	☐	☐	☐	☐	☐	£90	☐	☐	☐	☐	☐
£41	☐	☐	☐	☐	☐	£91	☐	☐	☐	☐	☐
£42	☐	☐	☐	☐	☐	£92	☐	☐	☐	☐	☐
£43	☐	☐	☐	☐	☐	£93	☐	☐	☐	☐	☐
£44	☐	☐	☐	☐	☐	£94	☐	☐	☐	☐	☐
£45	☐	☐	☐	☐	☐	£95	☐	☐	☐	☐	☐
£46	☐	☐	☐	☐	☐	£96	☐	☐	☐	☐	☐
£47	☐	☐	☐	☐	☐	£97	☐	☐	☐	☐	☐
£48	☐	☐	☐	☐	☐	£98	☐	☐	☐	☐	☐
£49	☐	☐	☐	☐	☐	£99	☐	☐	☐	☐	☐
£50	☐	☐	☐	☐	☐	£100	☐	☐	☐	☐	☐

Value	Definitely Yes	Probably Yes	Not Sure	Probably No	Definitely No	Value	Definitely Yes	Probably Yes	Not Sure	Probably No	Definitely No
£101	☐	☐	☐	☐	☐	£151	☐	☐	☐	☐	☐
£102	☐	☐	☐	☐	☐	£152	☐	☐	☐	☐	☐
£103	☐	☐	☐	☐	☐	£153	☐	☐	☐	☐	☐
£ 104	☐	☐	☐	☐	☐	£154	☐	☐	☐	☐	☐
£ 105	☐	☐	☐	☐	☐	£155	☐	☐	☐	☐	☐
£ 106	☐	☐	☐	☐	☐	£156	☐	☐	☐	☐	☐
£ 107	☐	☐	☐	☐	☐	£157	☐	☐	☐	☐	☐
£ 108	☐	☐	☐	☐	☐	£158	☐	☐	☐	☐	☐
£ 109	☐	☐	☐	☐	☐	£159	☐	☐	☐	☐	☐
£110	☐	☐	☐	☐	☐	£160	☐	☐	☐	☐	☐
£111	☐	☐	☐	☐	☐	£161	☐	☐	☐	☐	☐
£112	☐	☐	☐	☐	☐	£162	☐	☐	☐	☐	☐
£113	☐	☐	☐	☐	☐	£163	☐	☐	☐	☐	☐
£114	☐	☐	☐	☐	☐	£164	☐	☐	☐	☐	☐
£115	☐	☐	☐	☐	☐	£165	☐	☐	☐	☐	☐
£116	☐	☐	☐	☐	☐	£166	☐	☐	☐	☐	☐
£117	☐	☐	☐	☐	☐	£167	☐	☐	☐	☐	☐
£118	☐	☐	☐	☐	☐	£168	☐	☐	☐	☐	☐
£119	☐	☐	☐	☐	☐	£169	☐	☐	☐	☐	☐
£120	☐	☐	☐	☐	☐	£170	☐	☐	☐	☐	☐
£121	☐	☐	☐	☐	☐	£171	☐	☐	☐	☐	☐
£122	☐	☐	☐	☐	☐	£172	☐	☐	☐	☐	☐
£123	☐	☐	☐	☐	☐	£173	☐	☐	☐	☐	☐
£124	☐	☐	☐	☐	☐	£174	☐	☐	☐	☐	☐
£125	☐	☐	☐	☐	☐	£175	☐	☐	☐	☐	☐
£126	☐	☐	☐	☐	☐	£176	☐	☐	☐	☐	☐
£127	☐	☐	☐	☐	☐	£177	☐	☐	☐	☐	☐
£128	☐	☐	☐	☐	☐	£178	☐	☐	☐	☐	☐
£129	☐	☐	☐	☐	☐	£179	☐	☐	☐	☐	☐
£130	☐	☐	☐	☐	☐	£180	☐	☐	☐	☐	☐
£131	☐	☐	☐	☐	☐	£181	☐	☐	☐	☐	☐
£132	☐	☐	☐	☐	☐	£182	☐	☐	☐	☐	☐
£133	☐	☐	☐	☐	☐	£183	☐	☐	☐	☐	☐
£134	☐	☐	☐	☐	☐	£184	☐	☐	☐	☐	☐
£135	☐	☐	☐	☐	☐	£185	☐	☐	☐	☐	☐
£136	☐	☐	☐	☐	☐	£186	☐	☐	☐	☐	☐
£137	☐	☐	☐	☐	☐	£187	☐	☐	☐	☐	☐
£138	☐	☐	☐	☐	☐	£188	☐	☐	☐	☐	☐
£139	☐	☐	☐	☐	☐	£189	☐	☐	☐	☐	☐
£140	☐	☐	☐	☐	☐	£190	☐	☐	☐	☐	☐
£141	☐	☐	☐	☐	☐	£191	☐	☐	☐	☐	☐
£142	☐	☐	☐	☐	☐	£192	☐	☐	☐	☐	☐
£143	☐	☐	☐	☐	☐	£193	☐	☐	☐	☐	☐
£144	☐	☐	☐	☐	☐	£194	☐	☐	☐	☐	☐
£145	☐	☐	☐	☐	☐	£195	☐	☐	☐	☐	☐
£146	☐	☐	☐	☐	☐	£196	☐	☐	☐	☐	☐
£147	☐	☐	☐	☐	☐	£197	☐	☐	☐	☐	☐
£148	☐	☐	☐	☐	☐	£198	☐	☐	☐	☐	☐
£149	☐	☐	☐	☐	☐	£199	☐	☐	☐	☐	☐
£150	☐	☐	☐	☐	☐	£200	☐	☐	☐	☐	☐

Value	Definitely Yes	Probably Yes	Not Sure	Probably No	Definitely No	Value	Definitely Yes	Probably Yes	Not Sure	Probably No	Definitely No
£201	☐	☐	☐	☐	☐	£251	☐	☐	☐	☐	☐
£202	☐	☐	☐	☐	☐	£252	☐	☐	☐	☐	☐
£203	☐	☐	☐	☐	☐	£253	☐	☐	☐	☐	☐
£ 204	☐	☐	☐	☐	☐	£254	☐	☐	☐	☐	☐
£ 205	☐	☐	☐	☐	☐	£255	☐	☐	☐	☐	☐
£ 206	☐	☐	☐	☐	☐	£256	☐	☐	☐	☐	☐
£ 207	☐	☐	☐	☐	☐	£257	☐	☐	☐	☐	☐
£ 208	☐	☐	☐	☐	☐	£258	☐	☐	☐	☐	☐
£ 209	☐	☐	☐	☐	☐	£259	☐	☐	☐	☐	☐
£210	☐	☐	☐	☐	☐	£260	☐	☐	☐	☐	☐
£211	☐	☐	☐	☐	☐	£261	☐	☐	☐	☐	☐
£212	☐	☐	☐	☐	☐	£262	☐	☐	☐	☐	☐
£213	☐	☐	☐	☐	☐	£263	☐	☐	☐	☐	☐
£214	☐	☐	☐	☐	☐	£264	☐	☐	☐	☐	☐
£215	☐	☐	☐	☐	☐	£265	☐	☐	☐	☐	☐
£216	☐	☐	☐	☐	☐	£266	☐	☐	☐	☐	☐
£217	☐	☐	☐	☐	☐	£267	☐	☐	☐	☐	☐
£218	☐	☐	☐	☐	☐	£268	☐	☐	☐	☐	☐
£219	☐	☐	☐	☐	☐	£269	☐	☐	☐	☐	☐
£220	☐	☐	☐	☐	☐	£270	☐	☐	☐	☐	☐
£221	☐	☐	☐	☐	☐	£271	☐	☐	☐	☐	☐
£222	☐	☐	☐	☐	☐	£272	☐	☐	☐	☐	☐
£223	☐	☐	☐	☐	☐	£273	☐	☐	☐	☐	☐
£224	☐	☐	☐	☐	☐	£274	☐	☐	☐	☐	☐
£225	☐	☐	☐	☐	☐	£275	☐	☐	☐	☐	☐
£226	☐	☐	☐	☐	☐	£276	☐	☐	☐	☐	☐
£227	☐	☐	☐	☐	☐	£277	☐	☐	☐	☐	☐
£228	☐	☐	☐	☐	☐	£278	☐	☐	☐	☐	☐
£229	☐	☐	☐	☐	☐	£279	☐	☐	☐	☐	☐
£230	☐	☐	☐	☐	☐	£280	☐	☐	☐	☐	☐
£231	☐	☐	☐	☐	☐	£281	☐	☐	☐	☐	☐
£232	☐	☐	☐	☐	☐	£282	☐	☐	☐	☐	☐
£233	☐	☐	☐	☐	☐	£283	☐	☐	☐	☐	☐
£234	☐	☐	☐	☐	☐	£284	☐	☐	☐	☐	☐
£235	☐	☐	☐	☐	☐	£285	☐	☐	☐	☐	☐
£236	☐	☐	☐	☐	☐	£286	☐	☐	☐	☐	☐
£237	☐	☐	☐	☐	☐	£287	☐	☐	☐	☐	☐
£238	☐	☐	☐	☐	☐	£288	☐	☐	☐	☐	☐
£239	☐	☐	☐	☐	☐	£289	☐	☐	☐	☐	☐
£240	☐	☐	☐	☐	☐	£290	☐	☐	☐	☐	☐
£241	☐	☐	☐	☐	☐	£291	☐	☐	☐	☐	☐
£242	☐	☐	☐	☐	☐	£292	☐	☐	☐	☐	☐
£243	☐	☐	☐	☐	☐	£293	☐	☐	☐	☐	☐
£244	☐	☐	☐	☐	☐	£294	☐	☐	☐	☐	☐
£245	☐	☐	☐	☐	☐	£295	☐	☐	☐	☐	☐
£246	☐	☐	☐	☐	☐	£296	☐	☐	☐	☐	☐
£247	☐	☐	☐	☐	☐	£297	☐	☐	☐	☐	☐
£248	☐	☐	☐	☐	☐	£298	☐	☐	☐	☐	☐
£249	☐	☐	☐	☐	☐	£299	☐	☐	☐	☐	☐
£250	☐	☐	☐	☐	☐	£300	☐	☐	☐	☐	☐

Value	Definitely Yes	Probably Yes	Not Sure	Probably No	Definitely No	Value	Definitely Yes	Probably Yes	Not Sure	Probably No	Definitely No
£301	☐	☐	☐	☐	☐	£351	☐	☐	☐	☐	☐
£302	☐	☐	☐	☐	☐	£352	☐	☐	☐	☐	☐
£303	☐	☐	☐	☐	☐	£353	☐	☐	☐	☐	☐
£ 304	☐	☐	☐	☐	☐	£354	☐	☐	☐	☐	☐
£ 305	☐	☐	☐	☐	☐	£355	☐	☐	☐	☐	☐
£ 306	☐	☐	☐	☐	☐	£356	☐	☐	☐	☐	☐
£ 307	☐	☐	☐	☐	☐	£357	☐	☐	☐	☐	☐
£ 308	☐	☐	☐	☐	☐	£358	☐	☐	☐	☐	☐
£ 309	☐	☐	☐	☐	☐	£359	☐	☐	☐	☐	☐
£310	☐	☐	☐	☐	☐	£360	☐	☐	☐	☐	☐
£311	☐	☐	☐	☐	☐	£361	☐	☐	☐	☐	☐
£312	☐	☐	☐	☐	☐	£362	☐	☐	☐	☐	☐
£313	☐	☐	☐	☐	☐	£363	☐	☐	☐	☐	☐
£314	☐	☐	☐	☐	☐	£364	☐	☐	☐	☐	☐
£315	☐	☐	☐	☐	☐	£365	☐	☐	☐	☐	☐
£316	☐	☐	☐	☐	☐	£366	☐	☐	☐	☐	☐
£317	☐	☐	☐	☐	☐	£367	☐	☐	☐	☐	☐
£318	☐	☐	☐	☐	☐	£368	☐	☐	☐	☐	☐
£319	☐	☐	☐	☐	☐	£369	☐	☐	☐	☐	☐
£320	☐	☐	☐	☐	☐	£370	☐	☐	☐	☐	☐
£321	☐	☐	☐	☐	☐	£371	☐	☐	☐	☐	☐
£322	☐	☐	☐	☐	☐	£372	☐	☐	☐	☐	☐
£323	☐	☐	☐	☐	☐	£373	☐	☐	☐	☐	☐
£324	☐	☐	☐	☐	☐	£374	☐	☐	☐	☐	☐
£325	☐	☐	☐	☐	☐	£375	☐	☐	☐	☐	☐
£326	☐	☐	☐	☐	☐	£376	☐	☐	☐	☐	☐
£327	☐	☐	☐	☐	☐	£377	☐	☐	☐	☐	☐
£328	☐	☐	☐	☐	☐	£378	☐	☐	☐	☐	☐
£329	☐	☐	☐	☐	☐	£379	☐	☐	☐	☐	☐
£330	☐	☐	☐	☐	☐	£380	☐	☐	☐	☐	☐
£331	☐	☐	☐	☐	☐	£381	☐	☐	☐	☐	☐
£332	☐	☐	☐	☐	☐	£382	☐	☐	☐	☐	☐
£333	☐	☐	☐	☐	☐	£383	☐	☐	☐	☐	☐
£334	☐	☐	☐	☐	☐	£384	☐	☐	☐	☐	☐
£335	☐	☐	☐	☐	☐	£385	☐	☐	☐	☐	☐
£336	☐	☐	☐	☐	☐	£386	☐	☐	☐	☐	☐
£337	☐	☐	☐	☐	☐	£387	☐	☐	☐	☐	☐
£338	☐	☐	☐	☐	☐	£388	☐	☐	☐	☐	☐
£339	☐	☐	☐	☐	☐	£389	☐	☐	☐	☐	☐
£340	☐	☐	☐	☐	☐	£390	☐	☐	☐	☐	☐
£341	☐	☐	☐	☐	☐	£391	☐	☐	☐	☐	☐
£342	☐	☐	☐	☐	☐	£392	☐	☐	☐	☐	☐
£343	☐	☐	☐	☐	☐	£393	☐	☐	☐	☐	☐
£344	☐	☐	☐	☐	☐	£394	☐	☐	☐	☐	☐
£345	☐	☐	☐	☐	☐	£395	☐	☐	☐	☐	☐
£346	☐	☐	☐	☐	☐	£396	☐	☐	☐	☐	☐
£347	☐	☐	☐	☐	☐	£397	☐	☐	☐	☐	☐
£348	☐	☐	☐	☐	☐	£398	☐	☐	☐	☐	☐
£349	☐	☐	☐	☐	☐	£399	☐	☐	☐	☐	☐
£350	☐	☐	☐	☐	☐	£400	☐	☐	☐	☐	☐

Value	Definitely Yes	Probably Yes	Not Sure	Probably No	Definitely No	Value	Definitely Yes	Probably Yes	Not Sure	Probably No	Definitely No
£401	☐	☐	☐	☐	☐	£451	☐	☐	☐	☐	☐
£402	☐	☐	☐	☐	☐	£452	☐	☐	☐	☐	☐
£403	☐	☐	☐	☐	☐	£453	☐	☐	☐	☐	☐
£ 404	☐	☐	☐	☐	☐	£454	☐	☐	☐	☐	☐
£ 405	☐	☐	☐	☐	☐	£455	☐	☐	☐	☐	☐
£ 406	☐	☐	☐	☐	☐	£456	☐	☐	☐	☐	☐
£ 407	☐	☐	☐	☐	☐	£457	☐	☐	☐	☐	☐
£ 408	☐	☐	☐	☐	☐	£458	☐	☐	☐	☐	☐
£ 409	☐	☐	☐	☐	☐	£459	☐	☐	☐	☐	☐
£410	☐	☐	☐	☐	☐	£460	☐	☐	☐	☐	☐
£411	☐	☐	☐	☐	☐	£461	☐	☐	☐	☐	☐
£412	☐	☐	☐	☐	☐	£462	☐	☐	☐	☐	☐
£413	☐	☐	☐	☐	☐	£463	☐	☐	☐	☐	☐
£414	☐	☐	☐	☐	☐	£464	☐	☐	☐	☐	☐
£415	☐	☐	☐	☐	☐	£465	☐	☐	☐	☐	☐
£416	☐	☐	☐	☐	☐	£466	☐	☐	☐	☐	☐
£417	☐	☐	☐	☐	☐	£467	☐	☐	☐	☐	☐
£418	☐	☐	☐	☐	☐	£468	☐	☐	☐	☐	☐
£419	☐	☐	☐	☐	☐	£469	☐	☐	☐	☐	☐
£420	☐	☐	☐	☐	☐	£470	☐	☐	☐	☐	☐
£421	☐	☐	☐	☐	☐	£471	☐	☐	☐	☐	☐
£422	☐	☐	☐	☐	☐	£472	☐	☐	☐	☐	☐
£423	☐	☐	☐	☐	☐	£473	☐	☐	☐	☐	☐
£424	☐	☐	☐	☐	☐	£474	☐	☐	☐	☐	☐
£425	☐	☐	☐	☐	☐	£475	☐	☐	☐	☐	☐
£426	☐	☐	☐	☐	☐	£476	☐	☐	☐	☐	☐
£427	☐	☐	☐	☐	☐	£477	☐	☐	☐	☐	☐
£428	☐	☐	☐	☐	☐	£478	☐	☐	☐	☐	☐
£429	☐	☐	☐	☐	☐	£479	☐	☐	☐	☐	☐
£430	☐	☐	☐	☐	☐	£480	☐	☐	☐	☐	☐
£431	☐	☐	☐	☐	☐	£481	☐	☐	☐	☐	☐
£432	☐	☐	☐	☐	☐	£482	☐	☐	☐	☐	☐
£433	☐	☐	☐	☐	☐	£483	☐	☐	☐	☐	☐
£434	☐	☐	☐	☐	☐	£484	☐	☐	☐	☐	☐
£435	☐	☐	☐	☐	☐	£485	☐	☐	☐	☐	☐
£436	☐	☐	☐	☐	☐	£486	☐	☐	☐	☐	☐
£437	☐	☐	☐	☐	☐	£487	☐	☐	☐	☐	☐
£438	☐	☐	☐	☐	☐	£488	☐	☐	☐	☐	☐
£439	☐	☐	☐	☐	☐	£489	☐	☐	☐	☐	☐
£440	☐	☐	☐	☐	☐	£490	☐	☐	☐	☐	☐
£441	☐	☐	☐	☐	☐	£491	☐	☐	☐	☐	☐
£442	☐	☐	☐	☐	☐	£492	☐	☐	☐	☐	☐
£443	☐	☐	☐	☐	☐	£493	☐	☐	☐	☐	☐
£444	☐	☐	☐	☐	☐	£494	☐	☐	☐	☐	☐
£445	☐	☐	☐	☐	☐	£495	☐	☐	☐	☐	☐
£446	☐	☐	☐	☐	☐	£496	☐	☐	☐	☐	☐
£447	☐	☐	☐	☐	☐	£497	☐	☐	☐	☐	☐
£448	☐	☐	☐	☐	☐	£498	☐	☐	☐	☐	☐
£449	☐	☐	☐	☐	☐	£499	☐	☐	☐	☐	☐
£450	☐	☐	☐	☐	☐	£500	☐	☐	☐	☐	☐

APPENDIX B1

The decision problems of purchasers of ecosystem services can be modelled as budget constrained, target constrained or profit maximising, depending on the motivations of the buyers. Moreover, the buyers' problems can be included in a broad range of methods available in the conservation biology literature. Two key problem designs have taken prominence in that literature: the species set covering problem (SCP), and the species maximal covering problem (MCP) (Williams, et al. 2005). In the species SCP the objective is to select the minimum number of land parcels (or area) whilst selecting at least one land parcel containing each species or other features. In the MCP the objective is to maximise the number of species (or other features) represented in the solution whilst setting a limit on the number of land parcels selected. Such models can be used instead of SCP where appropriate and ReVelle, Williams and Boland (2002) give a good introduction to how these models are used in the reserve selection literature as well as grounding the subject in terms of a common problem (facility location) in operations research.

Here, as in the main text for the Chapter 2, we focus on the species SCP. In particular we show methods for expanding our framework of methods to be able to include target constrained and profit maximising buyers, highlighting a method (tangent line approximation) to deal with the inherent non-linearity of profit maximising problems.

As a reminder, a species set covering problem can be represented in the following model:

$$\min_{x} \quad \sum_{j=1}^{N} x_j \tag{B1}$$

$$s.t. \quad \sum_{j=1}^{N} a_{ij}x_j \geq 1 \qquad i = 1,2,\dots,m$$

$$x_j \in \{0,1\} \qquad j = 1,2,\dots,N$$

where N is the number of land parcels and m is the number of species. If $a_{ij} = 1$ then species i is present at land parcel j and 0 otherwise and if x_j = 1 then land parcel j is selected. The species SCP objective (B1) minimises the number of land parcels selected while the constraint ensures that each species is represented at least once.

Target-constrained buyers

A number of potential purchasers for ecosystem services are more concerned with achieving a certain target level rather than spending a budget, as an example let us imagine a national government concerned about meeting their carbon reduction targets[41]. The primary consideration in meeting ecosystem service targets is to keep costs as low as possible, therefore the objectives of such buyers can be modelled using a variation of the species SCP. Specifically, instead of simply minimising the number of land parcels or area of land parcels selected we can minimise the total cost of land parcels selected whilst meeting both the species representation constraint and the target constraint (Williams, et al. 2005). For example in the following objective:

$$\min_{x} \quad \sum_{j=1}^{N} c_j x_j \tag{B2}$$

$$\begin{aligned} s.t. \ & \sum_{j=1}^{N} a_{ij} x_j \geq 1 && i = 1,2,\dots,m \\ & \sum_{j=1}^{N} b_j x_j \geq \bar{T} \\ & x_j \in \{0,1\} && j = 1,2,\dots,N \end{aligned}$$

Where c_j is the cost of land parcel j, b_j is the benefit of land parcel j, and $\bar{T}$ is the carbon reduction target. As the buyers' problems can be represented in a linear form they can easily be included within the framework of methods set out in Chapter 2.

Profit maximising buyers

Other buyers' decision problems may be better represented by a profit maximising problem, in other words the buyers are looking to maximise the difference between

[41] A further example of a potential purchaser prioritising a target at the minimum cost would be the offsetting of environmental harmful activities in one locality by purchasing ecosystem services elsewhere. In that situation, it is easy to imagine that the motivation of the developers can again to be modelled by meeting a target of ecosystem services for the minimum cost.

their costs and benefits. To illustrate, let us imagine a water company paying for upstream catchment management. Measuring the water companies benefits could be represented using a number of metrics, for example, cubic metres of clean water or reduction in a particular pollutant entering the watercourse, however to calculate the difference in the costs and benefits requires the benefits to be measured in monetary terms. The relationship between monetary benefits and ecosystem service benefits need not be linear, indeed by assuming diminishing marginal benefits, the benefits are inherently non-linear. To fix ideas, imagine a water company has a water abstraction point downstream to a number of farms, and agricultural activity from these farms leads to the runoff of a variety of pollutants. We can imagine that the benefit to the water company from each upstream farm changing to an alternative land-management practice is independent but that the monetary benefits are dependent on the quantity of farms that have already converted to an alternative land-management practice (from the classifications of benefits in the main text this would be an example of spatial interdependence - quantity). For example, if the level of ecosystem services have already been increased such that the water quality of the river meets drinking water standards then there is minimal benefit to the water company from a further land-management change. Furthermore, if the water quality of the river is just above drinking water standards then the water company may have relatively cheap methods for dealing with low levels of pollutants, such as dilution; however, if the quantity is a long way from drinking water standards, the water company may have use expensive methods of cleaning the water, such as active carbon techniques.

A profit maximising buyer's decision problem can be represented in the following:

$$\max_{x} \quad \sum_{j=1}^{N} f(b_j x_j) - \ (c_j x_j) \tag{B3}$$

$$x_j \in \{0,1\} \qquad\qquad j = 1,2,\dots,N$$

where f describes the relationship that the benefits have to money. By maximising the difference between costs and benefits the buyer will select all the land parcels that provide a positive contribution to overall profit.

An important aspect of the profit maximising objective is the non-linearity in the benefits of the buyer, $f(b_j x_j)$. Assuming diminishing marginal financial benefits leads to a difficult combinatorial problem to solve. To the best of the authors' knowledge little work has been done on these problems within the ecosystem service literature, however, similar problems have been studied in other areas. For example, the competitive facility location problem uses spatial interactive models to study the best place to locate new facilities dependent on the location of other facilities and the location of customers (Aboolian, et al. 2007,2009). Aboolian et al. (2009) show that incorporating non-linear spatially dependent benefits into objectives can be solved by off-the-shelf linear optimisation programmes by using the Tangent Line Approximation (TLA) technique. The TLA techniques forms a piece-wise approximation of any non-decreasing concave function which goes through the origin and is a twice differentiable. This type of function is commonly used to represent diminishing marginal benefits.

Tangent line approximation (TLA)

The TLA procedure is based on the piece-wise approximation of a non-separable concave objective function, for further details see (Aboolian, et al. 2007,2009).

Let $f(b_j x_j)$ be a concave, non-decreasing and twice-differentiable function with $f(0) = 0$. The aim of the tangent line approximation technique is to create a piecewise linear approximation $f^{\alpha}(b_j x_j)$ with α the bound on the error of the approximation, such that $f(b_j x_j) \leq f^{\alpha}(b_j x_j)$. The benefits function and the piecewise linear approximation are presented in Figure B1.

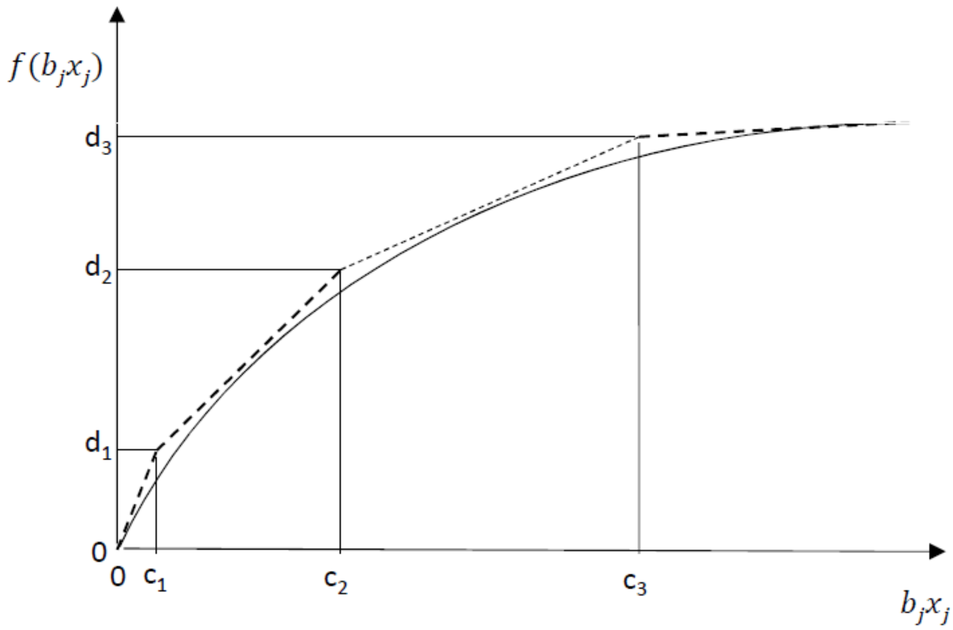

Figure B1. A concave function with a piecewise linear function approximated using the TLA technique.

The dotted line in Figure B1 represents the piecewise linear approximation of the concave benefits function. The approximation identifies L break points along the $b_j x_j$ axis: $c_1, c_2, \dots, c_L$, and L corresponding points along the $f(b_j x_j)$ axis: $d_1, d_2, \dots, d_L$.

To describe the TLA technique we make us of the following notation: each line segment is indexed by l with $l = 1, \dots, L$, in addition, s_l represents the slope of line segment l, the starting point of segment l is represented by c_l and the end point of a segment c_{l+1}.

To represent the functional form of the profit maximising purchasers diminishing benefits curve we use the following:

$$f(b_j x_j) = \frac{b_j x_j}{A + b_j x_j} \tag{B4}$$

where A denotes a constant.

The derivative of equation B1 is:

$$f'(b_j x_j) \ = \ \frac{A}{\left(A + b_j x_j\right)^2} \tag{B5}$$

Step 1: $l = 1,\ c_0 = d_0 = 0,\ s_1 = f\,'(0)$

We start by setting $f^{\alpha}(b_j x_j) = 0$ and using the point 0 as the starting point for the first segment. Furthermore, we set the slope s_1 of the first segment equal to $f'(0)$. To find the end point c_1 of the first line segment we need to find the value for $b_j x_j$ where relative error $\left(\frac{f^{\alpha}(b_j x_j) - f(b_j x_j)}{f(b_j x_j)}\right) = \alpha$. To do that we calculate the point $f^{\alpha}(c_1)$ on the ray originating at 0 and with the slope s_1 that gives a relative error of α:

$$s_0 + s_1 c_1 = f(c_1)(1 + \alpha) \tag{B6}$$

where s_0 is the intercept and for the first segment $s_0 = 0$. Equation B6 can be re-written as:

$$s_0 + \ s_1 c_1 = \frac{c_1}{A + \ c_1}(1 + \alpha) \tag{B7}$$

And the value of c_1 can be found by solving the following quadratic equation:

$$c_1 = \frac{-(As_1 + s_0 - 1 - \alpha) \ \pm \sqrt{(s_0 + As_1 - 1 - \alpha)^2 - 4(s_1)(As_0)}}{2s_1} \tag{B8}$$

To find the value of the corresponding point d_1: $d_1 = d_0 + s_1 c_1 - c_0$.

Step 2: $l = l + 1$

To calculate the slope of segment l for $l = 2, \dots, L$ we find the slope of the ray originating at point (c_{l-1}, d_{l-1}) that is tangent to $f(b_j x_j)$. The point of tangency c_τ, has two requirements: the ray originating at point (c_{l-1}, d_{l-1}) and the curve $f(b_j x_j)$ meet and at that point in space the derivatives are equal. The slope s_l is calculated as $f'(c_\tau)$. The end point c_l is calculated as in step 1 such that:

$$c_l = \frac{-(As_l + s_{l-1} - 1 - \alpha) \ \pm \ \sqrt{(As_l + s_{l-1} - 1 - \alpha)^2 - 4(s_l)(As_{l-1})}}{2s_l} \tag{B9}$$

The procedure continues until $f^\alpha(b_j x_j)$ has been defined for all points along the $b_j x_j$ axis.

The TLA technique defines $f^\alpha(b_j x_j)$ as L piecewise linear functions, those linear functions can be included in our framework of methods because it consists of the sum of L linear functions:

$$f^\alpha(b_j x_j) = \sum_{l=1}^{L} q_l s_l z_l \tag{B10}$$

where q_l is the length of each segment, s_l the slope and z_l denote L new continuous decision variables. Thus to obtain a linear model, one has to price to pay in terms of the increasing the size of the problem.

The water company's profit maximisation decision problem can now be modelled as an integer linear program:

$$\begin{aligned}
&Max \sum_{j=1}^{N} \sum_{l=1}^{L} q_l s_l z_l - c_j x_j \\
&s.t. \sum_{j=1}^{N} b_j x_j = \sum_{l=1}^{L} q_l \ \ z_l \\
&x_j \in \{0,1\} \qquad j = 1,2, \dots, N \\
&0 \le z_l \le 1 \qquad l = 1,2, \dots, L
\end{aligned} \tag{B11}$$

where the first constraint relates the z_l decision variables back to the original x_j decision variables and the concave shape of the marginal benefit function ensures the segments enter the solution to the integer linear programme in the correct order.

APPENDIX B2

Here we provide example code from the genetic algorithm and integer programming models from the negotiated PES purchasing institution.

```
%% Genetic Algorithm

%% I. Setup the GA parameters
ff          ='fitnessfunc';                    % 2 Absolute Buyers
ff2         ='fitnessfunc2';                   % 1 Absolute 1 Relative
Buyer
maxit       =100;                              % maximum number of
iterations
maxcost     =9999999;                          % maximum allowable cost
popsize     =100;                              % set population size
mutrate     =0.001;                            % set mutation rate
selection   =0.5;                              % fraction of population kept

%% II. Create the initial population

% Generate a random population
if RndCount == 1
   p0 = zeros(bigN - Nsites_L,1) ; p1 = ones(Nsites_L,1); p = vertcat (p0, p1);
   pop = zeros(popsize,bigN);
   for pindex = (1:1:popsize)
      loopsites  = reshape(p(randperm(size(p,1))), 1, bigN);
      pop(pindex,:) = loopsites;
   end
end
% randperm(n) returns a row vector containing a random permutation of the integers from 1 to n
inclusive.

%% III. Main genetic Algorithm Loop

iga=0; % generation counter initialized

while iga<maxit

%% IV. Call the integer programmes which model the buyers' decision problems
         if Relative == 0
                 [sol, value, valuefollower] = feval(ff,pop,bigN, follower_benefits, leader_benefits,
                 cost_final, Nsites_L, BudgetF, BudgetL, sol, value, valuefollower, constraint);
                 % Spatially independent - water quality buyer
        elseif Relative == 1
                 [sol, value, valuefollower] = feval(ff2,pop,bigN, follower_benefits, leader_benefits,
                 cost_final, Nsites_L, BudgetF, BudgetL, sol, value, valuefollower, constraint,
                 pair_dist, pair_dist_len, speciesbysite_mat, Nspecies, z_constraint, x_constraint,
                 cutoff, m, n, alt_LF);
                 % Spatially interdependent – Biodiversity buyer
        end

        iga=iga+1; % increments generation counter

%% V. Stopping criteria
        if iga>maxit || cost(1)>maxcost
                 break
        end
```

```
%% VI. Selection criteria
    % Keep top chromosome
    Top = pop(1,:);
    for ic=1:2:popsize/2
        off1 = zeros(1, bigN); off2 = zeros(1, bigN); off3 = zeros(1, bigN); off4 = zeros(1,
        bigN);
        A = pop(ic,:); B = pop(ic+1,:); % select mates
        [AB1, AB1ind] = find((A==1)&(B==1));
        [ABx, ABxind] = find((A==0)&(B==1)|(A==1)&(B==0));
        RndABxind = randperm(size(ABxind,2), size(ABxind,2)/2);
        % random integers half the length of ABxind
        RndABxind2 = randperm(size(ABxind,2), size(ABxind,2)/2);
        % random integers half the length of ABxind
        % Update offspring
        off1(AB1ind)=1; off1(ABxind(RndABxind))=1;
        off2(AB1ind)=1; off2(ABxind)=1; off2(ABxind(RndABxind))=0;
        off3(AB1ind)=1; off3(ABxind(RndABxind2))=1;
        off4(AB1ind)=1; off4(ABxind)=1; off4(ABxind(RndABxind2))=0;
        % Assign to pop
        pop(ic,:) = off1;
        pop(ic+1,:) = off2;
        pop(popsize + 1 - ic,:) = off3;
        pop(popsize - ic,:) = off4;
    end
    % Insert the top chromosome kept from last iteration
    pop(popsize,:) = Top;

%% VII. Mutate the population
    nmut=ceil(popsize*bigN*mutrate);
    for ic = 1:nmut
        row1=ceil(rand*(popsize-1))+1;
        col1=ceil(rand*bigN);
        col2=ceil(rand*bigN);
        temp=pop(row1,col1);
        pop(row1,col1)=pop(row1,col2);
        pop(row1,col2)=temp;
        im(ic)=row1;
    end

    % Check for infeasible populations
    for fcheckind = 1:popsize
        if sum(pop(fcheckind,:),2) > Nsites_L;      % row sum
            pop(fcheckind,:) = reshape(p(randperm(size(p,1))), 1, bigN);
                            % Replace infeasible populations with random
                            populations
        end
    end

end
```

Below are examples of the integer linear programming used for the buyers' decision problems in the negotiation process called in step IV of the genetic algorithm. They are programmed as CPLEX objects within Matlab.

```
%% Water quality buyer:

    % Initialize the CPLEX object
    cplex              = Cplex('CplexMIP');                    % Cplex object
    cplex.Model.sense  = 'maximize';                           % Maximise or Minimise

    % Use arrays to populate the model
    cplex.Model.obj    = follower_benefits;                    % objective
    cplex.Model.lb     = zeros(bigN,1);                        % lower bound
    cplex.Model.ub     = ones(bigN,1);                         % upper bound
    cplex.Model.ctype  = repmat('B',1,bigN);                   % variable type (binary,
continuous)
    cplex.Model.A      = [cost_final;
                  pop(pop_ind,:)
                  follower_benefits];
                                                               % constraints
    cplex.Model.lhs    = horzcat(BudgetF + BudgetL,Nsites_L, constraint);
                                                               % left hand side constraints
    cplex.Model.rhs    = horzcat(0,Nsites_L, inf);             % Right hand side constraints
    % Optimize the problem
    cplex.solve();                                             % call the solver


%% Biodiversity buyer:

    % Initialize the CPLEX object
    cplex              = Cplex('CplexMIP');                    % Cplex object
    cplex.Model.sense  = 'maximize';                           % Maximise or Minimise

    % Use arrays to populate the model
    cplex.Model.obj    = horzcat(zeros(1,bigN),pair_dist);     % objective
    cplex.Model.lb     = zeros(bigN+pair_dist_len,1);          % lower bound
    cplex.Model.ub     = ones(bigN+pair_dist_len,1);           % upper bound
    cplex.Model.ctype  = horzcat(repmat('B',1,bigN),repmat('C',1,pair_dist_len));
                                                               % variable type (binary,
                                                        continuous)
    cplex.Model.A      =  [horzcat(speciesbysite_mat,zeros(Nspecies,length(pair_dist)));
                       horzcat(x_constraint,-z_constraint)
                       horzcat(cost_final,zeros(1,pair_dist_len))
                       horzcat(pop(pop_ind,:), zeros(1,pair_dist_len))];
                                                               % constraints
    cplex.Model.lhs    =
                       horzcat(ones(1,Nspecies),zeros(1,pair_dist_len),zeros(1,pair_dist_len),Bu
                       dgetF + BudgetL,Nsites_L);
                                                               % left hand side constraints
    cplex.Model.rhs    =
horzcat(inf(1,Nspecies),inf(1,pair_dist_len),inf(1,pair_dist_len),0,Nsites_L);
                                                               % right hand side constraints
    % Optimize the problem
    cplex.solve();                                             % call the solver
```

APPENDIX C

Chapter 3: Experiment walkthrough

Welcome

- Welcome to the experiment! We are about start.

- Before we start, please could you put away anything that you have on your desks and turn off and put away your mobile phones. We will be paying you for your participation in this experiment and in return we expect that you will focus on that task for the next hour to an hour and a half.

- During that time you and the other participants in the room will be undertaking a series of 7 tasks on the computers.

- In those tasks you will be teamed-up with 2 other participants to make a group of THREE people. You won't know who the other people are in your group and the members of your group will change from task to task.

- Each task will involve you negotiating with the other members of your group in an attempt to agree on a DEAL. Whether you reach a deal and what particular deal you agree upon will determine how much money you will be entitled to from that task.

- On your desk, you should have a document outlining the key elements of each task. You can refer to that as we walk you through how a task will be played out on your computer.

Task and Round Counters and Timer

- To do that, we are going to begin by introducing you to the basic elements you will see on the screen in each task. So, to the top left of your screen you should see a task counter, this will update as you work through each of the seven tasks.

- In each task you will go through a series of rounds of negotiation with the other members of your group. The counter to the top right will tell you which round of negotiation you have reached.

- The final element at the top of the screen is a timer. You should now be able to see that counting down.

- During the negotiations you will have to make decisions, but you will only have limited time to come to those decisions … sometimes as little as 10 seconds. As soon as it is your turn to make a decision, the COUNT DOWN begins. If the countdown reaches zero then you will TIME OUT and forfeit your opportunity to make that decision … which may have an impact on how much you get paid. As a result, you will have to think quickly during the experiment.

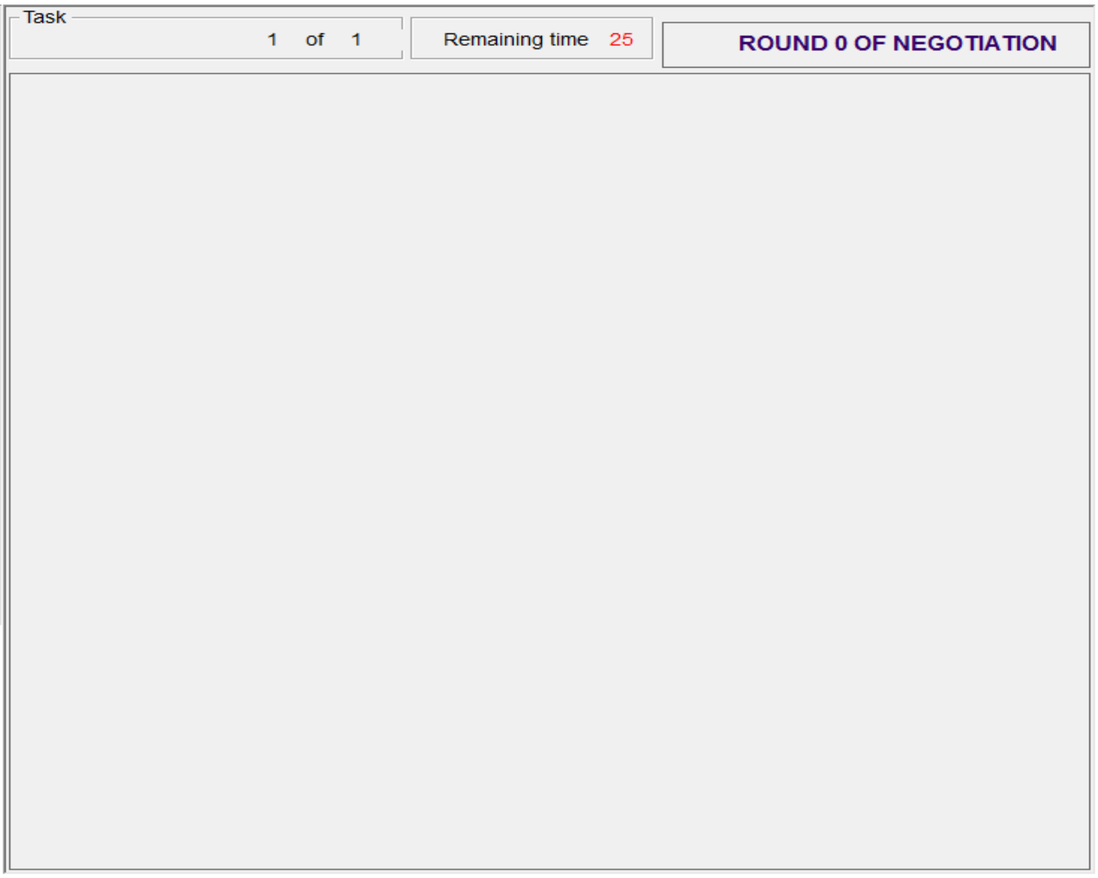

Default Payment

- In each task, each person in a Group of 3 is allocated to take the role of PLAYER 1, PLAYER 2 or PLAYER 3 … which particular role you take on will change from task to task.

- Each player in your group starts out entitled to the same DEFAULT PAYMENT. This is the payment that you will get if your negotiations fail and the three members of your group cannot agree to a DEAL.

- The first box on your screen shows this payment. For all the tasks you will undertake today your DEFAULT PAYMENT will be £7.50.

- Alternatively, provided each of the 3 players in your group agrees to the idea, then instead of each member getting their DEFAULT PAYMENT they will get their DEAL PAYMENT instead.

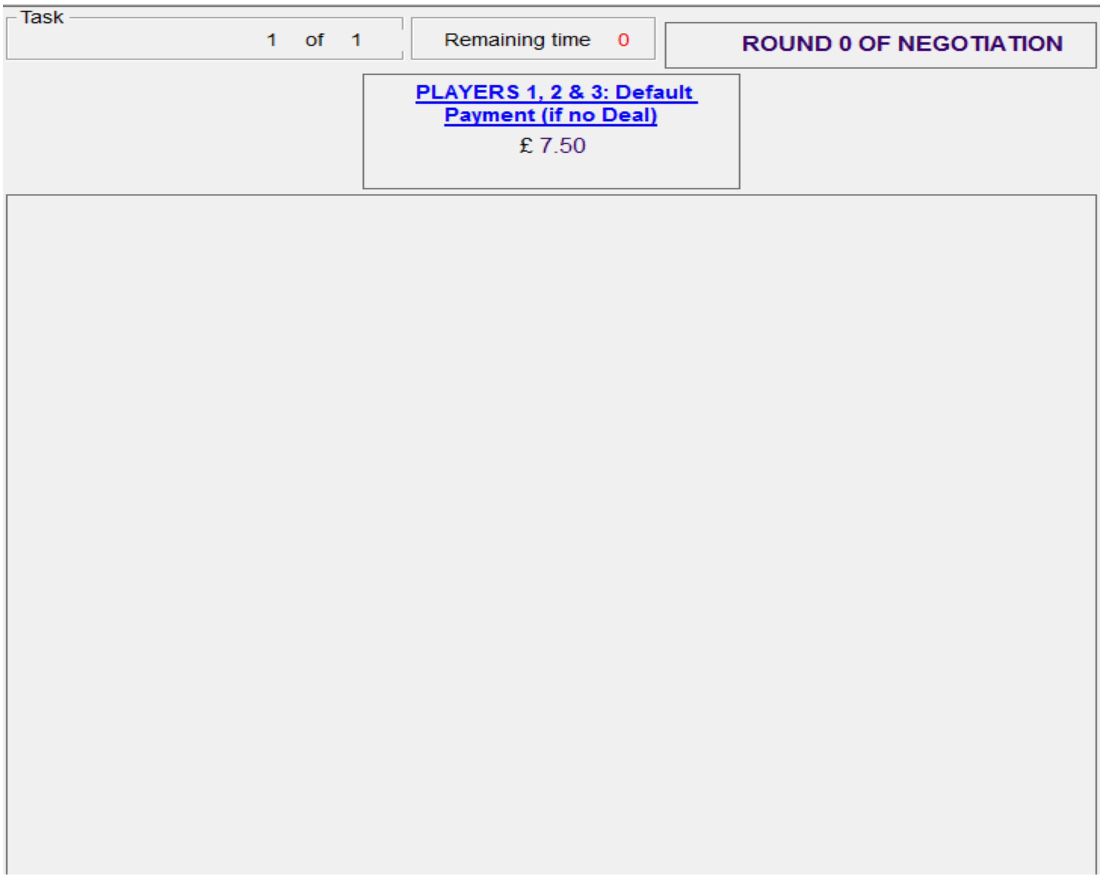

Deal Payment Players 1 & 2

- The Deal payments for Players 1 and 2 are shown in the two boxes that have now appeared on your screen.

- Notice that the role that you will be playing in any particular Task will be indicated to you by highlighting the title of that Player's Deal payment box in Blue. In this example, you are Player 2 … though remember which role you play will change from task to task.

- If you agree to a deal, then in some tasks there may be more than one possible outcome and you won't know which of those outcomes will turn out to be your actual Deal payment until after the negotiations have finished. We label those different Outcomes A and B. In the task we are considering here, however, there is only one outcome … outcome A. Accordingly, as Player 2 in this task, if your group were to agree to a Deal then you can be 100% certain that you will be due a Deal payment of £19. The first few tasks will be just like this. Don't worry, we'll come back and talk you through Deals with more than one Outcome before you start on those tasks.

- Notice that as Player 2, your Deal payment of £19 is considerably higher than your Default payment of £7.50. The same is true for Player 1 who stands to make £14 if a deal is agreed as opposed to the Default payment of £7.50. Indeed, in all the tasks Players 1 and 2 will always have Deal payments that are larger than their Default payments and, therefore, will be keen for all 3 Players to reach an agreement that allows them to claim their Default payments.

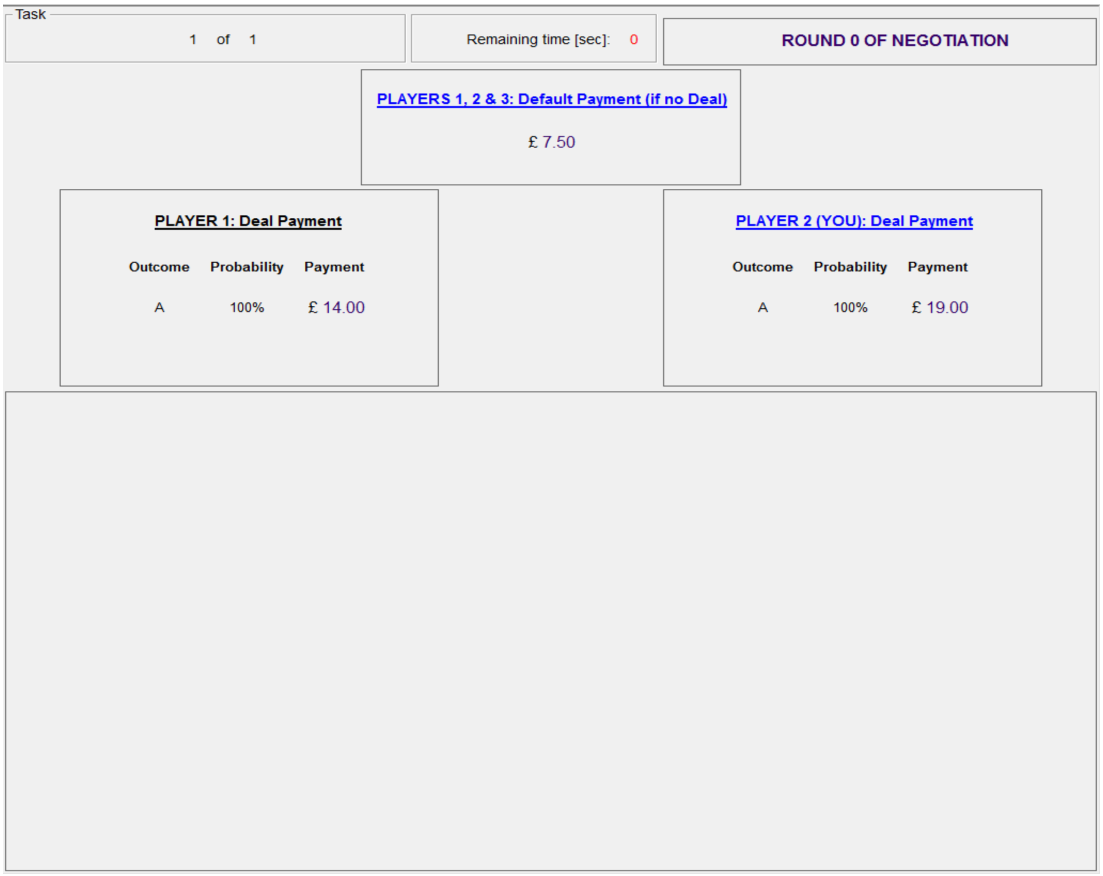

Deal payment Player 3

- The key obstacle in reaching an agreement, however, is that the Deal payment for Player 3 is always zero. You should be able to see the Deal payment box for Player 3 at the bottom of your screen.

- For that reason, the only way agreement can be reached is if the Players can negotiate a DEAL. That Deal involves Players 1 and 2 committing to share enough of their Deal payments with Player 3 so as to convince Player 3 that it is worth their while agreeing to the Deal.

- Negotiations in a task always begin with Player 1 making a PROPOSAL to Player 2. In that Proposal Player 1 suggests how much of their own Deal payment and how much of Player 2's Deal payment should be offered to Player 3. The sum of those two suggested contributions is the proposed payment to be made to Player 3.

- Player 1's Proposal is passed on to Player 2 who must decide whether to REJECT or ACCEPT it. If Player 2 rejects the Proposal, they may get the chance to offer a COUNTER PROPOSAL ... and negotiations may go back and forth between Players 1 and 2 until they finally agree on a Proposal to offer to Player 3.

Reject Proposal

- In this walk through, we join the negotiations part way through. In the left hand side of the Proposal Accept/Reject Box that has just appeared on your screen the Proposal History table which lists the last 5 proposals that have passed between Players 1 and 2.

- In this case, Player 1 started the negotiation by proposing that she pay £2.50 towards a payment to Player 3 while you, as Player 2, should contribute a further £7.50 … giving a total Payment to Player 3 of £10. In this case, you rejected that proposal (which is why it is coloured red in the table) and suggested a counter proposal in which Player 1 paid £5 and you paid £4 (a total Payment to Player 3 of £9). Unfortunately, Player 1 wasn’t happy with that proposal and rejected it, coming back with another proposal in which she pays £4 and you pay £8. Since that is the current proposal on the table, it is coloured black in the Proposal History table.

- You can see that current proposal written out in large in the right hand side of the box in red text and next to it a red decision button with the word “rejected” written on it. This is exactly what the screen will look like when you first receive a proposal to consider. The fact that the decision button says “rejected” and the text is red indicates that you currently intend to reject that proposal.

- By clicking on that decision button you can indicate, instead, that you would like to Accept the Proposal instead. Go ahead and try that now.

- Notice that the text has gone green and the decision button is now grey with the word “accepted” on it. That indicates that you currently intend to accept the proposal.

- To register your decision you MUST press on OK. Whatever word is written on the button when you press OK will be the decision you register in the negotiation. If you fail to press OK then the computer will not register your decision and just assume that you have rejected the Proposal.

- Now click on the decision button again to change it back to rejected. Observe the text that appears next to the OK button. This text is a warning, informing you that should you click OK and thereby reject the proposal then you run the risk that the NEGOTIATIONS WILL FAIL. If that happens then no Deal is reached and each player will have to content themselves with their DEFAULT PAYMENT.

- In this case the Probability of such a failure happening if you decide to reject the proposal is 1 in 30. If you were to go ahead and press OK, the computer would use its random functions and those odds to establish whether you have been unlucky and the negotiations have failed. During a negotiation those probabilities start out low at 1 chance in 500 ... by the time 5 rejections have been made in a negotiation that probability is up to 1 in 100 ... by the tenth rejection 1 in 15 ... and by the fifteenth more than 1 in 2.

- Making this decision even harder is the fact that you have to do it against the clock. We've disabled that for the purposes of this walk through, but when you start the real tasks, as soon as you see this screen the countdown clock will start clicking down. If you haven't pressed OK to register your decision by the time the countdown clock reaches zero, the computer will simply assume that you are rejecting the proposal.

- Let's, assume that you are sufficiently unhappy with this Proposal that you are prepared to take the risk of rejecting it. Make sure the decision button says "rejected" then press OK. Actually we've disabled that button as well, but we will move you on automatically from our master program.

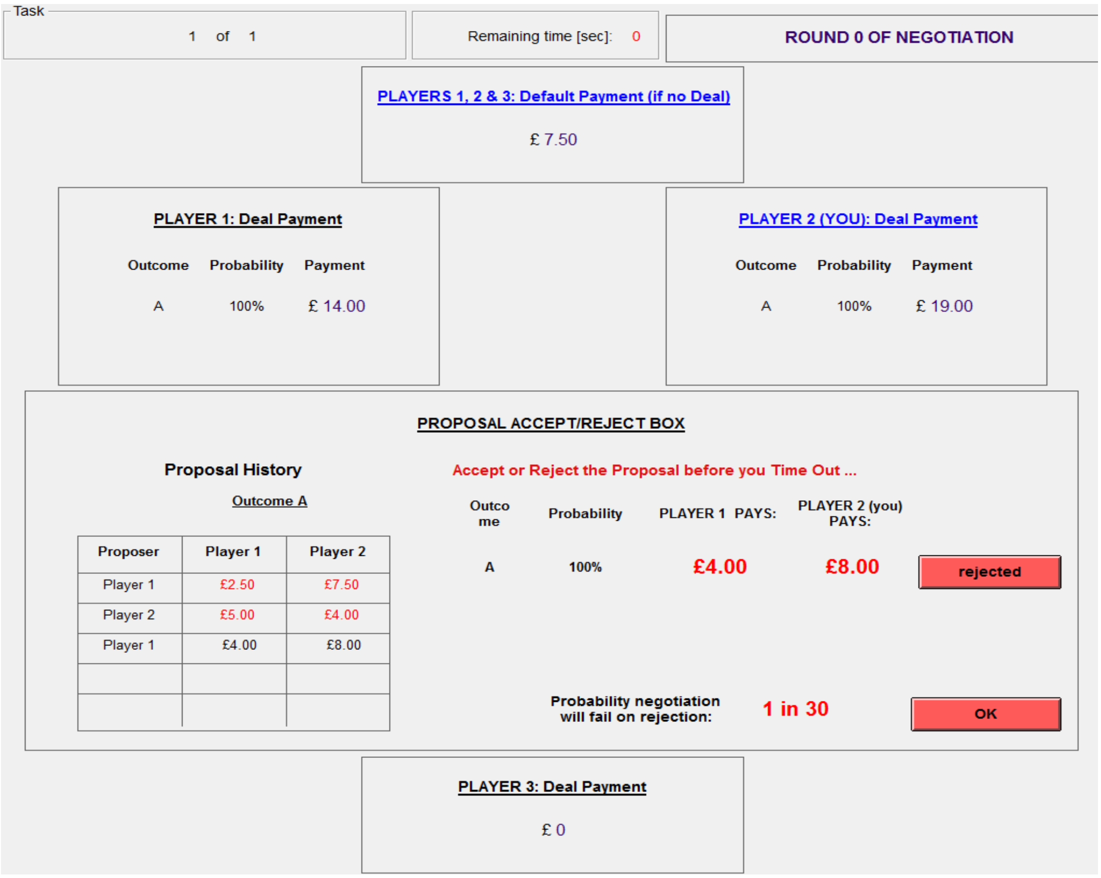

Make a Proposal

- In this case you got lucky and the negotiations did not fail. You now have the opportunity to make a proposal of your own.

- To make your counter proposal you fill in the amounts that you think that you and the other player should make to Player 3 in the boxes provided. In entering those amounts, be aware that the units are in £s … if you want to include a pence amount (which you are perfectly entitled to do) you will have to enter it after a decimal point.

- Since these amounts will be paid for out of Deal payments you will never be able to suggest an amount that exceeds a Player's Deal payment.

- Please fill in the boxes with the following proposal: Player 1 pays £5, you pay £6. Now press enter to register your proposal. Notice that in deciding on a proposal, you will again be up against the clock. If the countdown reaches zero before you have pressed the Enter button, then the computer will simply assume that your proposal is that you both pay £0 to Player 3.

- Go ahead and press Enter to send you proposal off to Player 1 for them to consider … though again notice that for the purposes of this walk through that we have disabled the Enter button and also have got the computer to ensure that you entered the amounts £5 and £6.

Waiting for the Other Player to Consider a Proposal

- Once you have sent your proposal to Player 1, you will move on to a waiting screen. This shouldn't take too long, but please do remain patient … there's nothing more you can do until Player 1 decides whether to accept or reject your proposal.

Accept Proposal

- In this case Player 1 decided to reject your proposal. You all got lucky in that the negotiation did not fail when they pressed reject … and now Player 1 has come back with a new Proposal. In this Proposal Player 1 pays £4.50 and you contribute £6.50 towards a combined payment to Player 3 of £11.

- Notice that the Proposal History table has been updated to show your last proposal … which Player 1 rejected ... and Player 1's new proposal.

- Notice also that the risk of the negotiations failing if you reject has also gone up from 1 in 30 to 1 in 15.

- Let's imagine that you are now happy with the proposal and don't want to take the risk of rejecting.

- Toggle the decision button so it reads "accepted" and then press OK to accept the proposal and then we will move you on to the next screen.

Waiting for Player 3 to Consider an Offer

- The proposal that you have agreed to with Player 1 is now sent over to Player 3 to consider. You will now have to wait to see whether Player 3 is going to accept your Offer.

- If Player 3 does accept your offer, then everyone has consented and a Deal is done. If Player 3 rejects your Offer then you may get the chance to enter into fresh negotiations with Player 1 to see if you can agree to another Offer to put before Player 3.

Player 3 waiting for an Offer

- Again, if you are Player 3, please be patient … it may take Players 1 and 2 a few rounds of proposal and counter-proposal before agreeing on an offer … provided negotiations don't fail before they reach an agreement.

Accept/Reject Offer

- When (and if) an offer arrives, Player 3 will see a screen containing an Offer Accept/Reject Box just like this. To the left is a table listing the offers that have been made to Player 3 in this task. To the right you can see the current offer and buttons allowing Player 3 to accept or reject that offer.

- Notice that Player 3 only sees the total amount that Players 1 and 2 have agreed to pay, not their individual contributions.

- The buttons on this screen work in much the same way as those we looked at previously. When the screen first appears to Player 3, the decision button will say "rejected" indicating an intention to reject the Offer. Again, rejecting an offer comes with a risk of the negotiation failing. That risk is written next to the OK button.

- For the sake of argument, let us assume that you, as Player 3, are happy enough with this offer to think that it is not worth taking the risk of rejecting. Toggle the decision button so that it goes from "rejected" to "accepted", the text of the Offer should go green and, since you are now not planning to reject, the risk information disappears.

- Again in the real tasks you will be making this decision against the clock and you will have to hit the OK button to register your decision before the countdown times out.

- Press OK now … though remember for the walk through we've disabled that button and will move you on from our control program.

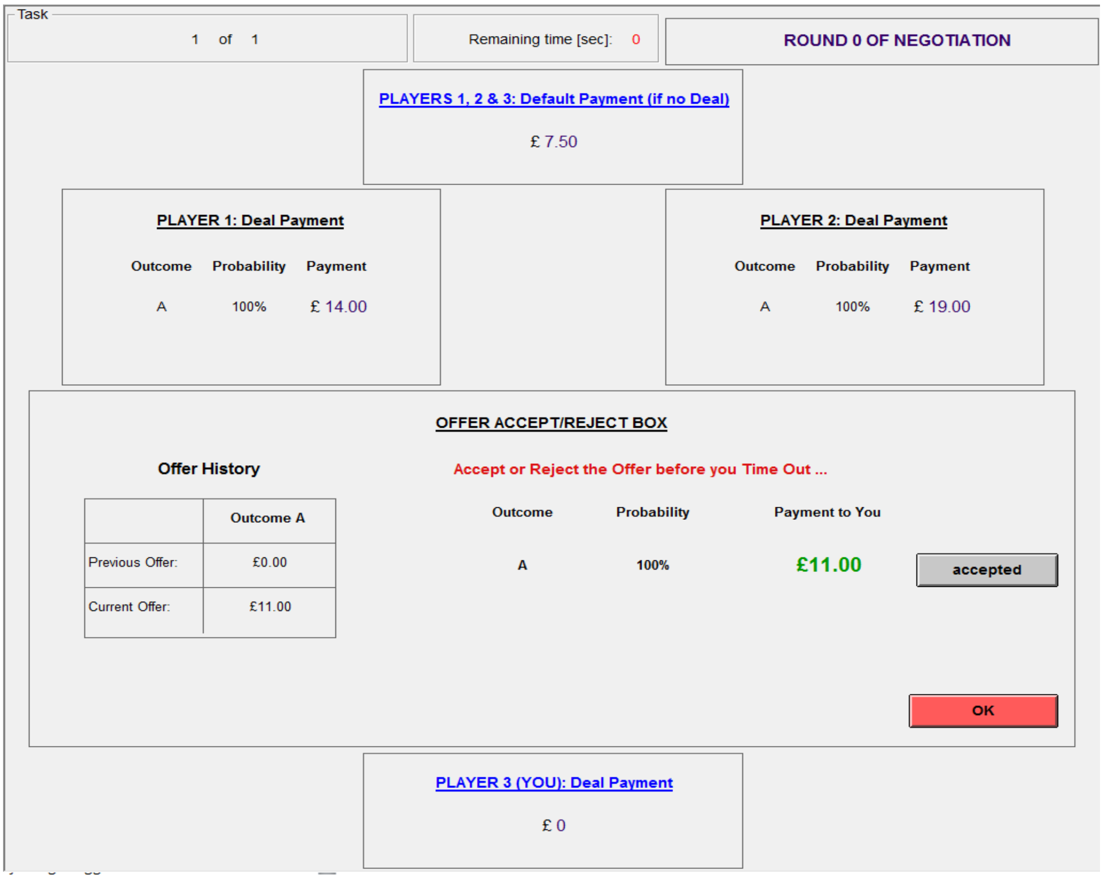

Deal Done Screen

- Since Player 3 has accepted an offer made by Players 1 and 2, this task ends with all three players agreeing to a deal. In this case, Player 3 has foregone the fall-back payment of £7.50 in favour of the £11 offered by Players 1 and 2. If, at the end of the experiment, this task was picked as the one as the one that counts for real for Player 3, then this will be the amount of money they earn from participating in the experiment.

- As a deal was done by this group, each will now see a screen showing what they stand to gain from that task.

- Before they can move on to the next task, they will have to wait until all the other groups finish. Once everyone has completed the task, the next task begins by teaming you up with a different set of two people to form a new group.

Stochastic treatment walkthrough

- We are now moving on to another set of tasks which are slightly more complicated than those you have just done in so much as the Deal payments for Players 1 and 2 can take one of two possible values and which of those is the actual value is not known during the negotiations.

Information Boxes

- Those two possible Deal payments are shown as Outcome A and Outcome B in the information boxes for Players 1 and 2. Notice that Player 3 always has a Deal payment of £0 whatever the outcome.

- While you do not know which outcome will be the actual outcome, you do know that there is exactly half a chance (50% chance) it will be Outcome A and half a chance it will be Outcome B. In this case, if it turns out to be Outcome A, your Deal payment as Player 2 will be £17 while the Deal payment for Player 3 will be £20. If, on the other hand it turns out to be Outcome B, your Deal payment will be £11 and Player 1's will be £10.

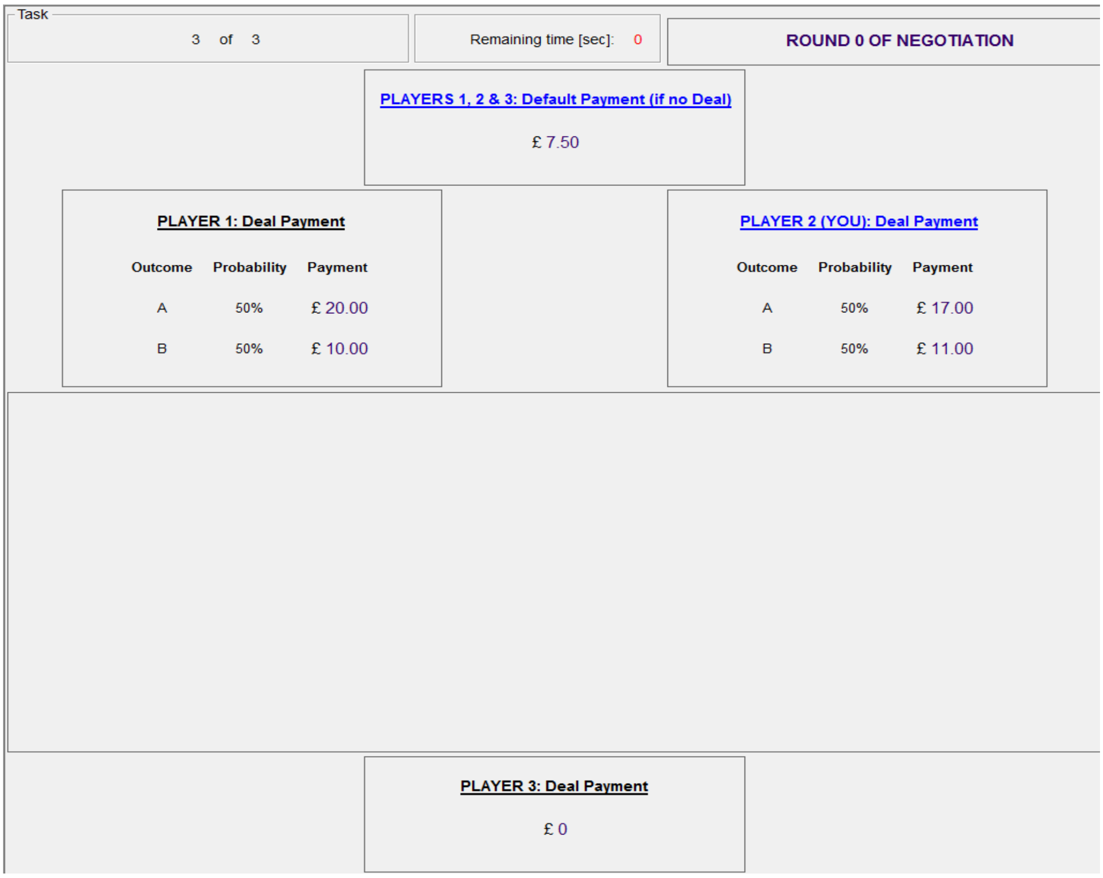

Proposal Screen

- Now in making a Proposal you must consider what payments you think should be made to Player 3 in the event of Outcome A and what payments should be made to Player 3 in the event of Outcome B. You must fill those amounts in the boxes provided and then press Enter to send you Proposal to the other Player to consider. Again you will be making your decisions against the clock.

- In deciding on those amounts bear in mind that Player 3 might be prepared to take some of the risk and accept an offer in which payment in one of the outcomes is below £7.50, provided the payment in the other outcome was sufficiently high that they thought taking that risk was worthwhile.

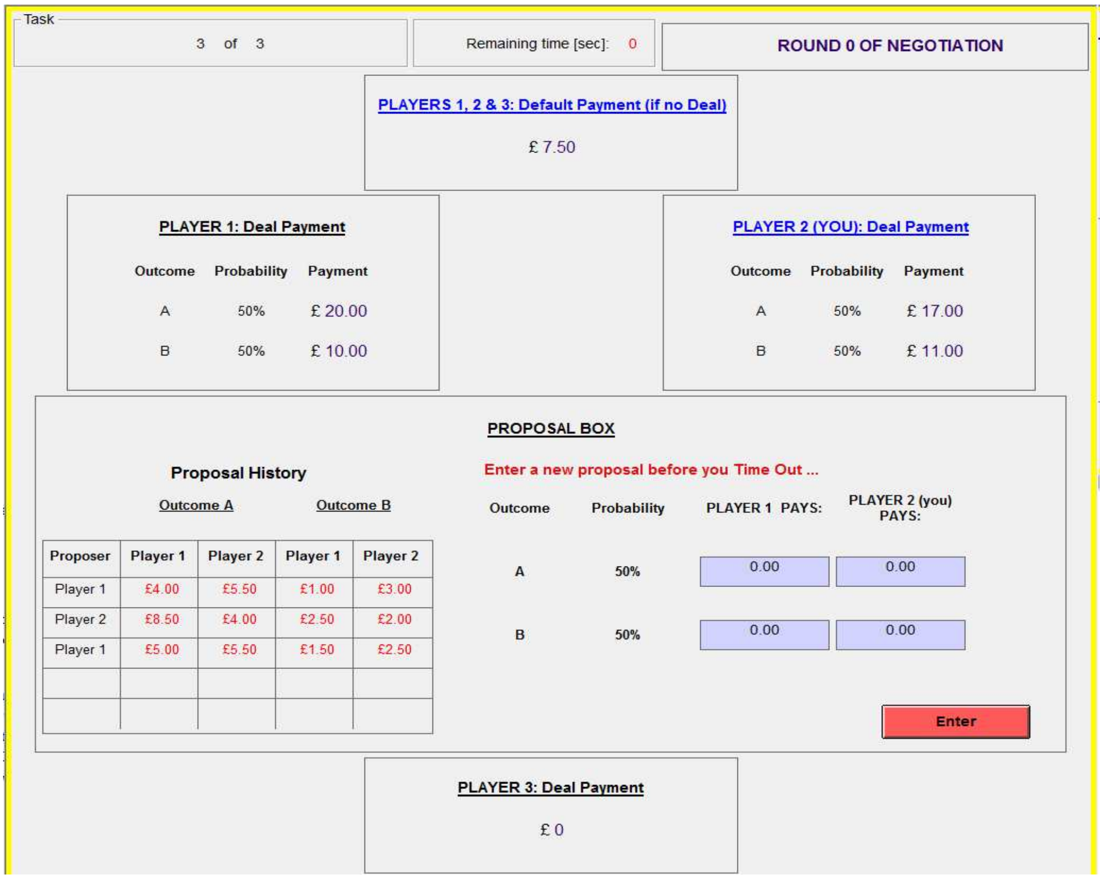

Screen: Proposal Screen

- When you receive a Proposal the screen will now look like this. In making your decision you have a separate decision button for each Outcome. Accordingly, you could accept the proposed payments in Outcome A, but reject those for Outcome B.

- Try toggling the decision buttons for the two Outcomes to “accepted”. You should see the text of the proposal for an Outcome going green, when you toggle the decision button for that outcome to “accepted”.

- Notice that the risk associated with reject only disappears when you have accepted the Proposal for both Outcomes. If you were to do that then, the proposal would be passed on to Player 3 for their consideration.

- Alternatively, you may decide that you are happy with the payments for one Outcome, but not those for the other. In that case, the Proposal would be rejected and your decisions would be recorded in the Proposal History table to the left of the screen.

- For example, in the second to last row of the Proposal History table, you can see how Player 1 reacted to your last Offer. You suggested that in the event of Outcome A that they should pay £7.30 and that you should pay £4.50. They didn’t agree with that and hence those numbers are coloured red in the table. In contrast, you suggested that in the event of Outcome B that you should both pay £2.50. They accepted that part of the proposal and hence those numbers are coloured green in the table. Indeed, they have kept those suggested payments as part of the proposal they have sent back to you to consider.

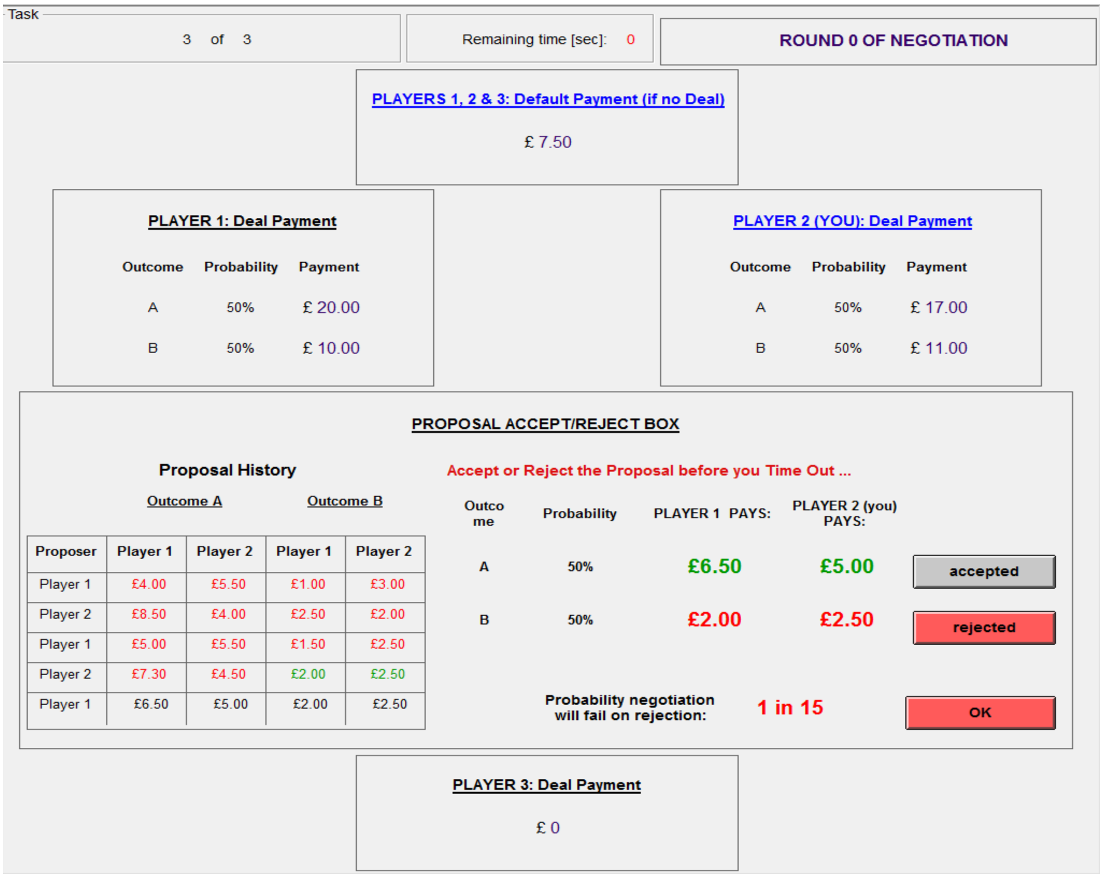

Screen: Offer Screen

- Finally, when an Offer arrives with Player 3 they will see a screen like this, showing the payments that are being offered by Players 1 and 2 in the event of Outcome A and in the event of Outcome B.

- Player 3 can express their opinions on that Offer by toggling the decision buttons from “rejected” to “accepted”. Try that now.

- Of course, Player 3 only avoids the risk associated with making a rejection if the payments proposed for both Outcomes are accepted. If that happens then all 3 Players have agreed and a DEAL is done.

- At the end of the experiment, if this task is the one that is chosen for real, then we need to find out whether Outcome A or Outcome B is the actual outcome. To do that we will simply toss a coin. If the coin comes up Heads then Outcome A is the actual outcome and you will get paid what you agreed to in the Deal under that outcome. Alternatively, if she coin comes up Tails then

Outcome B is the actual outcome and you will get paid what you agreed to in the Deal under that outcome.

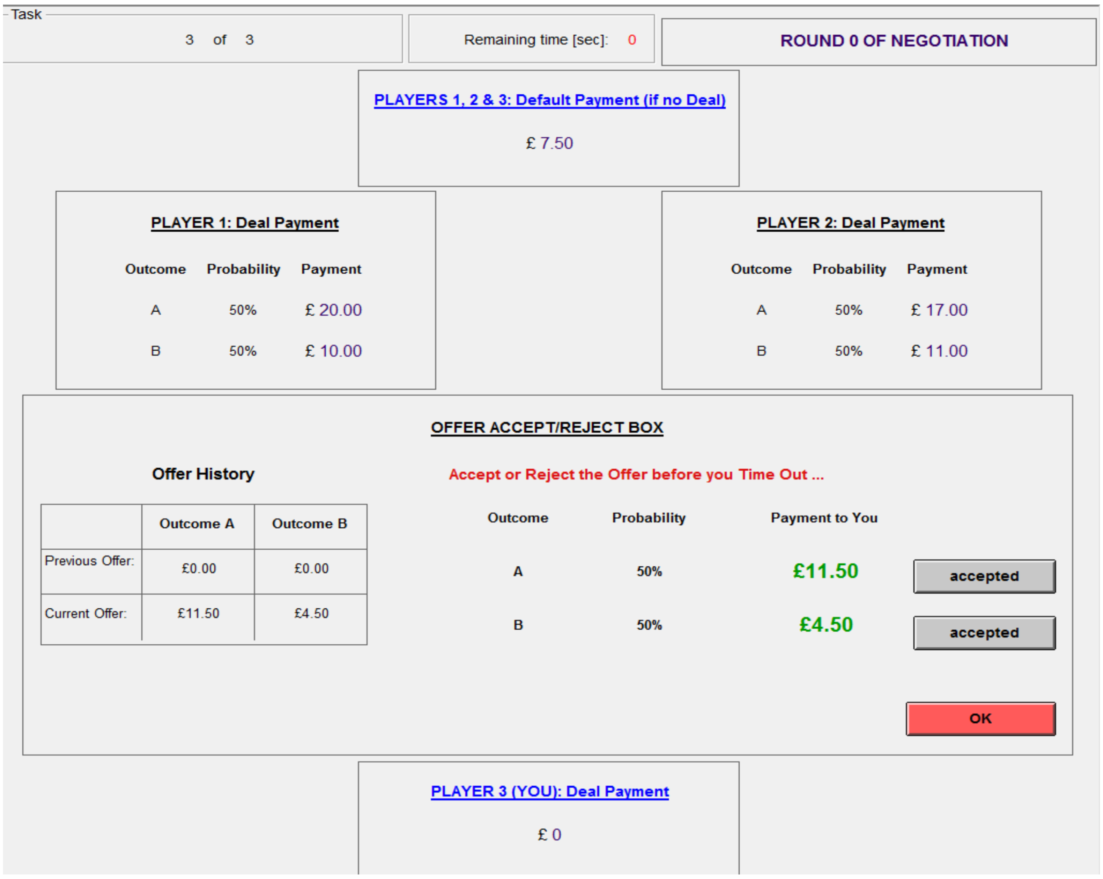

REFERENCES

ABOOLIAN, R., O. BERMAN, AND D. KRASS (2007): "Competitive Facility Location Model with Concave Demand," *European Journal of Operational Research*, 181, 598-619.

— (2009): "Efficient Solution Approaches for a Discrete Multi-Facility Competitive Interaction Model," *Annals of Operations Research*, 167, 297-306.

ACKERBERG, D. A. (2003). Advertising, learning, and consumer choice in experience good markets: an empirical examination*. International Economic Review, 44(3), 1007-1040.

AKTER, S., J. W. BENNETT, AND S. AKHTER (2008): "Preference Uncertainty in Contingent Valuation," *Ecological Economics*, 67, 345-351.

ALBERINI, A., K. BOYLE, AND M. WELSH (2003): "Analysis of Contingent Valuation Data with Multiple Bids and Response Options Allowing Respondents to Express Uncertainty," *Journal of Environmental Economics and Management*, 45, 40-62.

ARIELY, D., G. LOEWENSTEIN, AND D. PRELEC (2003): ""Coherent Arbitrariness": Stable Demand Curves without Stable Preferences," *Quarterly Journal of Economics*, 118, 73-105.

ARMSWORTH, P. R., S. ACS, M. DALLIMER, K.J. GASTON, N. HANLEY, & P. WILSON, (2012). The Cost Of Policy Simplification In Conservation Incentive Programs. *Ecology Letters,* 15(5), 406-414.

AUSUBEL, L. M., P. CRAMTON, AND R. DENECKERE (2002): "Bargaining with Incomplete Information," in *Handbook of Game Theory: Vol. 3*, ed. by R. J. Aumann, and S. Hart. Amsterdam: Elsevier, 1897-1945.

BAJARI, P., R. MCMILLAN, AND S. TADELIS (2009): "Auctions Versus Negotiations in Procurement: An Empirical Analysis," *Journal of Law, Economics, and Organization*, 25, 372-399.

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.

BANERJEE, S., KWASNICA, A. M., & SHORTLE, J. S. (2012). Agglomeration bonus in small and large local networks: A laboratory examination of spatial coordination. *Ecological Economics*, 84, 142-152.

BAJARI, P., AND S. TADELIS (2001): "Incentives Versus Transaction Costs: A Theory of Procurement Contracts," *RAND Journal of Economics*, 32, 387-407.

BARRETT, S. (2001): 'International Cooperation for Sale', *European Economic Review* 45(10), 1835–1850.

Barrett, S. (1994). Self-enforcing international environmental agreements. *Oxford Economic Papers*, 878-894.

BARRATT, S., & R. STAVINS (2003): Increasing participation and compliance in international climate change agreements. *International Environmental Agreements*, 3(4), 349-376.

BATEMAN, I. J., D. BURGESS, W. G. HUTCHINSON, AND D. I. MATTHEWS (2008): "Learning Design Contingent Valuation (Ldcv): Noaa Guidelines, Preference Learning

and Coherent Arbitrariness," *Journal of Environmental Economics and Management*, 55, 127-141.
BATEMAN, I. J., B. H. DAY, D. P. DUPONT, AND S. GEORGIOU (2009): "Procedural Invariance Testing of the One-and-One-Half-Bound Dichotomous Choice Elicitation Method," *The Review of Economics and Statistics*, 91, 806 - 820.
BERGER, L. A., AND J. C. HERSHEY (1994): "Moral Hazard, Risk Seeking, and Free Riding," *Journal of Risk and Uncertainty*, 9, 173-186.
BINMORE, K., P. MORGAN, A. SNAKED, AND J. SUTTON (1991): "Do People Exploit Their Bargaining Power? An Experimental Study," *Games and Economic Behavior*, 3, 295-322.
BOLTON, G., AND A. OCKENFELS (2000): "Erc: A Theory of Equity, Reciprocity, and Competition," *American Economic Review*, 90, 166-193.
BOYLE, K., F. R. JOHNSON, D. W. MCCOLLUM, W. H. DESVOUSGES, R. W. DUNFORD, AND S. P. HUDSON (1996): "Valuing Public Goods: Discrete Versus Continuous Contingent-Valuation Responses," *Land Economics*, 72, 381-396.
BROWN, G., T. PATTERSON, AND N. CAIN (2011): "The Devil in the Details: Non-Convexities in Ecosystem Service Provision," *Resource and Energy Economics*, 33, 355-365.
BROWN, T. C., P. A. CHAMP, R. C. BISHOP, AND D. W. MCCOLLUM (1996): "Which Response Format Reveals the Truth About Donations to a Public Good?," *Land Economics*, 72, 152-166.
BRUCE, C., AND J. CLARK (2010a): "The Effect of Entitlements and Equality on Cooperative Bargaining with Private, Unverifiable Information," Working Papers in Economics 10/68, University of Canterbury, Department of Economics and Finance.
— (2010b): "The Efficiency of Direct Public Involvement in Environmental Policymaking: An Experimental Test," *Environmental and Resource Economics*, 45, 157 - 182.
— (2012): "The Impact of Entitlements and Equity on Cooperative Bargaining: An Experiment," *Economic Inquiry*, 50, 867-879.
CADIGAN, J., P. M. SCHMITT, R. SHUPP, AND K. SWOPE (2009): "An Experimental Study of the Holdout Problem in a Multilateral Bargaining Game," *Southern Economic Journal*, 76, 444-457.
CAMERER, C. (2003): *Behavioral Game Theory, Experiments in Strategic Interaction*. New Jersey, USA: Princeton University Press.
CARSON, R. T., N. E. FLORES, AND N. MEADE (2001): "Contingent Valuation: Controversies and Evidence," *Environmental and Resource Economics*, 19, 173-210.
CASSAR, A. (2007). Coordination and cooperation in local, random and small world networks: Experimental evidence. *Games and Economic Behavior*, 58(2), 209-230.
CHAMP, P. A., R. C. BISHOP, T. C. BROWN, AND D. W. MCCOLLUM (1997): "Using Donation Mechanisms to Value Nonuse Benefits from Public Goods," *Journal of Environmental Economics and Management*, 33, 151-162.
CHAN, K. M., SHAW, M. R., CAMERON, D. R., UNDERWOOD, E. C., & DAILY, G. C. (2006). Conservation planning for ecosystem services. *PLoS Biol*, 4(11), E379.

CHANKONG, V., AND Y. Y. HAIMES (1983): *Multiobjective Decision Making: Theory and Methodology*. New York: North Holland.
COMMENGES, D. (2002): "Inference for Multi-State Models from Interval-Censored Data," *Statistical methods in medical research*, 11, 167 - 182.
CROSON, R. (1996): "Information in Ultimatum Games: An Experimental Study," *Journal of Economic Behavior & Organization*, 30, 197-212.
CROSON, R., T. BOLES, AND J. K. MURNIGHAN (2003): "Cheap Talk in Bargaining Experiments: Lying and Threats in Ultimatum Games," *Journal of Economic Behavior \& Organization*, 51, 143--159.
CZAJKOWSKI, M., HANLEY, N., & LARIVIERE, J. (2015). The effects of experience on preferences: theory and empirics for environmental public goods. American Journal of Agricultural Economics, *97 (1): 333-351*.
DAILY, G. C., AND K. ELLISON (2002): *The New Economy of Nature: The Quest to Make Conservation Profitable*. Washington DC: Island Press.
DAY, B., AND L. COULDRICK (2013): "Payment for Ecosystem Services Pilot Project: The Fowey River Improvement Auction," Defra, London.
DAY, B., J. ELLIOTT, G. JONES, A. BINNER, AND G. SMITH (Forthcoming): "Scoping the Strengths and Weaknesses of Different Auction and Pes Mechanisms for Countryside Stewardship " London, UK: Defra, Reference No: LM0105.
DEFRA (2013): "Developing the Potential for Payments for Ecosystem Services: An Action Plan," London, UK: Department for Environment, Food and Rural Affairs.
— (2014): "The New Common Agricultural Policy Schemes in England: December 2014 Update," London, UK.
DICKINSON, D. L. (1998): "The Voluntary Contributions Mechanism with Uncertain Group Payoffs," *Journal of Economic Behavior & Organization*, 35, 517-533.
DUBOURG, W. R., M. W. JONES-LEE, AND G. LOOMES (1994): "Imprecise Preferences and the Wtp-Wta Disparity," *Journal of Risk and Uncertainty*, 9, 115-133.
— (1997): "Imprecise Preferences and Survey Design in Contingent Valuation," *Econometrica*, 64, 681-702.
ENGEL, S., S. PAGIOLA, AND S. WUNDER (2008): "Designing Payments for Environmental Services in Theory and Practice: An Overview of the Issues," *Ecological Economics*, 65, 663-674.
EUROPEAN COMMISSION (2014):"Agricultural and Rural Development: Agriculture and the Environment,"08/09/14
EVANS, M. F., N. E. FLORES, AND K. BOYLE (2003): "Multiple-Bounded Uncertainty Choice Data as Probabilistic Intentions," *Land Economics*, 79, 549-560.
FAO (2007): "The State of Food and Agriculture 2007, Paying Farmers for Environmental Services," FAO Agriculture Series No. 38, ISSN 0081-4539, Rome.
FERRARO, P. J. (2008): "Asymmetric Information and Contract Design for Payments for Environmental Services," *Ecological Economics*, 65, 810-821.
FISCHBACHER, U. (2007): "Z-Tree: Zurich Toolbox for Ready-Made Economic Experiments," *Experimental Economics*, 10, 171-178.
FLACHAIRE, E., AND G. HOLLARD (2006): "Controlling Starting-Point Bias in Double-Bounded Contingent Valuation Surveys," *Land Economics*, 82, 103-111.

— (2007): "Starting Point Bias and Respondent Uncertainty in Dichotomous Choice Contingent Valuation Surveys," *Resource and Energy Economics*, 29, 183-194.
FRYDMAN, H. (1995): "Semiparametric Estimation in a Three-State Duration-Dependent Markov Model from Interval-Censored Observations with Application to Aids Data," *Biometrics*, 51, 502 - 511.
FRYDMAN, H., AND M. SZAREK (2009): "Nonparametric Estimation in a Markov "Illness-Death" Process from Interval Censored Observations with Missing Intermediate Transition Status," *Biometrics*, 65, 143 - 151.
GOLDBERG, V. P. (1977): "Competitive Bidding and the Production of Precontract Information," *Bell Journal of Economics*, 8, 250-261.
GOLDMAN, R. L., B. H. THOMPSON, AND G. C. DAILY (2007): "Institutional Incentives for Managing the Landscape: Inducing Cooperation for the Production of Ecosystem Services," *Ecological Economics*, 64, 333-343.
GREINER, B. (2004): "The Online Recruitment System Orsee 2.0 - a Guide for the Organization of Experiments in Economics," University of Cologne, Department of Economics.
GUTH, W., S. HUCK, AND P. OCKENFELS (1996): "Two-Level Ultimatum Bargaining with Incomplete Information: An Experimental Study," *Economic Journal*, 106, 593--604.
GÜTH, W., AND E. VAN DAMME (1998): "Information, Strategic Behavior, and Fairness in Ultimatum Bargaining: An Experimental Study," *Journal of Mathematical Psychology*, 42, 227-247.
HANLEY, N., B. KRISTRÖM, AND J. F. SHOGREN (2009): "Coherent Arbitrariness: On Value Uncertainty for Environmental Goods," *Land Economics*, 85, 41-50.
HAUPT, R. L., AND S. E. HAUPT (2004): *Practical Genetic Algorithms*. New Jersey, USA: John Wiley & Sons.
HEIMAN, A., B. MCWILLIAMS, AND D. ZILBERMAN (2001): "Demonstrations and Money-Back Guarantees: Market Mechanisms to Reduce Uncertainty," *Journal of Business Research*, 54, 71-84.
HERRERO, M. (1985): "A Strategic Bargaining Approach to Market Institutions," Unpublished Ph.D. thesis, University of London (LSE).
HERRIGES, J. A., AND J. F. SHOGREN (1996): "Starting Point Bias in Dichotomous Choice Valuation with Follow-up Questioning," *Journal of Environmental Economics and Management*, 30, 112-131.
HOFFMAN, E., AND M. L. SPITZER (1985): "Entitlements, Rights, and Fairness: An Experimental Examination of Subjects' Concepts of Distributive Justice," *Journal of Legal Studies*, 14, 259-297.
HOROWITZ, J., AND K. MCCONNELL (2002): "A Review of Wta/Wtp Studies," *Journal of Environmental Economics and Management*, 44, 426-447.
HUTCHINSON, W. G., K. S. CARSON, AND S. M. CHILTON (2007): "Does Consequentiality Ensure Demand Revelation in Double Referenda," Kyoto, Japan.
JIN, Y., D. ZILBERMAN, AND A. HEIMAN (2005): "The Welfare Economics of a Money-Back Guarantee in Retailing," Department of Agricultural and Resource Economics, University of California at Berkeley.

KAGEL, J. H., AND K. W. WOLFE (2001): "Tests of Fairness Models Based on Equity Considerations in a Three-Person Ultimatum Game," *Experimental Economics*, 4, 203-219.

KAHNEMAN, D., AND J. L. KNETSCH (1992): "Valuing Public Goods: The Purchase of Moral Satisfaction," *Journal of Environmental Economics and Management*, 22, 57-70.

KAHNEMAN, D., AND A. TVERSKY (1982): "The Psychology of Preferences," *Scientific American*, 246, 160-173.

KEMKES, R. J., J. FARLEY, AND C. J. KOLIBA (2010): "Determining When Payments Are an Effective Policy Approach to Ecosystem Service Provision," *Ecological Economics*, 69, 2069-2074.

KENNAN, J., AND R. WILSON (1993): "Bargaining with Private Information," *Journal of Economic Literature*, 31, 45-104.

KESER, C., AND C. MONTMARQUETTE (2008): "Voluntary Contributions to Reduce Expected Public Losses," *Journal of Economic Behavior & Organization*, 66, 477-491.

KINZIG, A. P., C. PERRINGS, F. S. CHAPIN, S. POLASKY, V. K. SMITH, D. TILMAN, AND B. L. TURNER (2011): "Sustainability. Paying for Ecosystem Services--Promise and Peril," *Science*, 334, 603-604.

KIRKPATRICK, J. B. (1983): "An Iterative Method for Establishing Priorities for the Selection of Nature Reserves: An Example from Tasmania," *Biological Conservation*, 25, 127-134.

KLEIN, J., AND M. MOESCHBERGER (1997): *Survival Analysis: Techniques for Censored and Truncated Data*. Springer New York.

KLINE, J. D., AND M. J. MAZZOTTA (2012): "Evaluating Tradeoffs among Ecosystem Services in the Management of Public Lands," United States Department of Agriculture, Forest Service, Pacific Northwest Research Station, General Technical Report PNW-GTR-865.

KOBAYASHI, M., K. MOELTNER, AND K. ROLLINS (2012): "Latent Thresholds Analysis of Choice Data under Value Uncertainty," *American Journal of Agricultural Economics*, 94, 189-208.

KOBAYASHI, M., K. ROLLINS, AND M. D. R. EVANS (2010): "Sensitivity of Wtp Estimates to Definition of a "Yes": Reinterpreting Expressed Response Intensity," *Agricultural and Resource Economics Review*, 39.

LATACZ-LOHMANN, U., AND S. SCHILIZZI (2005): "Auctions for Conservation Contracts: A Review of the Theoretical and Empirical Literature," Report to the Scottish Executive Environment and Rural Affairs Department

LAURANCE, W. F., T. E. LOVEJOY, H. L. VASCONCELOS, E. M. BRUNA, R. K. DIDHAM, P. C. STOUFFER, C. GASCON, R. O. BIERREGAARD, S. G. LAURANCE, AND E. SAMPAIO (2002): "Ecosystem Decay of Amazonian Forest Fragments: A 22-Year Investigation," *Conservation Biology*, 16, 605-618.

LAWTON, J. H., P. N. M. BROTHERTON, V. K. BROWN, C. ELPHICK, A. H. FITTER, J. FORSHAW, R. W. HADDOW, S. HILBORNE, R. N. LEAFE, G. M. MACE, M. P. SOUTHGATE, W. J. SUTHERLAND, T. E. TEW, J. VARLEY, AND G. R. WYNNE (2010): "Making Space for Nature: A Review of England’s Wildlife Sites and Ecological Network," Report to Defra.

LESTER, S. E., COSTELLO, C., HALPERN, B. S., GAINES, S. D., WHITE, C., & BARTH, J. A. (2013). Evaluating tradeoffs among ecosystem services to inform marine spatial planning. *Marine Policy*, 38, 80-89.
LI, C.-Z., AND L. MATTSSON (1995): "Discrete Choice under Preference Uncertainty: An Improved Structural Model for Contingent Valuation," *Journal of Environmental Economics and Management*, 28, 256-269.
MAES, J., PARACCHINI, M. L., ZULIAN, G., DUNBAR, M. B., & ALKEMADE, R. (2012). Synergies and trade-offs between ecosystem service supply, biodiversity, and habitat conservation status in Europe. *Biological conservation*, 155, 1-12.
MAHIEU, P.-A., P. RIERA, B. KRISTRÖM, R. BRÄNNLUND, AND M. GIERGICZNY (2014): "Exploring the Determinants of Uncertainty in Contingent Valuation Surveys," *Journal of Environmental Economics and Policy*, 3, 186-200.
MCFADDEN, D. (1994): "Contingent Valuation and Social Choice," *American Journal of Agricultural Economics*, 76, 689-708.
MILLENNIUM ECOSYSTEM ASSESSMENT (2005): "Ecosystems and Human Well-Being: Synthesis," Island Press, Washington.
MUTHOO, A. (2002): "The Economics of Bargaining," in *Knowledge for Sustainable Development: An Insight into the Encyclopedia of Life Support Systems*: EOLSS Publishers/UNESCO.
NALLE, D. J., J. L. ARTHUR, C. MONTGOMERY, AND J. SESSIONS (2002): "Economic and Spatial Impacts of an Existing Reserve Network on Future Augmentation," *Environmental Modeling & Assessment*, 7, 99-105.
NALLE, D. J., C. MONTGOMERY, J. L. ARTHUR, S. POLASKY, AND N. H. SCHUMAKER (2004): "Modeling Joint Production of Wildlife and Timber," *Journal of Environmental Economics and Management*, 48, 997-1017.
NATURAL CAPITAL COMMITTEE (2015): "The State of Natural Capital: Protecting and Improving Natural Capital for Prosperity and Wellbeing, Third Report to the Economic Affairs Committee," Defra, London.
NATURAL ENGLAND (2014):"Land Management Update Issue 15 June 2014,"
NELSON, E., S. POLASKY, D. J. LEWIS, A. J. PLANTINGA, E. LONSDORF, D. WHITE, D. BAEL, AND J. J. LAWLER (2008): "Efficiency of Incentives to Jointly Increase Carbon Sequestration and Species Conservation on a Landscape.," *Proceedings of the National Academy of Sciences of the United States of America*, 105, 9471-9476.
ONAL, H., AND R. A. BRIERS (2002): "Incorporating Spatial Criteria in Optimum Reserve Network Selection," *Proceedings of the Royal Society B: Biological Sciences*, 269, 2437-2441.
ÖNAL, H., AND R. A. BRIERS (2005): "Designing a Conservation Reserve Network with Minimal Fragmentation: A Linear Integer Programming Approach," *Environmental Modeling & Assessment*, 10, 193-202.
OSBORNE, M., AND A. RUBINSTEIN (1994): *A Course in Game Theory*. London, UK: The MIT Press.
PAGIOLA, S. (2008): "Payments for Environmental Services in Costa Rica," *Ecological Economics*, 65, 712-724.
PARKHURST, G. M., J. F. SHOGREN, C. BASTIAN, P. KIVI, J. DONNER, AND R. B. W. SMITH (2002): "Agglomeration Bonus: An Incentive Mechanism to Reunite Fragmented Habitat for Biodiversity Conservation," *Ecological Economics*, 41, 305-328.

PERROT-MAÎTRE, D. (2006): "The Vittel Payments for Ecosystem Services: A "Perfect" Pes Case?," London, UK.: International Institute for Environment and Development.

PILLUTLA, M. M., AND J. K. MURNIGHAN (1996): "Unfairness, Anger, and Spite: Emotional Rejections of Ultimatum Offers," *Organizational Behavior and Human Decision Processes*, 68, 208-224.

POE, G. L., AND C. A. VOSSLER (2009): "Consequentiality and Contingent Values: An Emerging Paradigm," University Library of Munich, Germany.

POLASKY, S., LEWIS, D. J., PLANTINGA, A. J., & NELSON, E. (2014). Implementing the optimal provision of ecosystem services. Proceedings of the National Academy of Sciences, 111(17), 6248-6253.

POLASKY, S., E. NELSON, J. D. CAMM, B. CSUTI, P. FACKLER, E. LONSDORF, C. MONTGOMERY, D. WHITE, J. ARTHUR, B. GARBER-YONTS, R. HAIGHT, J. KAGAN, A. STARFIELD, AND C. TOBALSKE (2008): "Where to Put Things? Spatial Land Management to Sustain Biodiversity and Economic Returns," *Biological Conservation*, 141, 1505-1524.

POLASKY, S., E. NELSON, E. LONSDORF, P. FACKLER, AND A. STARFIELD (2005): "Conserving Species in a Working Landscape: Land Use with Biological and Economic Objectives," *Ecological Applications*, 15, 1387-1401.

PRESSEY, R. L., H. P. POSSINGHAM, AND J. R. DAY (1997): "Effectiveness of Alternative Heuristic Algorithms for Identifying Indicative Minimum Requirements for Conservation Reserves," *Biological Conservation*, 80, 207-219.

QUICK, T., M. REED, M. SMYTH, R. BIRNIE, C. BAIN, P. ROWCROFT, AND A. WHITE (2013): "Developing Place-Based Approaches for Payments for Ecosystem Services," URS London.

RAPOPORT, A., J. A. SUNDALI, AND D. A. SEALE (1996): "Ultimatums in Two-Person Bargaining with One-Sided Uncertainty: Demand Games," *Journal of Economic Behavior \& Organization*, 30, 173--196.

READY, R. C., J. C. BUZBY, AND D. HU (1996): "Differences between Continuous and Discrete Contingent Value Estimates," *Land Economics*, 72, 397-411.

READY, R. C., S. NAVRUD, AND W. R. DUBOURG (2001): "How Do Respondents with Uncertain Willingness to Pay Answer Contingent Valuation Questions?," *Land Economics*, 77, 315-326.

READY, R. C., J. C. WHITEHEAD, AND G. C. BLOMQUIST (1995): "Contingent Valuation When Respondents Are Ambivalent," *Journal of Environmental Economics and Management*, 29, 181-196.

REESON, A. F., RODRIGUEZ, L. C., WHITTEN, S. M., WILLIAMS, K., NOLLES, K., WINDLE, J., & ROLFE, J. (2011). Adapting auctions for the provision of ecosystem services at the landscape scale. *Ecological Economics*, 70(9), 1621-1627.

REVELLE, C. S., J. C. WILLIAMS, AND J. J. BOLAND (2002): "Counterpart Models in Facility Location Science and Reserve Selection Science," *Environmental Modeling & Assessment*, 7, 71-80.

ROACH, B., K. BOYLE, AND M. WELSH (2002): "Testing Bid Design Effects in Multiple-Bounded, Contingent-Valuation Questions," *Land Economics*, 78, 121-131.

ROSELIUS, T. (1971): "Consumer Rankings of Risk Reduction Methods," *Journal of Marketing*, 35, 56-61.

ROTH, A. E. (1995): "Bargaining Experiments," in *The Handbook of Experimental Economics*, ed. by J. H. Kagel, and A. E. Roth. New Jersey: Princeton University Press, 253-348.

ROTH, A. E., AND M. W. K. MALOUF (1979): "Game-Theoretic Models and the Role of Information in Bargaining," *Psychological Review*, 86, 574–594.

ROTH, A. E., M. W. K. MALOUF, AND J. K. MURNIGHAN (1981): "Sociological Versus Strategic Factors in Bargaining," *Journal of Economic Behavior & Organization*, 2, 153-177.

ROTH, A. E., AND J. K. MURNIGHAN (1982): "The Role of Information in Bargaining: An Experimental Study," *Econometrica*, 50, 1123-1142.

RUBINSTEIN, A. (1982): "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 50, 97-109.

SALZMAN, J. (2005): "The Promise and Perils of Payments for Ecosystem Services," *Int. J. Innovation and Sustainable Development*, 1, 5-20.

— (2009): "A Policy Maker's Guide to Designing Payments for Ecosystem Services," Duke Law Faculty Scholarship, Paper 2081.

— (2010): "Designing Payments for Ecosystem Services."

SÁNCHEZ-AZOFEIFA, G. A., A. PFAFF, J. ROBALINO, AND J. P. BOOMHOWER (2007): "Costa Rica's Payment for Environmental Services Program: Intention, Implementation, and Impact," *Conservation Biology*, 21, 1165-1173.

SARKAR, S., R. L. PRESSEY, D. P. FAITH, C. R. MARGULES, T. FULLER, D. M. STOMS, A. MOFFETT, K. A. WILSON, K. J. WILLIAMS, P. H. WILLIAMS, AND S. ANDELMAN (2006): "Biodiversity Conservation Planning Tools: Present Status and Challenges for the Future," *Annual Review of Environment and Resources*, 31, 123-159.

SCHMITT, P. M. (2004): "On Perceptions of Fairness: The Role of Valuations, Outside Options, and Information in Ultimatum Bargaining Games," *Experimental Economics*, 7, 49--73.

SCHMITT, P. M., R. SHUPP, K. SWOPE, AND J. MAYER (2008): "Pre-Commitment and Personality: Behavioral Explanations in Ultimatum Games," *Journal of Economic Behavior & Organization*, 66, 597-605.

SCHOMERS, S., AND B. MATZDORF (2013): "Payments for Ecosystem Services: A Review and Comparison of Developing and Industrialized Countries," *Ecosystem Services*, 6, 16-30.

SHUPP, R., J. CADIGAN, P. M. SCHMITT, AND K. J. SWOPE (2013): "Institutions and Information in Multilateral Bargaining Experiments," *The B.E. Journal of Economic Analysis & Policy*, 13, 485-524.

SINHA, A., P. MALO, A. FRANTSEV, AND K. DEB (2014): "Finding Optimal Strategies in a Multi-Period Multi-Leader–Follower Stackelberg Game Using an Evolutionary Algorithm," *Computers & Operations Research*, 41, 374-385.

SMITH, S., P. ROWCROFT, M. EVERARD, L. COULDRICK, M. REED, H. ROGERS, T. QUICK, C. EVES, AND C. WHITE (2013): "Payments for Ecosystem Services: A Best Practice Guide," Defra, London.

SNYDER, S. A., J. R. MILLER, A. M. SKIBBE, AND R. G. HAIGHT (2007): "Habitat Acquisition Strategies for Grassland Birds in an Urbanizing Landscape," *Environmental management*, 40, 981-992.

STAHL, I. (1972): "Bargaining Theory," Stockholm, Sweden.

STONEHAM, G., V. CHAUDHRI, A. HA, AND L. STRAPPAZZON (2003): "Auctions for Conservation Contracts: An Empirical Examination of Victoria's Bushtender Trial," *Australian Journal of Agricultural and Resource Economics*, 47, 477–500.

SUTTON, J. (1986): "Non-Cooperative Bargaining Theory: An Introduction," *Review of Economic Studies*, 53, 709-724.

TEEB (2012): "The Economics of Ecosystems and Biodiversity in Business and Enterprise," Edited by Joshua Bishop, Earthscan, London and New York.

TÓTH, S. F., R. G. HAIGHT, S. A. SNYDER, S. GEORGE, J. R. MILLER, M. S. GREGORY, AND A. M. SKIBBE (2009): "Reserve Selection with Minimum Contiguous Area Restrictions: An Application to Open Space Protection Planning in Suburban Chicago," *Biological Conservation*, 142, 1617-1627.

TÓTH, S. F., AND M. E. MCDILL (2008): "Promoting Large, Compact Mature Forest Patches in Harvest Scheduling Models," *Environmental Modeling & Assessment*, 13, 1-15.

TVERSKY, A., AND D. KAHNEMAN (1974): "Judgment under Uncertainty: Heuristics and Biases," *Science*, 185, 1124-1131.

UK NATIONAL ECOSYSTEM ASSESSMENT (2011): "The Uk National Ecosystem Assessment: Synthesis of the Key Findings," UNEP-WCMC, Cambridge.

USDA (2014):"Conservation Programs,"United States Department of Agriculture Economic Research Service,http://www.ers.usda.gov/topics/natural-resources-environment/conservation-programs/background.aspx, 12/01/2015

VALLÉE, T., AND T. BAŞAR (1999): "Off-Line Computation of Stackelberg Solutions with the Genetic Algorithm," *Computational Economics*, 13, 201-209.

VAN KOOTEN, G. C., E. KRCMAR, AND E. H. BULTE (2001): "Preference Uncertainty in Non-Market Valuation: A Fuzzy Approach," *American Journal of Agricultural Economics*, 83, 487-500.

VANEXEL, N., W. BROUWER, B. VANDENBERG, AND M. KOOPMANSCHAP (2006): "With a Little Help from an Anchor: Discussion and Evidence of Anchoring Effects in Contingent Valuation," *Journal of Socio-Economics*, 35, 836-853.

VENTER, O., W. F. LAURANCE, T. IWAMURA, K. A. WILSON, R. A. FULLER, AND H. P. POSSINGHAM (2009): "Harnessing Carbon Payments to Protect Biodiversity," *Science*, 326, 1368.

VOSSLER, C. A., R. G. ETHIER, G. L. POE, AND M. WELSH (2003): "Payment Certainty in Discrete Choice Contingent Valuation Responses: Results from a Field Validity Test," *Southern Economic Journal*, 69, 886-902.

VOSSLER, C. A., G. L. POE, M. WELSH, AND R. G. ETHIER (2004): "Bid Design Effects in Multiple Bounded Discrete Choice Contingent Valuation," *Environmental and Resource Economics*, 29, 401-418.

WANG, H. (1997): "Treatment of "Don't Know" Responses in Contingent Valuation Surveys: A Random Valuation Model," *Journal of Environmental Economics and Management*, 32, 219-232.

WÄTZOLD, F., & DRECHSLER, M. (2014). Agglomeration payment, agglomeration bonus or homogeneous payment?. *Resource and Energy Economics*, 37, 85-101.

WELSH, M., AND R. C. BISHOP (1993): "Multiple Bounded Discrete Choice Models," in *W-133, Benefits & Costs Transfer in Natural Resource Planning*: Western Regional Research Publication, Sixth Interim Report, 331-352.

WELSH, M., AND G. L. POE (1998): "Elicitation Effects in Contingent Valuation: Comparisons to a Multiple Bounded Discrete Choice Approach," *Journal of Environmental Economics and Management*, 36, 170-185.

WHITEHEAD, J. C. (2002): "Incentive Incompatibility and Starting-Point Bias in Iterative Valuation Questions," *Land Economics*, 78, 285-297.

WILLIAMS, J. C., C. S. REVELLE, AND S. A. LEVIN (2005): "Spatial Attributes and Reserve Design Models: A Review," *Environmental Modeling & Assessment*, 10, 163-181.

WILLIAMS, J. C., AND S. A. SNYDER (2005): "Restoring Habitat Corridors in Fragmented Landscapes Using Optimization and Percolation Models," *Environmental Modeling & Assessment*, 10, 239-250.

WOODWARD, R. T. (2011): "Double-Dipping in Environmental Markets," *Journal of Environmental Economics and Management*, 61, 153-169.

WUNDER, S. (2008): "Payments for Environmental Services and the Poor: Concepts and Preliminary Evidence," *Environment and Development Economics*, 13, 279-297.

ZWICK, R., AND X.-P. CHEN (1999): "What Price Fairness? A Bargaining Study," *Management Science*, 45, 804-823.

ZWICK, R., AND V. MAK (2012): "Gaming with Fairness: Some Conjectures on Behavior in Alternating Offer Bargaining Experiments," in *The Oxford Handbook of Economic Conflict Resolution*, ed. by G. Bolton, and R. Croson: Oxford University Press.