Ecological Economics 70 (2011) 1621–1627

Contents lists available at ScienceDirect

Ecological Economics

journal homepage: www.elsevier.com/locate/ecolecon

Analysis

Adapting auctions for the provision of ecosystem services at the landscape scale

Andrew F. Reeson [a,b,*], Luis C. Rodriguez [a], Stuart M. Whitten [a], Kristen Williams [a], Karel Nolles [b,c], Jill Windle [d], John Rolfe [d]

[a] CSIRO Ecosystem Sciences, Canberra, Australia
[b] Faculty of Agriculture Food and Natural Resources, University of Sydney, Australia
[c] Aton Experimental Economics, Sydney, Australia
[d] Centre for Environmental Management, CQUniversity, Rockhampton, Australia

ARTICLE INFO

Article history:
Received 26 February 2010
Received in revised form 24 March 2011
Accepted 28 March 2011
Available online 5 May 2011

Keywords:
Biodiversity conservation
Auction
Competitive tender
Experimental economics
Ecosystem services
Market-based instrument

ABSTRACT

Auctions, or competitive tenders, can overcome information asymmetries to efficiently allocate limited funding for ecosystem services. Most auctions focus on ecosystem services on individual properties to maximise the total amount provided. However, for many services it is not just the total quantity but their location in the landscape relative to other sites that matters. For example, biodiversity conservation may be much more effective if conserved sites are connected. Adapting auctions to address ecosystem services at the landscape scale requires an auction mechanism which can promote coordination while maintaining competition. Multi-round auctions, in which bidding is spread over a number of rounds with information provided between rounds on the location of other bids in the landscape, offer an approach to cost effectively deliver landscape-scale ecosystem services. Experimental economic testing shows these auctions deliver the most cost effective environmental outcomes when the number of rounds is unknown in advance, which minimises rent-seeking behaviour. It also shows that a form of bid-improvement rule facilitates coordination and reduces rent seeking. Where the biophysical science is well developed, such auctions should be relatively straightforward to implement and participate in, and have the potential to provide significantly better outcomes than standard 'one-shot' tenders.

1. Introduction

Payments for ecosystem services (ES) are increasingly being applied to promote biodiversity conservation and other environmental policy goals. Auctions, or competitive tenders, are a proven method of overcoming information asymmetries concerning landholders' private costs and ensuring the efficient allocation of limited ES payments (Latacz-Lohmann and Van der Hamsvoort, 1997; Stoneham et al., 2003). In an ES auction (which is a form of procurement auction), landholders submit bids to provide ES in return for a payment. Landholders are free to choose the level of their payment. However the auction mechanism is competitive, with only those that offer the best value for money (quantity of ES provided per dollar requested) likely to be successful. Most ES auctions adopt a sealed bid, discriminatory price mechanism, in which successful landholders are paid their bid price (e.g. Stoneham et al., 2003; Windle et al., 2009).

In order to rank the bids made by landholders in an auction, a metric is required to measure and compare the level of ES provided by alternative bids. A number of metrics have been developed for conservation auctions, such as habitat hectares and the biodiversity benefits index (e.g. Chomitz et al., 2006; Oliver et al., 2005; Parkes et al., 2003; Wünscher et al., 2008). These calculate the value of each bid in terms of ecological outcomes, and express it as a single unit. This means the auction mechanism can select the individual projects which provide the best value for money. However, by focussing on individual bids this approach will not necessarily select the optimal spatial configuration of conservation projects across a landscape (Gole et al., 2005).

In many cases the effective provision of ecosystem services requires a landscape-scale approach, rather than a focus on individual properties (Goldman et al., 2007). For example, connectivity between biodiversity conservation sites facilitates dispersal of biota, potentially increasing the contribution that individual management actions make toward the goal of viable populations. Although different species respond to connectivity in different ways (e.g. Hostetler, 1999; Lindborg and Eriksson, 2004), the spatial configuration of sites is often critical to the biological success of conservation efforts (e.g. Drielsma and Ferrier, 2009; Jiang et al., 2007; McAlpine et al., 2006) and the selection of projects should be considered at a landscape scale in order to achieve lasting biodiversity outcomes. Some ecological metrics do assign a value to connectivity. For example, a conservation auction in Australia's Desert Uplands region had connectivity as a major focus, and applied a metric which included a significant weighting for proximity to conserved patches of remnant vegetation

* Corresponding author at: GPO Box 284, Canberra 2601, Australia. Tel.: +61 2 6242 1640; fax: +61 2 6242 1705.
E-mail address: andrew.reeson@csiro.au (A.F. Reeson).

0921-8009/$ – see front matter © 2011 Published by Elsevier B.V.
doi:10.1016/j.ecolecon.2011.03.022

within the landscape and proximity to other bids (Windle et al., 2009).

The relative value of connectivity or permeability compared to other ecological attributes such as habitat area and condition will depend on the characteristics of the target species or community, such as dispersal ability and range requirements. Species which are poor dispersers may require connected habitat, while others may be able to make use of stepping stones across a fragmented landscape. Some degree of habitat connectivity is required for most conservation outcomes in the short term. In the medium and long term it is likely to be of even greater importance, allowing species and communities to progressively adjust their ranges in response to climate change (Mawdsley et al., 2009). The highly modified and fragmented nature of agricultural landscapes means that adapting to climate change may be particularly problematic for many species and communities.

Where there are landscape-scale objectives such as habitat connectivity the ecological metric required to prioritise proposed conservation projects becomes more complex. As the value of any one bid depends on which other bids end up in the final package, it is not possible to come up with a meaningful independent biodiversity value for an individual bid. Rather it is necessary to consider each possible combination of bids, and work out which combination provides the best biodiversity outcomes within the budget constraint. That is, an effective metric should provide a measure of combined value rather than individual value. An alternative, less computationally intensive, approach is to select projects iteratively, incorporating each newly selected site into the landscape context within which the remaining proposed sites are assessed (Barton et al., 2009). This interdependency between sites is not new to conservation biologists who have long worked within the principle of biodiversity complementarity, a calculus for the marginal contribution each site makes toward global biodiversity values (Faith, 1994; Sarkar et al., 2006).

This paper considers how ES auction mechanisms may be modified to address the combinatorial values inherent in landscape-scale biodiversity conservation. The following section considers the design of incentive mechanisms which can cost effectively deliver the coordination required for landscape-scale outcomes. Section 3 describes the experimental testing of some proposed alternative auction mechanisms, with the results presented in Section 4, followed by discussion of the policy implications in Section 5.

2. Auction Mechanisms

To address landscape-scale objectives in conservation auctions it is necessary to have a mechanism for coordinating the actions of individual landholders in order to maximise landscape synergies, for example by offering adjoining parcels of land to form a wildlife corridor. Coordinating the actions of autonomous agents is difficult as it requires them to have both information about the actions of others and an incentive to coordinate with them. A series of studies by Parkhurst, Shogren and others investigate the use of a 'smart subsidy', which is a fixed payment with an agglomeration bonus, to provide an incentive for neighbouring landholders to coordinate their bids (Parkhurst et al., 2002; Parkhurst and Shogren, 2005, 2007). In laboratory experiments the bonus mechanism was successful in prompting experimental participants to coordinate their actions for a number of simple spatial configurations. These approaches build on game theory in which the complete payoff matrix is known and/or private information of other agents' costs and benefits is available. With complete information, coordination may occur if it is a clear Nash equilibrium.

In more complex and realistic coordination experiments the bonus mechanism proved less effective (Parkhurst and Shogren, 2007). Where there is no clear equilibrium, agents will require an additional mechanism in order to coordinate their actions. In experimental games, iteration can promote coordination as agents acquire information on the strategies of others. For example, in diverse experimental designs subjects generally fail to attain the desired outcome in a one-shot game, but are successful in achieving the goal as the game is repeated (e.g. Clark and Sefton, 2001). Iteration has been shown to promote coordination by neighbouring landholders in economic experiments; coordination was more likely in later rounds of the experiment, when participants were able to use their experience from previous rounds (Parkhurst and Shogren, 2007). Iteration combined with incentives for coordination therefore has the potential to facilitate coordination among autonomous agents.

A conservation auction with multiple bidding rounds, in which landholders are provided with information on the location of bids from the previous round, offers a mechanism through which landholders can identify potential synergies with other bids and adjust their own bids accordingly (Rolfe et al., 2009; Windle et al., 2009). It could allow landholders to converge on a coordinated solution without having advance knowledge of each others' costs and likely strategies. In an auction setting, as opposed to a fixed payment scheme, landholders have an incentive to coordinate their bids even in the absence of a bonus. Provided the bid assessment process places a positive value on connectivity, bids which coordinate with others will have a greater chance of success. All things being equal, landholders should therefore attempt to submit bids which align with those of their neighbours.

However, auctions work by compelling landholders to compete, thereby revealing their costs and enabling the purchaser to select those projects with the lowest cost per unit of biodiversity. In a discriminatory price auction, bidders have incentives to inflate their bid prices above their true costs, depending on their expectations of their costs relative to other bidders, in order to seek a surplus (Abbink et al., 2006; Latacz-Lohmann and Van der Hamsvoort, 1997). If an auction is repeated, bidders' expectations will become more accurate and those with low costs may increase their surplus request. Experimental studies show that bidders' prices tend to rise over repeated discriminatory price auctions (Cason and Gangadharan, 2005; Cummings et al., 2004; Schilizzi and Latacz-Lohmann, 2007). This will compromise the ability of the auction to reveal low cost providers, eroding the efficiency benefits. There is evidence of this occurring in the US Conservation Reserve Program (Kirwan et al., 2005; Reichelderfer and Boggess, 1988).

There is also a danger that a mechanism intended to promote coordination among landholders may at the same time promote strategic behaviour. As information on other bids is revealed, some individuals will learn that their bid has particularly high value, for example by virtue of being integral to a potential corridor. This is likely to result in such bidders raising their prices and extracting more rent based on this information (Cason et al., 2003). The multi-round auction format also increases the likelihood of collusion among bidders (e.g. Burtraw et al., 2009; Fabra 2003). Therefore while multi-round auctions may overcome the coordination problem inherent in landscape-scale conservation, they also offer greater potential for collusion and rent seeking by bidders. As bidders inflate their prices the auction becomes less effective at identifying low cost suppliers and a budget-constrained buyer is able to secure fewer ES. There may therefore be a trade-off between promoting coordination over multiple rounds and minimising collusion and learned strategic behaviour. The more rounds the better the coordination of bids across the landscape, but the greater the learning (both of equilibrium prices and one's value in the landscape) and potential for collusion.

It is well established that relatively minor details in the design of auctions and other market institutions can have a major impact on market performance (e.g. Klemperer, 2002). The limited theoretical guidance on the design of multi-round auctions for conservation necessitates an experimental approach. Economic experiments allow alternative auction formats to be tested and compared. We set out to experimentally test this trade-off between coordination and collusion in a discriminatory price, multi-round auction in a scenario with landscape-scale combinatorial conservation values. Providing information is critical to enabling landholders to coordinate their bids. Identifying the locations of the most competitive bids can provide a basis for other participants to coordinate with. In our scenario, the most cost effective package of bids in each round are identified as provisional winners, with their locations revealed publicly to provide coordination opportunities in the subsequent bidding round.

We examine bidder behaviour over repeated rounds to determine to what extent efficiency improvements resulting from increased spatial coordination are offset by increased strategic bidding. Our null hypothesis is that the net efficiency of the auction will increase as the number of rounds increases, as spatial coordination improves while strategic behaviour remains constant. We measure the efficiency of the auction in terms of the simulated landscape conservation value—including combinatorial spatial values—it delivers within a budget constraint. Rent seeking by bidders (the degree to which their asking price exceeds their opportunity cost) provides a measure of strategic behaviour.

The multi-round auction process may cause participants to focus more on price competition than on modifying the configuration of their bids to coordinate with their neighbours. There is potential for those who find a corridor forming around them to try raising their price in order to extract some extra rent, behaviour which could hamper coordination and erode any efficiency benefits of the auction process. This strategic behaviour can be avoided if participants are prevented from modifying provisionally successful bids between rounds (a form of bid-improvement rule), though this may prompt some to submit higher prices initially. We test the impact of such a lock-in rule on auction efficiency and strategic bidding behaviour.

The multi-round auction format may allow participants greater opportunity to try out alternative bid strategies. A participant may initially submit a high bid price in the knowledge that if they are not provisionally successful they will have the opportunity to reduce it subsequently. If they know how many rounds there are, they know exactly how many opportunities they have to submit inflated prices before coming down closer to their opportunity cost if necessary. However, if the number of rounds is unknown this strategy becomes risky. An unknown end-point may result in some participants missing opportunities to make or modify bids; on the other hand it may reduce strategic behaviour, as there is always the chance that the auction will close and a participant who is holding out will end up missing out. As an additional treatment we compare multi-round auctions with known and unknown end-points, again measuring overall efficiency (conservation value achieved) and strategic behaviour (rent seeking by bidders).

3. Experimental Methods

Experimental economics was applied to test and compare a number of variations of multi-round auctions under controlled laboratory conditions. Software was developed to create a simulated landscape linked to an auction for land-use change (converting from agricultural production to conservation), with a simple combinatorial metric for selecting the optimal package of bids within a budget constraint. Experiments were run in university computer laboratories, with students taking on the role of landholders. The landscape consisted of 400 cells, with each cell assigned to one of ten properties. Participants were presented with a map showing the whole landscape, with the various property boundaries marked out (see Fig. 1). The landscape was homogeneous, with the same production and conservation values for every cell. This provided a simple, context-free landscape in which to test and compare alternative auction mechanisms. If a landholder chose to do nothing in the experiment they would receive the production values of the cells in their property at the end of each experimental 'year'—this represents a baseline income from agriculture. According to standard experimental economics protocols, participants were paid based on the income they 'earned' in the experiment, meaning decisions had real financial consequences.

To test the auction mechanisms, participants were told that they had the opportunity to rent out some, or all, of their land. Terms such as 'conservation' were avoided to keep the context as neutral as possible. If land was successfully rented out, the landholder would not receive its production value (100 experimental dollars per cell), but they were free to determine the payment they required for renting it. To offer their land for rent, participants could click on the cells they wished to offer, and then enter a price. To promote connectivity, the instructions stated 'The more cells your offer connects to, the more likely you are to be successful. This includes cells that your neighbours offer for rent." And "The more top-bottom connections in your offer, the more likely you are to be successful. This includes cells that your neighbours offer for rent." This provides an incentive to coordinate with neighbours. They were told in the initial instructions that if the price they asked for was less than the production value of the land they could lose money from entering the auction. They were also told that it was a competitive auction, so the higher their price, the less chance they would have of successfully renting their land. In each auction round participants had three minutes in which to enter their bids.

A global optimisation was used to select the package of bids which provided the best overall landscape conservation value within the budget constraint. The budget available was 25,000 experimental dollars (this parameter was not revealed to participants). The assessment metric consisted of a fixed value for each cell conserved plus a connectivity bonus, which added a weighting for connections between conserved cells. There was also a 'north–south' bonus, an extra weighting for connectivity in a north–south (i.e. top to bottom on the map) direction, to reflect situations in which connectivity in a particular direction is preferred. The overall conservation value for a landscape with a particular package of bids is the sum of the conservation value of each conserved cell plus its connectivity weighting bonuses. The conservation value for each cell was 2, with a weighting of 1 for each conserved cell it was connected to (either contiguously or through other conserved cells) and an additional weighting of 1 for top to bottom connections. The component values and weightings of the assessment metric were not revealed to participants. Participants were restricted to two bids each per round in order to limit the number of possible packages of bids which had to be searched for the optimum combination (ensuring that the optimisation took no more than a minute or two on a standard computer).

Once participants had submitted their bids, the combinatorial bid assessment metric was applied to select the package of bids that provided the best overall value, considering conservation value and connectivity. Once the calculation was complete, participants' screens were updated to show the results (Fig. 1). If the auction was not yet complete, bids that formed part of the best package were identified as 'provisional winners'. Participants could see the location of all provisional winners in the landscape, and their screen also labelled their own bids as either provisional winners or unsuccessful. This provided participants with information which they could use to coordinate their bids with those of their neighbours, for example by submitting a part of their property which aligned with a provisionally successful bid on a neighbouring property. The auction was then re-opened, and participants had the opportunity to modify their bids (by changing the price or the cells offered) or enter additional bids (still with a maximum of two). If participants chose to do nothing, their bids remained live. The bid assessment metric was then re-run to select the optimal package of bids. Participants could not communicate directly in any way during the experiment.

Auctions were run for two, three or four rounds. In the known end-point treatment the total number of rounds was announced prior to

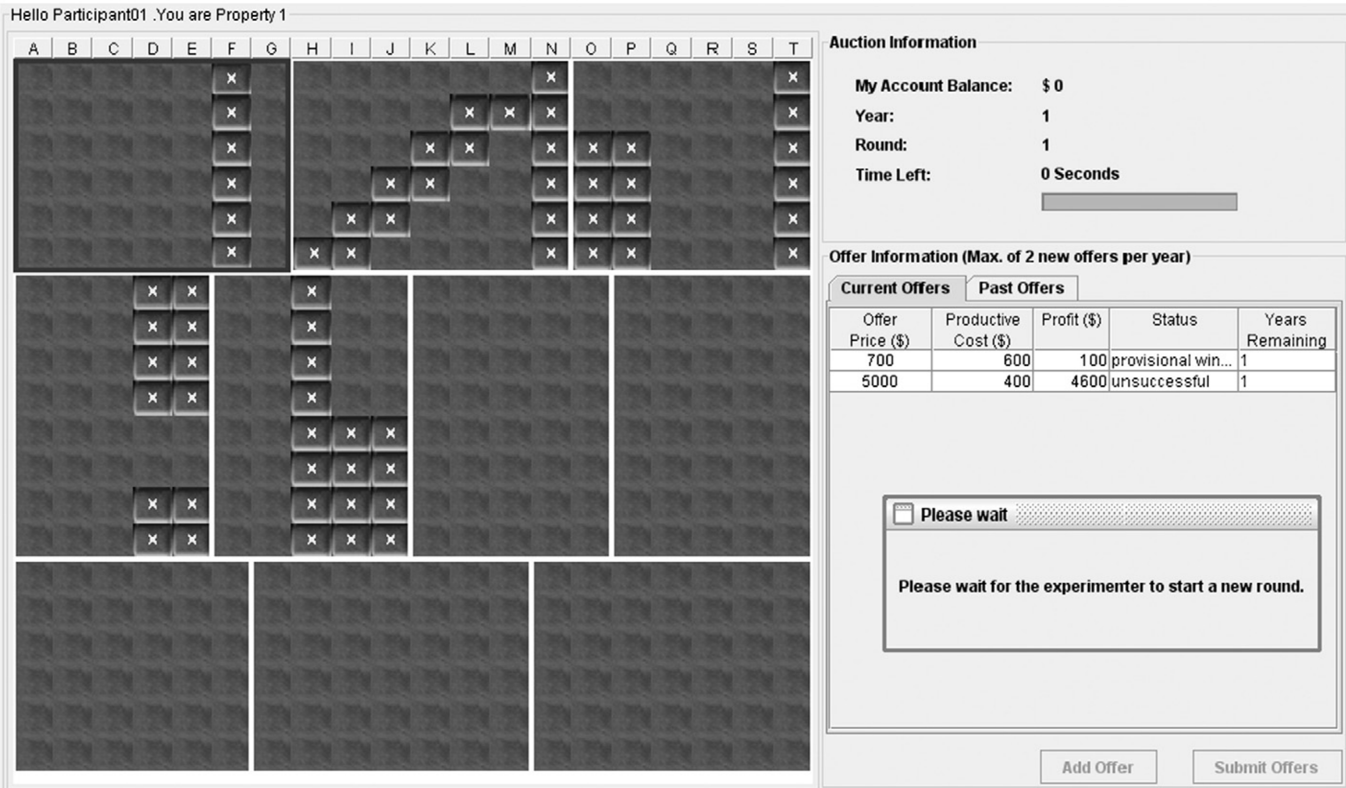

Fig. 1. Screen shot from the experimental scenario. The 'landscape' consists of 400 cells across 10 properties. In this screen shot the locations of the provisionally winning bids from the previous round (marked by crosses) are revealed to all participants prior to the start of the next round.

the start of the auction. In the unknown end-point treatment the number of rounds was not known to participants, with the auction ending without warning at the completion of a round. Another pair of treatments examined the effect of a rule in which provisional winners were prevented from modifying their bids between rounds. In the absence of this lock-in rule all bids could be modified between rounds, while with the lock-in rule neither the price nor the area offered could be adjusted for provisionally winning bids, nor could the bid be withdrawn. Data were collected from a total of 28 independent auctions, each with a different group of 10 participants. There were four principal treatment combinations with the known versus unknown end-point and the lock-in rule on or off; between five and nine replicates of each treatment combination were run (see Table 1)[1].

Table 1
Mean profit sought per bid submitted in the first and last rounds of each auction, by treatment combination (with standard errors); n = number of independent replicates of each treatment combination.

	No lock-in rule		Lock-in rule	
	First round	Last round	First round	Last round
Unknown end-point	341.4 (45.7) n = 5	349.4 (52.8)	384.1 (85.5) n = 6	248.7 (35.6)
Known end-point	414.0 (27.6) n = 8	356.8 (23.4)	431.9 (37.8) n = 9	363.3 (34.7)

[1] Numbers of replicates ended up uneven as not all experimental sessions successfully collected data.

4. Experimental Results

Individual bidding behaviour, and overall simulated conservation outcomes, were analysed with generalized linear models (GLM) using Genstat (13th edition). The level of rent seeking and simulated conservation values were used as measures of efficiency in the various auction mechanisms. The lower the rent seeking the better the auction performs in terms of revealing costs and minimising collusion, and hence efficiently allocating funding. The overall efficiency of each auction can be measured by comparing the conservation value achieved with the maximum possible conservation value. Mechanisms which promote increased connectivity will give more efficient outcomes, as will mechanisms which result in lower rent seeking (as more land can be acquired within the budget constraint).

4.1. Rent Seeking

Rent seeking was assessed by considering the profit (price requested—opportunity cost) in each bid. Table 1 shows the mean profit per bid submitted in the various treatment combinations, including all bids submitted in the auction. The area of each bid was included in the GLM to account for differences in profit between bids covering larger and smaller areas. Dummy variables were included in the models for the known/unknown end-point and the lock-in rule (on/off). The total number of rounds was included as a continuous variable. To avoid problems of repeated measures, analyses used a single bidding round of each auction replicate. Regression models are detailed in Table 2.

Considering bids from the first round of each auction, rent seeking was significantly greater when the number of rounds was known in advance ($p < 0.001$) (see Table 1). In the final round of each auction, rent seeking remained significantly higher when the total number of rounds was known to participants ($p = 0.016$). This is a surprising observation, as in the known end-point treatment participants were

Table 2
Regression models (standard errors in brackets).

	Rent seeking (rd 1)	Rent seeking (final rd)	Rent seeking (rd 1, known)	Rent seeking (final rd, known)	Auction efficiency[a] (final rd)
Constant	4.100** (0.185)	4.263** (0.317)	4.332** (0.289)	4.375** (0.321)	−70.6* (27.2)
Known end-point[b]	0.442** (0.109)	0.277* (0.115)	–	–	−4.74* (2.20)
Lock-in rule[c]	−0.145 (0.102)	−0.285** (0.108)	−0.160 (0.107)	−0.209 (0.121)	4.34* (1.99)
Area of bid	0.061** (0.008)	0.050** (0.008)	0.048** (0.008)	0.053** (0.009)	–
Number of rounds	–	0.016 (0.078)	0.166* (0.080)	0.041 (0.090)	0.57 (1.48)
Funds spent	–	–	–	–	0.007** (0.001)

[a] Auction efficiency expressed as % of the maximum.
[b] 0 = unknown end-point; 1 = known end-point.
[c] 0 = lock-in rule absent, 1 = lock-in rule present.
* Significant at 0.05.
** Significant at 0.01.

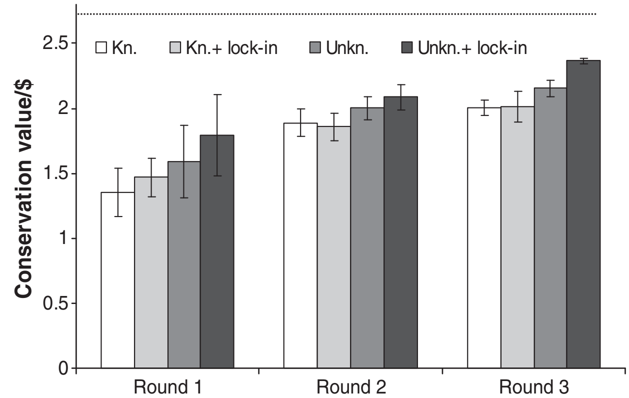

Fig. 2. Landscape conservation value (means +/− standard errors) by treatment combination (known and unknown end-point; with and without the lock-in rule), over the first three rounds of the auctions. The dashed line indicates the theoretical maximum.

fully aware that this was the final round, yet rent seeking remained higher than in the unknown end-point treatment, where there remained the possibility of additional rounds. The total number of rounds had no effect ($p = 0.84$).

Considering only the first round of each experiment in which the number of rounds was known, rent seeking showed a significant positive relationship with the total number of rounds ($p = 0.039$). By the final round this effect had disappeared ($p = 0.652$). This suggests that participants in longer auctions initially ask for higher prices in the knowledge that they will have more opportunity to subsequently reduce their price if they are not competitive. Therefore increasing the number of rounds will not necessarily improve overall efficiency.

The lock-in rule for provisional winners had no effect on rent seeking in the initial round ($p = 0.156$). By the final round, rent seeking was significantly lower where provisional winners were locked-in ($p = 0.009$). These results suggest that the lock-in rule does not cause people to raise their prices initially, even though they are prevented from subsequently raising their price if their bid is a provisional winner, but it does succeed in preventing provisional winners from seeking greater profits in subsequent rounds.

4.2. Auction Efficiency

The combinatorial metric used in the experiments provides a measure for the overall simulated landscape conservation value achieved under the various experimental treatments. Conservation (biodiversity) values were higher in the final round of each auction than in the initial round (paired *t*-test, $p = 0.03$). Fig. 2 shows the mean overall landscape conservation value achievable within the budget constraint, by treatment across the first three rounds of all the auctions.[2]

Auction efficiency data were analysed by GLM using the same treatment variables described above. 'Funds spent' was included as a covariate to account for small differences in the amount of available funding that was allocated to each optimal package, as the optimisation did not accept fractions of bids (any unspent budget was retained by the experimenter). The regression model is shown in Table 2. In the last round of each auction, efficiency was significantly higher when the total number of rounds was unknown to participants in advance ($p = 0.042$) and with the lock-in rule ($p = 0.040$). Efficiency achieved increased over rounds (see Fig. 2), though the total number of rounds did not prove significant in explaining the overall efficiency (or conservation value) in the final round (Table 2).

Participants did appear to be using the information provided during the auction to coordinate with other bids. The location of a bid within the landscape was adjusted on 39% of occasions (excluding the first round and locked-in bids); given opportunity cost did not vary across the landscape this is likely to reflect attempts to coordinate. To determine whether coordination did increase with repetition we compared the numbers of connectivity bonus points in the selected package of bids in the first and last rounds of our auctions. More bonus points were awarded in the last round than in the first (paired *t*-test, $p = 0.03$); the average increase was 8% from the first to the last round.

The decrease in rent seeking from the first to the last round was 19% on average across all the treatments. These results suggest that increased coordination and decreased rent seeking are both contributing to the improved efficiency seen in latter rounds of the auction. However, in the absence of a treatment in which coordination is prevented it is not possible to separate the impacts of reduced rent seeking in latter rounds from the benefits of improved coordination, as with reduced rent seeking more land is purchased so there is greater potential for bonus points.

5. Discussion

Multi-round auctions have the potential to address the key issues around the design of incentives for efficient provision of ecosystem services requiring complementary site actions to achieve landscape-scale outcomes. By spreading the auction over a number of rounds, with information about the location of other bids in the landscape provided between rounds, coordination can occur across the landscape without the need for advance knowledge of others' likely actions or additional incentives. Coordination by individual landholders is rewarded by an increased likelihood of success in the auction. At the same time the competitive nature of the auction mechanism encourages landholders to reveal their opportunity costs of carrying out conservation or other environmental management projects.

However, there are also risks in using multi-round auction mechanisms. Landholders are likely to learn more about their costs relative to other bidders over the course of the auction, resulting in those with lower than average costs inflating their prices. Conversely those who initially submit higher prices may learn to reduce their prices in order to be more competitive. The auction process is likely to be unfamiliar to most participants, so multiple rounds may provide an opportunity for them to learn about the auction mechanism itself and

[2] Round four is omitted due to its smaller sample sizes.

submit more considered bids in later rounds (even independent of strategic considerations) (see List and Shogren, 1999). In real-world applications (though not in our laboratory scenario), bidders are also likely to be uncertain about their true costs. In order to reduce the risk of submitting a bid which does not meet their costs they may initially submit higher prices, which they may revise down in the light of others' estimates of their own costs (Rolfe et al., 2009).

In our experiments, some participants appeared to focus more on price competition than coordination, strategically inflating their bid prices in order to seek additional rent. This was countered by the use of the lock-in rule, and not revealing the total number of rounds in advance. Having an unknown end-point means that engaging in strategic behaviour becomes a risky business. As the auction can end at any time, someone who enters an inflated bid may not get the chance to reduce it if they are unsuccessful, and so may miss out entirely. Rent seeking was substantially lower in the initial rounds under this treatment, which is likely to reflect uncertainty about the end-point. Lower bid prices increase the amount of ES (in this case, land conserved) that can be purchased in the auction, and hence improve overall cost effectiveness. A more surprising result was that rent seeking was lower when the number of rounds was unknown than in even the final round where the end-point was known. This suggests that the initial uncertainty has reduced strategic behaviour and prompted intense price competition.

The lock-in rule for provisionally winning bids also resulted in more cost effective outcomes. It effectively ensures that the impacts of learning are unidirectional—participants may learn to lower their prices if their bids are relatively expensive, but are restricted in their ability to increase their price if their initial bid is relatively cheap. Freeman and Woodward (2010) found that a similar rule led to higher initial bid prices, but this was not the case in our experiment. In our scenario this treatment increased overall simulated landscape conservation outcomes as it prevented provisional winners from increasing their prices without prompting them to be higher in the first place. The lock-in rule worked particularly well when the number of rounds was unknown—initial bid prices were lower, and the lock-in rule ensured they could not creep upwards.

A fruitful area for future research will be considering how the auction process may be spread over a number of years, similar in concept to the scheduling of priority conservation actions. In reality a funding agency will often have insufficient resources to achieve any significant degree of landscape connectivity in a single auction. Conservation plans and strategies are often implemented partially, in stages over a planning period with the number of sites and management actions constrained by a budget (Cowling and Pressey, 2001; Pressey and Taffs, 2001; Pressey et al., 2007). Conservation action scheduling can be represented as a multistage constrained optimization problem which aims to maximize the number of biodiversity types that have met their management targets by the end of a funding program. The choice of auction design will depend on the availability of resources, the urgency of reaching the objective and the importance of learning by participants over repeated auctions (Freeman and Woodward 2010).

An auction for ecosystem services is only as good as the metric used to assess bids. Further work is required to develop metric frameworks capable of assessing alternative landscape configurations at the appropriate level of detail (Ferrier and Drielsma, 2010). In many cases detailed ecological or biophysical knowledge is also needed for effective landscape-scale metrics. The combinatorial nature of the problem also creates a challenge for assessing bids, since the number of possible combinations rises rapidly. For example with just three bids (A, B, C) there are seven possible packages (A; B; C; AB; AC; BC; ABC). For five bids there are 31 possible packages, rising to over one thousand for 10 bids, one million for 20 and one billion for 30. Considerable computational power will therefore be required to assess even relatively small numbers of bids, and for larger numbers the problem becomes NP-hard and cannot be solved. Search heuristics such as genetic algorithms or simulated annealing can be applied to find approximate solutions in such cases (Hajkowicz et al., 2007; Seddon et al., 2010). A number of solutions to this problem have been developed for combinatorial auctions (see Cramton et al., 2006).

The policy recommendations from these initial experiments are clear. Multi-round auctions can deliver coordinated outcomes most efficiently where the number of rounds is unknown to participants in advance, and provisional winners cannot raise their prices. These simple rules should be applicable in the field, although it will clearly be more complex than a traditional conservation auction; they will also be applicable to other ecosystem services which have landscape-scale synergies (see Goldman et al., 2007). An agency may run the auction over a number of rounds, stopping once a desired ES target is reached. Clearly the transaction costs will increase with the number of rounds. Allowing bids to automatically carry over from one round to the next would minimise the extra transaction costs for participants imposed by the iterated process. Uncertainty about whether any particular round will be the last has been shown to reduce strategic bidding, which reduces the number of rounds required. These simple rules can enable complex landscape-scale objectives to be achieved in a relatively straightforward and cost effective manner.

Acknowledgements

We are grateful to Land and Water Australia for financially supporting this research, John Zhou and Raymond Xiao for software development, Mark Hughes, John Tisdell and Brian Scott for support and two anonymous referees for helpful comments on an earlier version of this manuscript.

References

Abbink, K., Brandts, J., Pezanis-Christou, P., 2006. Auctions for government securities: a laboratory comparison of uniform, discriminatory and Spanish designs. Journal of Economic Behavior and Organization 61, 284–303.
Barton, D.N., Faith, D.P., Rusch, G.M., Acevedo, H., Paniagua, L., Castro, M., 2009. Environmental service payments: evaluating biodiversity conservation trade-offs and cost-efficiency in the Osa Conservation Area, Costa Rica. Journal of Environmental Management 90, 901–911.
Burtraw, D., Goeree, J., Holt, C.A., Myers, E., Palmer, K., Shobe, W., 2009. Collusion in auctions for emission permits: an experimental analysis. Journal of Policy Analysis and Management 28, 672–691.
Cason, T.N., Gangadharan, L., Duke, C., 2003. A laboratory study of auctions for reducing non-point source pollution. Journal of Environmental Economics and Management 46, 446–471.
Cason, T.N., Gangadharan, L., 2005. A laboratory comparison of uniform and discriminative price auctions for reducing non-point source pollution. Land Economics 81, 51–70.
Chomitz, K.M., da Fonseca, G.A.B., Alger, K., Stoms, D.M., Honzák, M., Landau, E.C., Thomas, T.S., Thomas, W.W., Davis, F., 2006. Viable reserve networks arise from individual landholder responses to conservation incentives. Ecology and Society 11, 40.
Clark, K., Sefton, M., 2001. Repetition and signalling: experimental evidence from games with efficient equilibria. Economics Letters 70, 357–362.
Cowling, R.M., Pressey, R.L., 2001. Rapid plant diversification: planning for an evolutionary future. Proceedings of the National Academy of Sciences 98, 5452–5457.
Cramton, P., Shoham, Y., Steinberg, R., 2006. Combinatorial Auctions. MIT Press, Cambridge, MA.
Cummings, R.G., Holt, C.A., Laury, S.K., 2004. Using laboratory experiments for policymaking: an example from the Georgia irrigation reduction auction. Journal of Policy Analysis and Management 23, 341–363.
Drielsma, M., Ferrier, S., 2009. Rapid evaluation of metapopulation persistence in highly variegated landscapes. Biological Conservation 142, 529–540.
Fabra, N., 2003. Tacit collusion in repeated auctions: uniform versus discriminatory. Journal of Industrial Economics 51, 271–293.
Faith, D.P., 1994. Phylogenetic pattern and the quantification of organismal biodiversity. Philosophical Transactions of the Royal Society of London, Series B. 345, 45–58.
Ferrier, S., Drielsma, M., 2010. Synthesis of pattern and process in biodiversity conservation assessment: a flexible whole-landscape modelling framework. Diversity and Distributions 16, 386–402.
Freeman, M.A., Woodward, R.T., 2010. Sequential license buyback auctions: an experimental analysis. Marine Resource Economics 25, 355–372.

Goldman, R.L., Thompson, B.H., Daily, G.C., 2007. Institutional incentives for managing the landscape: inducing cooperation for the production of ecosystem services. Ecological Economics 64, 333–343.
Gole, C., Burton, M., Williams, K.J., Clayton, H., Faith, D.P., White, B., Huggett, A., Margules, C., 2005. Auction for landscape recovery. WWF-Australia. http://www.napswq.gov.au/publications/books/mbi/round1-project21.html 2005.
Hajkowicz, S.A., Higgins, A., Williams, K., Faith, D.P., Burton, M., 2007. Optimisation and the selection of conservation contracts. Australian Journal of Agricultural and Resource Economics 51, 39–56.
Hostetler, M., 1999. Scale, birds, and human decisions: a potential for integrative research in urban ecosystems. Landscape and Urban Planning 45, 15–19.
Jiang, Y., Swallow, S.K., Paton, P.W.C., 2007. Designing a spatially-explicit nature reserve network based on ecological functions: an integer programming approach. Biological Conservation 140, 236–249.
Klemperer, P., 2002. What really matters in auction design? Journal of Economic Perspectives 16, 169–189.
Kirwan, B., Lubowski, R.N., Roberts, M.J., 2005. How cost-effective are land retirement auctions? Estimating the difference between payments and willingness to accept in the Conservation Reserve Program. American Journal of Agricultural Economics 87, 1239–1247.
Latacz-Lohmann, U., Van der Hamsvoort, C., 1997. Auctioning conservation contracts: a theoretical analysis and an application. American Journal of Agricultural Economics 79, 407–418.
Lindborg, R., Eriksson, O., 2004. Historical landscape connectivity affects present plant species diversity. Ecology 85, 1840–1845.
List, J.A., Shogren, J.F., 1999. Price information and bidding behavior in repeated second-price auctions. American Journal of Agricultural Economics 81, 942–949.
McAlpine, C.A., Rhodes, J.R., Callaghan, J.G., Bowen, M.E., Lunney, D., Mitchell, D.L., Pullar, D.V., Possingham, H.P., 2006. The importance of forest area and configuration relative to local habitat factors for conserving forest mammals: a case study of koalas in Queensland, Australia. Biological Conservation 132, 153–165.
Mawdsley, J.R., O'Malley, R., Ojima, D.S., 2009. A review of climate-change adaptation strategies for wildlife management and biodiversity conservation. Conservation Biology 23, 1080–1089.
Oliver, I., Ede, A., Hawes, W., Grieve, A., 2005. The NSW Environmental Services Scheme: results for the biodiversity benefits index, lessons learned, and the way forward. Ecological Management and Restoration 6, 197–205.
Parkes, D., Newell, G., Cheal, D., 2003. Assessing the quality of native vegetation: the 'habitat hectares' approach. Ecological Management and Restoration 4, S29–S38.
Parkhurst, G.M., Shogren, J.F., Bastian, C., Kivi, P., Donner, J., Smith, R.B.W., 2002. Agglomeration bonus: an incentive mechanism to reunite fragmented habitat for biodiversity conservation. Ecological Economics 41, 305–328.
Parkhurst, G.M., Shogren, J.F., 2005. Does complexity reduce coordination? Applied Economics Letters 12, 447–452.
Parkhurst, G.M., Shogren, J.F., 2007. Spatial incentives to coordinate contiguous habitat. Ecological Economics 64, 344–355.
Pressey, R.L., Taffs, K.H., 2001. Scheduling conservation action in production landscapes: priority areas in western New South Wales defined by irreplaceability and vulnerability to vegetation loss. Biological Conservation 100, 355–376.
Pressey, R.L., Cabeza, M., Watts, M.E., Cowling, R.M., Wilson, K.A., 2007. Conservation planning in a changing world. Trends in Ecology and Evolution 22, 583–592.
Reichelderfer, K., Boggess, W.G., 1988. Government decision-making and program performance—the case of the conservation reserve program. American Journal of Agricultural Economics 70, 1–11.
Rolfe, J., Windle, J., McCosker, J., 2009. Testing and implementing the use of multiple bidding rounds in conservation auctions: a case study application. Canadian Journal of Agricultural Economics 57, 287–303.
Sarkar, S., Pressey, R.L., Faith, D.P., Margules, C.R., Fuller, T., Stoms, D.M., Moffett, A., Wilson, K., Williams, K.J., Williams, P.H., Andelman, S., 2006. Biodiversity conservation planning tools: present status and challenges for the future. Annual Review of Environment and Resources 31, 123–159.
Schilizzi, S., Latacz-Lohmann, U., 2007. Assessing the performance of conservation auctions: an experimental study. Land Economics 83, 497–515.
Seddon, J.A., Barratt, T.W., Love, J., Drielsma, M., Briggs, S.V., Gibbons, P., Ferrier, S., 2010. Linking site and regional scales of biodiversity assessment for delivery of conservation incentive payments. Conservation Letters 3, 415–424.
Stoneham, G., Chaudhri, V., Ha, A., Strappazzon, L., 2003. Auctions for conservation contracts: an empirical examination of Victoria's BushTender trial. Australian Journal of Agricultural and Resource Economics 47, 477–500.
Windle, J., Rolfe, J., McCosker, J., Lingard, A., 2009. A conservation auction for landscape linkage in the southern Desert Uplands, Queensland. Rangeland Journal 31, 127–135.
Wünscher, T., Engel, S., Wunder, S., 2008. Spatial targeting of payments for environmental services: a tool for boosting conservation benefits. Ecological Economics 65, 822–833.