Designing combinatorial exchanges for the reallocation of resource rights

Martin Bichler[a,1], **Vladimir Fux**[a], **and Jacob K. Goeree**[b]

[a]Department of Informatics, Technical University of Munich, 85748 Munich, Germany; and [b]AGORA Center for Market Design, University of New South Wales, Sydney, Kensington, NSW 2052, Australia

Edited by Ilya Segal, Stanford University, Stanford, CA, and accepted by Editorial Board Member Paul R. Milgrom December 4, 2018 (received for review February 6, 2018)

We describe the design and implementation of a combinatorial exchange for trading catch shares in New South Wales, Australia. The exchange ended a decades-long political debate by providing a market-based response to a major policy problem faced by fisheries worldwide: the reallocation of catch shares in cap-and-trade programs designed to prevent overfishing. The exchange was conducted over the Internet to lower participation costs and allowed for all-or-nothing orders to avoid fragmented share portfolios. A subsidy was distributed endogenously to facilitate the transfer of shares from inactive to active fishers. Finally, prices were linear and anonymous to ensure that sellers of identical packages received the same payments. These features were crucial to mitigate economic distortions from introducing catch shares and to gain broad acceptance of the program. However, they led to computationally challenging allocation and pricing problems. The exchange operated from May to July 2017 and effectively reallocated shares from inactive fishers to those who needed them most: 86% of active fishers' bids were matched and their share deficits were reduced by 95% in high-priority share classes. Similar reallocation problems arise in fisheries with catch-share systems worldwide as well as in other cap-and-trade systems for resource rights, e.g., water and pollution rights. The implemented exchange illustrates how computational optimization and market design can provide policy tools, able to solve complex policy problems considered intractable only a few years ago.

overfishing | fishery management | catch shares | market design

The sustainability of global fisheries is in danger. Advanced fishing technologies and increased demand for fish have resulted in substantial overfishing in the past (1). Overfishing has caused more ecological extinction than any other human influence on coastal ecosystems. Worm et al. (2) report that 29% of fish and seafood species have collapsed (i.e., their catch has declined by 90%) and 100% are projected to collapse by 2048, unless action is taken. Overfishing is a classic example of the tragedy of the commons (3) and is regularly listed as a prime environmental concern.

To counteract overfishing, policymakers worldwide have moved to programs that allocate shares of the catch quota to individuals or groups of fishers. First, a scientifically sound limit is set on the amount of fish that can be sustainably taken from a given fishery. Then a percentage of that limit is guaranteed to groups of individual fishermen. Recent empirical evidence shows that such catch-share programs successfully combat overfishing (4, 5).

Despite the advantages of catch-share programs, their implementation often poses challenges. Rosenberg (ref. 6, p. 165) writes that problems, "are usually related to the initial allocation of the shares, which inevitably results in winners and losers in the access to the fishery-catch quota… Further difficulties can arise when shares are traded after the initial allocation." One reason for flawed initial allocations is that they are typically based on historical catch numbers (7). This procedure, known as "grandfathering," helps to gain industry acceptance of the catch-share program as it does not impose immediate restrictions on fishers' activities. However, grandfathered allocations cause distortions when the industry undergoes structural change, when stocks deplete or recover, or when quotas are tightened. To restore efficiency, shares need to be reallocated, which often raises complex policy issues.

The fishing industry in New South Wales (NSW) exemplifies the need for the reallocation of catch shares. Over a decade ago, shares were assigned to more than 1,000 fishers. In some share classes, allocations were based on grandfathering while in others, shares were allocated "equitably" or uniformly; i.e., all fishers active in those share classes received the same number of shares. This assignment did not cause economic distortions at the time because the shares were not effectuated. All that was needed to fish was a business license.

Increasing concerns about overfishing led the NSW government to introduce the "linkage program" in 2017. This program links fishers' allowed catch to the shares they hold. There are over 100 different share classes describing different types of access rights across several regions. A share class may determine a permission to catch certain types and quantities of fish or stipulate allowed efforts (e.g., maximum number days of fishing per week, the number of hooks per line, or the number of nets).

The main goal of the linkage program was to establish an instrument that could curb overfishing in the long run (e.g., by tightening share quotas over time). Its goal was not to reduce overall catch in the short run. Indeed, to get the industry to accept the linkage program, it should have no immediate adverse consequences for fishers.

However, the initial allocation of shares, established over a decade earlier, would certainly cause major disruptions. The

Significance

Overfishing is a prime environmental concern. Catch share systems have been shown to be effective tools to combat overfishing. But flawed initial allocations and the subsequent need for reallocation of catch shares often create challenging policy problems. Market-based solutions have attracted interest but the complex institutional and political constraints necessitate market forms to avoid undesirable outcomes. The combinatorial exchange we designed and implemented is a market design for the reallocation of catch shares. The design was adopted because it addressed stakeholders' key requirements: fair pricing, avoiding fragmented share portfolios, and mitigating distortions from the introduction of catch shares via cost-effective and market-based subsidies.

Author contributions: M.B., V.F., and J.K.G. designed research, performed research, analyzed data, and wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission. I.S. is a guest editor invited by the Editorial Board.

[1]To whom correspondence should be addressed. Email: bichler@in.tum.de.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1802123116/-/DCSupplemental.

NSW fishing industry is characterized by the familiar 80–20 rule, i.e., less than 20% of the fishers do more than 80% of the catching. The top 20% most-active fishers would be severely constrained by the linkage program while inactive fishers would be unaffected, with adverse consequences for the industry's commercial viability. Some of the most-active fishers faced share deficits of up to 50%; i.e., they held only half the shares required to justify their catch levels. To mitigate economic distortions introduced by the linkage program, a reallocation of shares was needed. The policy problem was how to efficiently transfer shares from inactive fishers, some of whom wanted to exit the industry, to active fishers without causing financial distress for the latter.

One common approach is for the government to buy out shares from exiting fishers (8), which are then retired or redistributed among active fishers. Before our involvement this was the option the NSW government was considering. A total amount of A$15 million had been set aside to secure commitments from exiting fishers to forego the option of using their shares. These exiting fishers would then be allowed to sell their shares during a given period. If unsuccessful, their shares would be retired.

The idea of using a government buyout was dropped for several reasons. First, the government would have to use some "scoring rule" to compare sell offers across different share classes. Second, there would be severe adverse-selection problems: A buyout offer is more competitive when the shares have little value to their current owner, but they might have even less value to others. In other words, without any price guidance about the desirability of shares, there is no guarantee that the right shares are bought out. Third, decentralized bargaining would not likely have led to an efficient redistribution of shares. Fishers are geographically dispersed across the state, making it costly to find all efficient buyer–seller matches, especially since fishers hold portfolios of shares, and sellers thus need to be matched with multiple buyers. Finally, government buyouts are often challenged in court (9, 10) on the grounds that they favor specific market participants, e.g., because of the particular scoring rule that was used.

To avoid these issues we proposed a centralized exchange where fishers' supply and demand would determine share allocations (instead of some scoring rule based on imperfect knowledge of fishers' needs). As we discuss next, the exchange required several nonstandard features to ensure fair and efficient outcomes without causing financial distress for active fishers. Meeting these goals was crucial to gain industry acceptance of the linkage program, without which the government would lose the ability to curb overfishing in the future.

To minimize participation costs for geographically dispersed fishers, the exchange had to be run as an electronic market over the Internet. Creating separate markets for the different share classes would expose exiting fishers to the risk that they sell some, but not all, of their shares. This exposure problem might leave exiting fishers with a fragmented share portfolio and few proceeds. Also, regulation requires those fishers that remain in the industry to hold a minimum number of shares in a given share class, which creates additional exposure problems: If fishers purchase additional shares but are unable to meet the threshold, then their investments are lost.

To avoid exposure problems, the market had to offer the possibility of combinatorial buy and sell offers. A fisher wishing to exit the industry should be able to submit an "all-or-nothing" sell offer that specifies a total price for the entire share portfolio. Likewise, a fisher wishing to buy additional shares should be able to specify a minimum and a maximum quantity he or she would accept at the per-unit price specified.

While a combinatorial exchange avoids exposure problems, computational complexity becomes a challenge. In a combinatorial exchange, determining the optimal allocations is computationally hard (NP-hard). The size of NSW fisheries applications—600 participating fishers, more than 100 different share classes, and 1,280 bids—raised questions about computational tractability.

Another main requirement was to mitigate any distortions caused by the introduction of catch shares, especially for the most-active fishers. The government's linkage program forced them to purchase additional shares "simply to be able to catch tomorrow what they were catching today." Many of them were financially constrained and unable to cover the costs of eliminating their share deficits. To avoid financial distress for the most-active fishers, any use of the A$15 million fund had to be mainly targeted at them. Moreover, this had to be done in a market-based and cost-effective manner; i.e., subsidy levels had to be determined based on the bids and asks received and minimize total costs to the taxpayers. The introduction of endogenous subsidies led to significant challenges in the payment computations.

Finally, an important constraint was the use of linear and anonymous prices. The reason for using linear prices is that they are intuitive and tractable. With $L > 100$ different share classes, each containing $z \geq 1$ shares, there would be $(z+1)^L$ nonlinear prices in the market if packages were priced separately. The reason for using anonymous prices is that they are considered "fair" in the sense that they yield an "equal treatment of equals." It would have caused a political stir if fishers learned they paid more (received less) than a rival for the shares they bought (sold). Hence, anonymity and linearity of prices were key requirements. But, as is well known, in the presence of combinatorial or package bids such prices may preclude efficient outcomes.

To summarize, combinatorial offers, endogenous subsidies, and linear and anonymous pricing were nonstandard design desiderata that required a tailor-made design for the exchange.

Related Literature and Applications

The research reported in this paper belongs to the area of market design. This interdisciplinary field has had some marked successes, e.g., improving the matching of interns to hospitals, students to schools, and organ donors to patients, work that was awarded the Nobel Prize in Economic Sciences in 2012. Besides matching institutions without monetary transfers, market design has influenced the design of one-sided combinatorial markets—also known as combinatorial auctions. Such auctions are used in a variety of applications, including logistics and industrial procurement markets (11) and the sales of highly valuable spectra for wireless and mobile-phone applications (12, 13). There exists a significant body of research on combinatorial auctions (14), including those that use linear prices (15–17). In contrast, linear prices can lead to inefficiencies in combinatorial exchanges and little is known about the efficiency loss of linear approximations for specific applications.

Interestingly, anonymous and linear prices are being used in day-ahead electricity markets in Europe, which are important examples of combinatorial exchanges (see ref. 18 for a comprehensive overview of electricity market design and ref. 19 for a discussion of linear prices in European markets). Electricity production, for instance, exhibits substantial cost nonconvexities due to startup costs and minimum power output of power plants. Nonconvexities are also present on the demand side as buyers are typically interested in purchasing a certain quantity of electricity for several consecutive hours. This is why day-ahead energy markets allow for so-called "block bids," i.e., package bids that demand multiple units of single-hour supply.

The use of linear prices together with package bids in the context of electricity markets is reassuring in that it suggests efficiency losses might be small in large markets. There are important differences with the fisheries application, however. The endogenous distribution of a targeted subsidy, for instance, raised issues not present in electricity markets.

To evaluate the potential efficiency loss due to linear prices more generally, we conducted extensive computer simulations for an environment similar to the NSW fisheries application (20). These simulations also allowed us to evaluate the computational issues that arise in a large market with hundreds of fishers and over 100 share classes. We found that efficiency losses due to linear prices were less than 5% on average. Also, the allocation problems in all our simulations could be solved to (near) optimality after a few hours.

Designing the Exchange

This section describes how we dealt with the various market-design challenges (see *SI Appendix* for a formal model). Our design choices reflect the complex institutional and political constraints and were based on theoretical considerations as well as computer simulations.

Bid Language. A bid language describes the types of offers that can be submitted in the exchange. It should provide buyers and sellers with the flexibility to express their preferences adequately. A fisher intending to exit the industry should be able to submit an all-or-nothing offer that specifies a single ask price for the entire share portfolio. This ask price represents the minimum a seller wants for disposing of all shares. (In addition, exiting fishers received a fixed payment of A$20,000 for their business license.) An example for the case of three share classes, labeled A, B, and C, is shown in Fig. 1.

Fishers wishing to remain active might want to sell shares in unprofitable share classes while adding shares to other share classes. To sell part of their portfolio, a simple sell offer could be used. The seller specifies a quantity and a per-unit price for each of the share classes and the exchange treats the sell offers in the different share classes as independent; i.e., any subset of them can be successful. In contrast, the exit offer in Fig. 1 is either successful, in which case all shares are sold, or not, in which case no shares are sold.

Because synergies across share classes were less important to buyers, they were not allowed to submit all-or-nothing offers that included shares from multiple share classes. Instead, buyers could submit separate offers for different share classes, which specified, for each share class, a per-unit price and a quantity interval (e.g., to meet minimum-share requirements). Fig. 2 provides an example where a buyer competes in many share classes. In share class A, for instance, the buyer wants at least 125 units and at most 250 units and is willing to pay up to A$3 for each unit in this quantity interval.

This bid language was used in the computer simulations that showed (near) optimal solutions could be found in a matter of hours (20). In contrast, a bid language that allows for all-or-nothing buy offers involving multiple share classes would have caused serious computational issues as the matching between buyers and sellers is significantly more complex.

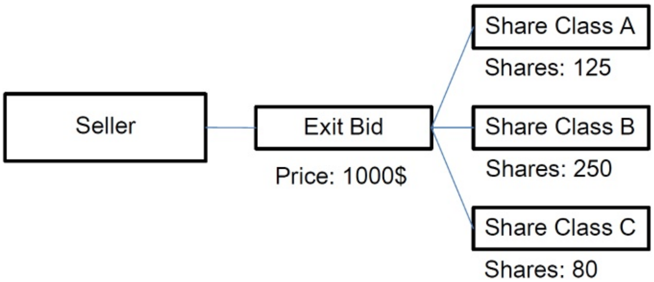

Fig. 1. All-or-nothing sell offer (exit offer).

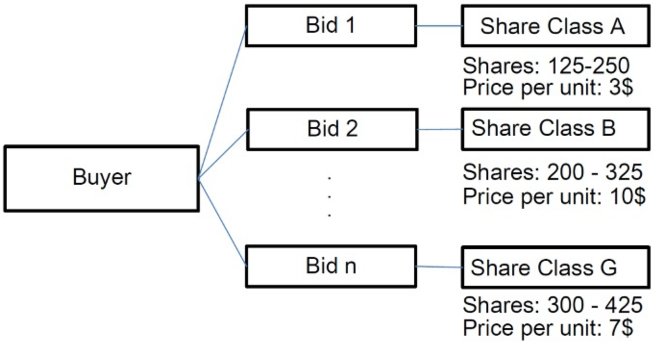

Fig. 2. Buy offer.

The proposed bid language also circumvents the "missing bids" problem that plagues some combinatorial auctions (21). This problem could have been severe in the fisheries exchange. As noted above, with $L > 100$ share classes, each containing $z \geq 1$ shares, there are $(z+1)^L$ possible packages. Buyers would be able to submit only a tiny fraction of all of the bids that an exclusive–or (XOR) bid language tries to elicit. The clearing mechanism would treat the missing bids as reflecting zero value, which results in inefficient allocations. Since synergies across share classes were low for the buyers, it sufficed to allow for separate buy offers for each share class. This vastly reduced the number of parameters a buyer had to specify and avoided the missing-bids problem.

Lexicographic Objective. The standard objective in two-sided markets is to maximize the weighted sum of buyer surplus, i.e., prices offered minus prices paid, and seller surplus, i.e., prices received minus prices asked. In the NSW case, full weight was given to the buyer side as the main purpose of the reallocation was to mitigate distortions for active fishers and bolster the industry's commercial viability. Many active fishers who needed to acquire large amounts of shares were financially constrained. To avoid outcomes with little or no trade, the NSW government decided to "grease the market" by injecting a subsidy Δ of up to A$15 million. Several priorities guided the distribution of the subsidy Δ.

Active fishers with deficits (priority P1). The main goal of the subsidy was to help active fishers reduce their share deficits. To this end, P1 maximized the number of shares bought by active fishers (up to their individual deficit levels) weighted by the per-unit prices they offered. The constraint that buyers paid no less than sellers asked was relaxed by augmenting buyers' per-unit prices with share-class–specific subsidies (under the constraint that the total amount of subsidy spent this way did not exceed Δ).

So active fishers with deficits were subsidized for any shares they bought toward reducing their share deficits. If they wanted additional shares, e.g., to expand their businesses, they were grouped together with other active fishers without deficits. Subsidizing this group was considered a next priority.

Active fishers without deficits (priority P2). The second goal was to help active fishers without share deficits to expand their businesses. To this end, P2 maximized the number of shares bought by these fishers weighted by the per-unit prices they offered. The constraint that buyers paid no less than sellers asked was again relaxed by adding share-class–specific subsidies to buyers' per-unit prices (under the restriction that the total amount of

ECONOMIC SCIENCES

subsidy spent this way did not exceed the residual subsidy left after running P1).

Any subsidy left after helping the active fishers was targeted at the exiting fishers.

Retiring fishers (priority P3). The third goal was to incentivize commercially nonviable businesses to exit the industry. Any residual subsidy left after running P1 and P2 was applied to fulfill as many additional exit offers as possible.

A final priority was to serve inactive fishers according to the competitiveness of their buy offers.

Inactive fishers (priority P4). The fourth goal was to maximize buyer surplus of the inactive fishers (this required no subsidy).

These priorities were accepted by all stakeholders. They led to a lexicographic objective that was maximized sequentially. After optimizing a certain priority, the resulting allocations were added as constraints in the optimization of the next priority (i.e., later priorities could only result in more shares being assigned to active fishers and to more successful exit bids). A number of other constraints were enforced in all of the optimization problems, as we discuss next.

Constraints. Demand–supply constraints guarantee that in every share class the total number of units bought does not exceed the total number of units sold. Individual rationality constraints make sure that no participant will incur a loss. This means that no winning seller should receive less than his ask price and no winning buyer should pay more than his bid. Also the government cannot spend more than the predefined subsidy, Δ, which we guaranteed via a total subsidy or budget constraint.

As detailed above, this subsidy was mainly targeted at active fishers by augmenting their bids with share-class–specific subsidies (priorities P1 and P2). As a result, active fishers paid less than the market price that sellers received, while inactive fishers paid the full market price. The share-class–specific subsidies were determined endogenously under the constraint they would fall into some share-class–specific range set by the government (e.g., 2–80% of the market price). Any residual subsidy was used to maximize the number of successful exit bids (priority P3), which required an additional A$20,000 to buy back the business license. A final use of the subsidy was to facilitate matches between exit offers and multiple buy offers. Artificial government buy bids were inserted in the market to complete imperfect matches but only if that raised the priority's objective function. To maintain anonymity and budget balancedness, the government paid the full market price for all of the shares bought. These shares were retired after the exchange was completed.

Linear and Anonymous Prices. While linear and anonymous prices were considered essential, they can lead to inefficiencies and paradoxically rejected bids.

To glean some intuition, consider a single seller who offers a package of three shares at a total price of A$9. Buyer 1 wants two shares and offers to pay A$8 while buyer 2 wants one share and offers A$2. Collectively, the buyers value the three shares more than the seller. However, supply is three at any price above A$3 while demand is three for prices lower than A$2 and demand is two for prices between A$2 and A$4. There is no linear and anonymous equilibrium price that clears the market and supports the welfare-maximizing allocation.

Besides efficiency losses, linear anonymous prices may also cause some offers to be "paradoxically rejected." Suppose, for instance, that a single seller offers two shares at a per-unit price of A$1. Buyer 1 wants two shares and offers to pay $5 while buyer 2 wants one share and offers A$3. Supply is two at any price above A$1 while demand is three for prices between A$1 and A$2.50 and demand is one for prices between A$2.50 and A$3. An exchange that maximizes total surplus, i.e., the total gains from trade, will assign two units to buyer 1 at a price of, say, A$1.75. But this means that buyer 2's offer is paradoxically rejected. It is important to understand the impact of such phenomena in large markets.

Prior literature provided little guidance as to how severe efficiency losses might be in large combinatorial markets; e.g., the economics literature mainly focuses on the restrictive conditions for existence of linear market-clearing prices (22). The decision to use linear and anonymous prices and accept the possibility of paradoxically rejected bids was mainly based on the computer simulations reported in ref. 20, which found small efficiency losses ($<5\%$) in large-scale markets. These design choices resemble those made for European day-ahead energy markets (19).

The computed prices are not necessarily unique, although the range of feasible prices is typically small. To resolve any indeterminacy, we ran a quadratic optimization program to determine unique prices such that the price differences across share classes were as small as possible.

Multiround Process. Most fishers were unfamiliar with any type of electronic trading, which raised concerns that a single-round process might result in poor outcomes. To allow for learning and equilibration, as well as price feedback, the government reserved the right to run the exchange for another round (up to three rounds). Fishers did not know whether there would be a next round, to mitigate gaming behavior.

Only minimal activity rules were imposed across rounds. Fishers had to submit (and confirm) an offer in this round to participate in the next round. They could modify their offers any number of times before the round's close and from one round to the next. This might include changing an exit offer to individual buy/sell offers, changing a buy offer to a sell offer (or vice versa), adding buy offers for new share classes, deleting previous offers, or changing the demanded quantities or per-unit prices. An offer submitted in a new round superseded offers from previous rounds. Offers that were left unchanged remained valid from one round to the next and could become successful in a later round even if they were not successful in an earlier round.

After a round closed, allocation, prices, and share-class–specific subsidies were computed for various scenarios. These scenarios were based on predetermined variations in parameters, e.g., total subsidy (Δ), the amount of subsidy reserved for buying out business licenses, and the bounds on the share-class–specific subsidies. An evaluation panel consisting of government officials and independent experts, with oversight from the project's independent probity advisor, then selected one of the scenarios based on the government's stated objectives and decided whether there would be another round.

After each round, fishers received feedback based on the selected scenario, e.g., which of their offers were successful, the market prices for each share class, and the subsidies they received. This feedback was sent via individual emails as well as on a secure website. If there was another round, fishers were

Table 1. Number of offers in each round

Offer	Round 1	Round 2	Round 3
Buy	747	739	740
Winning	158	322	446
Losing	589	417	294
Sell	421	439	432
Winning	35	90	131
Losing	386	349	301
Exit	95	101	107
Winning	34	51	62
Losing	61	50	45

Table 2. Changes in losing offers

Offer	Round 1 to round 2, %			Round 2 to round 3, %		
Buy	↑60.3	↓5.2	∼34.5	↑39.9	↓4.7	∼55.4
Sell	↑13.1	↓53.6	∼33.3	↑3.8	↓44.4	∼51.9
Exit	↑3.6	↓67.9	∼28.6	↑8.9	↓48.9	∼42.2

informed about its opening and closing times and were reminded that offers automatically carried over unless modified and that the new round might result in new share allocations and prices. After the final round, fishers were advised about the final outcomes (e.g., winning offers, share allocations, and market prices) and how to receive/submit payments.

Results

The exchange operated between May 1, 2017 and July 28, 2017 in three rounds. In each of these rounds buyers and sellers could revise their bids via a web-based bid submission system. (A description of the market rules and educational videos can be found at https://www.dpi.nsw.gov.au/fishing/commercial/reform/historical-docs/adjustment-subsidy-program.) Close to 600 fishing businesses registered for the market and placed 740 buy bids, 432 sell offers, and 107 package offers in the last round.

The exchange completely reformed the NSW fishery industry. Importantly, previously underused shares were successfully transferred to active fishers: 86% of their buy bids were matched and their overall share deficit was reduced by more than 75% overall and up to 95% in high-priority share classes. In addition, 60% of the package offers were matched: 62 businesses successfully sold all their shares and were able to exit the industry, receiving A$10.1 million in total (which included the A$20,000 for their business licenses). Around A$5.9 million went to sellers who sold only part of their share portfolio. While the government had set aside A$15 million, they spent only A$11.6 million, because higher subsidy levels did not significantly raise their primary objectives. This saved the taxpayer A$3.4 million.

Because of linear and anonymous prices and the subsidy, buyers paid less than their bids. In total, winning buyers bid A$6.2 million but paid A$3.2 million. Similarly, winning sellers asked for A$12.7 million in total, but received payments of A$14.8. A subsidy of A$11.2 million enabled these trades while implementing anonymous and linear prices (an additional A$0.4 million was spent on government buys used to complete matches between exit offers and multiple buy offers).

The submitted offers and the resulting market prices varied widely across share classes. A majority of the share classes traded at prices less than $200 but in a few share classes the share price was several thousand Australian dollars.

We next analyze the changes that occurred across the three rounds. Table 1 shows the number of offers submitted in each round. Interestingly, this number is almost the same in each round, although there were no activity rules used across rounds. Since bidders could not know which round would be the final round, they participated actively from the start. The number of winning offers increased because offers became more competitive and the government raised the total subsidy over the three rounds (from A$6 million in round 1 to A$11 million in round three).

Table 2 shows for each type of offer the fraction of losing offers that increased (↑), decreased (↓), or stayed the same (∼) across rounds. From round 1 to round 2, bids were raised for the majority of the losing buy offers while very few bids were lowered. Likewise, the asks were lowered for most of the losing sell offers (whether they were exit or simple sell offers) while very few asks were raised. The same is true from round 2 to round 3, although the effects are somewhat less pronounced as a large fraction of the offers were left unchanged. In contrast, the vast majority (65–85%) of the winning buy and sell offers were unchanged from round 1 to round 2 and from round 2 to round 3.

Table 3 shows the impact of these changes in offers on the objectives (*Lexicographic Objective*), assuming the parameters from the selected scenario in the final round apply to all three rounds so that any improvements in the objectives are due to offers becoming more competitive. The priorities are measured in Australian dollars except for priority P3, which corresponds to the number of successful exit bids. Note that this number increases almost 30% from round 1 to round 3. A similar conclusion holds for the other priorities. Priority P1 increases more than 50% and priority P2 almost 30% over the three rounds. Finally, priority P4 increases sixfold over the three rounds although the total value represented by this priority is rather small.

These results underscore the value of a multiround process. Unfamiliarity with the trading rules and overly optimistic expectations about the prices that shares can be sold (or bought) at resulted in relatively poor performance in the first round. The price feedback provided after each of the first two rounds stimulated more competitive offers (Table 2), resulting in more efficient allocations. In particular, after the third round, active fishers' deficits were successfully reduced (up to 95% in high-priority share classes).

Conclusion

Catch-share systems are an important policy tool to curb overfishing (5). However, their effectiveness can be undermined by a flawed initial allocation of shares, e.g., when shares are distributed uniformly or via a grandfathering process. Reallocating shares may result in substantial efficiency gains but also raises complex policy issues arising from institutional and political constraints.

The reallocation of catch shares in NSW, for instance, posed several major design challenges: avoiding fragmented share portfolios, mitigating financial distress for active fishers with large share deficits, and treating equals equally. The exchange we designed and implemented addressed these issues through the use of all-or-nothing offers, endogenously determined subsidies, and linear and anonymous prices.

Our design was informed by extensive computer simulations (20), which provided reassurance that even in large combinatorial markets, (near) optimal allocations could be found in a matter of hours and that efficiency losses due to linear prices would likely be less than 5%. These simulation results were corroborated by the real exchange where the allocation problem was solved to full optimality and the efficiency loss due to linear prices was only 1.1%.

The exchange successfully reformed the NSW fishing industry. Active fishers were able to reduce their share deficits and more than 60 businesses voluntarily exited the industry by selling their entire share portfolios. The government had set aside up to A$15 million to facilitate the transfer of shares from inactive to active fishers but ended up spending only A$11.6 million, saving the taxpayer A$3.4 million.

These results underscore how computational optimization and market design can provide tools to solve complex policy problems. The traditional approach to reallocating shares, i.e., a

Table 3. Changes in priorities

Priority, by fisher type	Round 1	Round 2	Round 3
P1) Active with deficit	A$2,950,177	A$3,516,733	A$4,445,633
P2) Active without deficit	A$1,267,387	A$1,586,368	A$1,634,024
P3) Exiting	48	55	62
P4) Inactive	A$12,203	A$24,727	A$71,358

ECONOMIC SCIENCES

government buyout, would have applied the full A$15 million to buy shares without having precise information about demand-side preferences. As a result, government buyouts yield unsatisfactory outcomes that are regularly challenged in court. The responsible Minister, the Honorable Niall Blair, put it as follows: "I feel very confident in stating that we have been able to achieve a result for the taxpayer of NSW and the shareholders in the commercial fishing industry that is far beyond what any traditional approach to industry reform would have achieved."*

The NSW fishing industry is not the only example where a reallocation of catch shares was needed to correct a flawed initial distribution. In 2014, there were 154 fisheries worldwide that had introduced catch-share systems, and this number is growing (7). There will be a need for the reallocation of catch shares in many of these fisheries and the exchange we designed could be used for this. More generally, the exchange can serve as a template for the reallocation of resource rights in other applications, e.g., water rights, pollution rights, and environmental offsets, where similar design challenges arise.

*Blair N (2017) Personal letter by Minister Niall Blair.

ACKNOWLEDGMENTS. We thank Doug Ferrell and the Department of Primary Industries in NSW for their support.

1. Jackson JBC, et al. (2001) Historical overfishing and the recent collapse of coastal ecosystems. *Science* 293:629–637.
2. Worm B, et al. (2006) Impacts of biodiversity loss on ocean ecosystem services. *Science* 314:787–790.
3. Hardin G (1968) The tragedy of the commons. *Science* 162:1243–1248.
4. Melnychuk MC, et al. (2012) Can catch share fisheries better track management targets? *Fish Fish* 13:267–290.
5. Birkenbach AM, Kaczan DJ, Smith MD (2017) Catch shares slow the race to fish. *Nature* 544:223–226.
6. Rosenberg AA (2017) Marine conservation: The race to fish slows down. *Nature* 544:165–166.
7. Lynham J (2014) How have catch shares been allocated? *Mar Policy* 44:42–48.
8. Squires D (2010) Fisheries buybacks: A review and guidelines. *Fish Fish* 11:366–387.
9. Iudicello S, Lueders SB (2016) A survey of litigation over catch shares and groundfish management in the pacific coast and northeast multispecies fisheries. *Environmental Law* 46:157–208.
10. Kompas T, Gooday P (2007) The failure of 'command and control' approaches to fisheries management: Lessons from Australia. *Int J Glob Environ Issues* 7:174–190.
11. Sheffi Y (2004) Combinatorial auctions in the procurement of transportation services. *Interfaces* 34:245–252.
12. Leyton-Brown K, Milgrom P, Segal I (2017) Economics and computer science of a radio spectrum reallocation. *Proc Natl Acad Sci USA* 114:7202–7209.
13. Bichler M, Goeree JK (2017) *Handbook of Spectrum Auction Design* (Cambridge Univ Press, Cambridge, UK).
14. Cramton PC, et al. (2006) *Combinatorial Auctions* (MIT Press, Cambridge, MA), Vol 475.
15. Kwasnica T, Ledyard JO, Porter D, DeMartini C (2005) A new and improved design for multi-objective iterative auctions. *Manage Sci* 51:419–434.
16. Brunner C, Goeree JK, Holt CA, Ledyard JO (2010) An experimental test of flexible combinatorial spectrum auction formats. *Am Econ J Microeconomics* 2:39–57.
17. Scheffel T, Pikovsky A, Bichler M, Guler K (2011) An experimental comparison of linear and nonlinear price combinatorial auctions. *Inf Syst Res* 22:346–368.
18. Cramton P (2017) Electricity market design. *Oxford Rev Econ Policy* 33:589–612.
19. Meeus L, Verhaegen K, Belmans R (2009) Block order restrictions in combinatorial electric energy auctions. *Eur J Oper Res* 196:1202–1206.
20. Bichler M, Fux V, Goeree J, A matter of equality: Linear pricing in combinatorial exchanges. *Inf Syst Res*, in press.
21. Bichler M, Goeree JK, Mayer S, Shabalin P (2014) Spectrum auction design: Simple auctions for complex sales. *Telecommun Policy* 38:613–622.
22. Gul F, Stacchetti E (1999) Walrasian equilibrium with gross substitutes. *J Econ Theor* 87:95–124.

Online Appendix: Designing Combinatorial Exchanges for the Reallocation of Resource Rights

Martin Bichler[a,1], **Vladimir Fux**[a], **and Jacob K. Goeree**[b]

[a]Department of Informatics, Technical University of Munich, Germany; [b]AGORA Center for Market Design, University of New South Wales, Australia

This manuscript was compiled on December 28, 2018

This online appendix provides the mathematical program to compute allocation and prices.

Overfishing | Fishery Management | Catch Shares | Market Design

Appendices.

1. Notation

We denote the set of all buy bids as $\mathcal{B}$; this set can be naturally divided into set of active buy bids $\mathcal{B}^{\mathcal{A}}$ and set of inactive buy bids $\mathcal{B}^{\mathcal{I}}$, which have no common elements, i.e. $\mathcal{B}^{\mathcal{A}} \cap \mathcal{B}^{\mathcal{I}} = \emptyset$. Each buy-side bid $j \in \mathcal{B}$ is related to only one share class $l \in \mathcal{L}$ and is a tuple $< l, {b_j}^l, \underline{\beta}_j^l, \overline{\beta}_j^l >$, where ${b_j}^l$ is unit price, which a fisher can pay for a single share, and $\underline{\beta}_j^l$ and $\overline{\beta}_j^l$ are the lower and the upper bounds on the number of units fisher wants to acquire. In addition, with each active buy-side bid $j \in \mathcal{B}^{\mathcal{A}}$ a deficit value d_j^l is associated. This deficit value is the shortage of shares for this particular fisher in a share class l, which is computed by the government prior to the auction.

Further, we denote by $\mathcal{S}$ the set of individual sell side bids and by $\mathcal{S}^E$ the set of exit bids (package bids). An individual bid $i \in \mathcal{S}$ in a share class $l \in \mathcal{L}$ consists of a quantity q_i^l that a fisher wants to sell and a unit price s_i^l she wants to get for a single share at a minimum.

An exit bid $i \in \mathcal{S}^E$ consists of a package vector $q_i = \{q_i^l\}_{l \in \mathcal{L}}$, which contains the number of units owned by a retiring fisher in each share class (some elements of the vector are zeros) and the total price s_i which a fisher wants to get for his package at a minimum. In addition, if an exit bid is accepted, a compensation for a fishing license p^E is payed from the subsidy that the government provides to the market.

With each bid we associate an allocation variable, which describes whether a particular bid is accepted, and, in case of buy bids, how many units are allocated. A buy bid is represented by an integer variable β_j^l. To have a feasible allocation, the value of β_j^l needs to be within the quantity bounds specified by the buyer, $[\underline{\beta}_j^l, \overline{\beta}_j^l]$ or zero. For this, we introduce an additional binary variable ζ_j^l which is 1 if and only if the corresponding β_j^l is positive:

$$\beta_j^l \le \overline{\beta}_j^l \zeta_j^l \qquad \forall j \in \mathcal{B}, l \in \mathcal{L} \qquad \text{[BQty1]}$$

$$\underline{\beta}_j^l \zeta_j^l \le \beta_j^l \qquad \forall j \in \mathcal{B}, l \in \mathcal{L} \qquad \text{[BQty2]}$$

Acceptance of an individual sell or exit bid is determined by binary variables σ_i^l where $i \in \mathcal{S}$ and $\sigma_i \in \{0,1\}$ where $i \in \mathcal{S}^E$ respectively.

2. Objective Function

The lexicographic policy goals are described as objective functions in a series of mathematical programs. This means that priority $k+1$ is taken into account after priority k. Let's describe each priority as an objective function for a mixed-integer linear program. In what follows "main constraints" refers to the constraints of the problem described in the subsequent sections.

Priority P1 focuses on buy bids with a deficit, and can be written as:

$$\max \sum_{j \in \mathcal{B}^{\mathcal{A}}} \sum_{l \in \mathcal{L}} (\kappa_j^l b_j^l) \qquad \text{[Priority 1]}$$

$$\text{s.t. } \kappa_j^l \le \beta_j^l, \kappa_j^l \le d_j^l \qquad \forall j \in \mathcal{B}^{\mathcal{A}}, l \in \mathcal{L} \qquad \text{[Deficit]}$$

(Main constraints)

So the goal is to maximize the number of units bought, up to a deficit level for active fishers, weighted by the bid price. The new variable κ_j^l is introduced in order to consider the buy bid of an active fisher only up to her deficit level d_j^l. As a result of the optimization, we get an allocation $(\{\beta_j^{l,P1}\}_{j \in \mathcal{B}^{\mathcal{A}}})$, which becomes a new constraint for priority 2. This new constraint ensures that the buyers with a deficit are satisfied as far as possible and also later priorities and optimization runs do not negatively impact this policy goal.

Priority P2 considers all active buyers, and thus we formulate the objective function as

$$\max \sum_{j \in \mathcal{B}^{\mathcal{A}}} \sum_{l \in \mathcal{L}} (\beta_j^l b_j^l), \qquad \text{[Priority 2]}$$

$$\beta_j^l \ge \beta_j^{l,P1} \qquad \forall j \in \mathcal{B}^{\mathcal{A}} : d_j^l > 0 \qquad \text{[Transition 1]}$$

(Main constraints)

Note that Transition 2 constraint implies Transition 1. Priority P3 aims to maximize number of accepted exit bids, i.e. package bids of bidders who want to leave the market:

$$\max \sum_{i \in \mathcal{S}^E} \sigma_i, \qquad \text{[Priority 3]}$$

$$\beta_j^l \ge \beta_j^{l,P2} \qquad \forall j \in \mathcal{B}^{\mathcal{A}} \qquad \text{[Transition 2]}$$

(Main constraints)

Priority P4 just maximizes the total volume of shares bought, such that even inactive fishers are taken into account:

$$\max \sum_{j \in \mathcal{B}} \sum_{l \in \mathcal{L}} (\beta_j^l b_j^l), \qquad \text{[Priority 4]}$$

$$\beta_j^l \ge \beta_j^{l,P2} \qquad \forall j \in \mathcal{B}^{\mathcal{A}} \qquad \text{[Transition 2]}$$

$$\sigma_i \ge \sigma_i^{P3} \qquad \forall i \in \mathcal{S}^E \qquad \text{[Transition 3]}$$

(Main constraints)

[1]Martin Bichler. E-mail: bichlerin.tum.de

Priorities P1-P4 were solved in a sequence. Overall, due to constraints on allocation, the solution to each priority should satisfy objective values of the previous priority solutions. To speed up the overall process, we allowed to "warm start" the solution process of later priorities. This means, we provide the solution of the last priority to the solver. The last solution is feasible also for the new priority, but may be improved with respect to the new objective.

Let us now discuss the main constraints that we need to consider in all models for all priorities.

3. Demand-Supply Constraint

We need to guarantee that in every share class the total number of units bought does not exceed the total number of units sold. This is implemented via the following demand-supply constraint:

$$\theta^l + \sum_{j \in \mathcal{B}} \beta_j^l = \sum_{i \in \mathcal{S}^E} q_i^l \sigma_i + \sum_{i \in \mathcal{S}} q_i^l \sigma_i^l, \qquad \forall l \in \mathcal{L} \qquad \text{[Demand-Supply]}$$

where θ^l is an integer-valued variable describing the artificial demand by the government, which facilitates trades in cases where sell-side bids do not have matching demand on the buy side. This variable can only be positive when it is beneficial for the current objective function (one of priorities 1-4).

For all units bought, the government has to pay the full market price, because we have anonymous prices and also require the budget to be balanced. Such shares bought by the government were sunset after the market and the money paid by the government to sellers had to come from the total subsidy.

4. Individual Rationality Constraints

An important requirement for the market is to guarantee that no participant will incur a loss. This means that no winning seller should receive less than his quoted ask price and no winning buyer is paying more than his bid. Thus for winning buyers, the following constraint should be satisfied

$$\rho^l \leq b_j^l, \qquad \forall j \in \mathcal{B} : \beta_j^l > 0, l \in \mathcal{L} \qquad \text{[IRB']}$$

where ρ^l is continuous variable representing linear price in share class l. Active fishers may pay a subsidized price, i.e. the government pays a discount on the market price δ^l in share class l:

$$\rho^l \leq b_j^l + \delta^l. \qquad \forall j \in \mathcal{B}^{\mathcal{A}} : \beta_j^l > 0, l \in \mathcal{L} \qquad \text{[IRBA']}$$

Note that these constraints should be satisfied only for winning buyers. We rewrite them as valid linear constraints using big-M constraints:

$$\rho^l \leq b_j^l + \delta^l + (1 - \zeta_j^l)M \qquad \forall j \in \mathcal{B}^{\mathcal{A}}, l \in \mathcal{L} \qquad \text{[IRBA]}$$

$$\rho^l \leq b_j^l + (1 - \zeta_j^l)M \qquad \forall j \in \mathcal{B}^{\mathcal{I}}, l \in \mathcal{L} \qquad \text{[IRBI]}$$

Similarly, we can formulate individual rationality constraints for the sell-side:

$$s_i \sigma_i \leq \sum_{l \in \mathcal{L}} q_i^l \rho_i^l \qquad \forall l \in \mathcal{L}, \forall i \in \mathcal{S}^E : \qquad \text{[IRP]}$$

$$s_i^l \sigma_i^l \leq \rho^l \qquad \forall i \in \mathcal{S} \qquad \text{[IRS]}$$

5. Budget Constraint

To encourage active participation, the government provides a subsidy Δ, which is distributed endogenously. Some amount Δ^E of this subsidy was determined for buying out fishing licenses of those fishers who exit the market:

$$\sum_{i \in \mathcal{S}^E} p^E \sigma_i \leq \Delta^E \qquad \text{[Exit subsidy]}$$

The residual subsidy Δ^R (i.e. $\Delta = \Delta^R + \Delta^E$) covered discounts for active buyers δ^l and the money necessary for buying shares through the government which helped facilitate sales of package bids where there was no matching buy-side quantity in some of the share classes (θ):

$$\sum_{j \in \mathcal{B}^A} \beta_j^l \delta^l + \sum_{l \in \mathcal{L}} (\rho^l - \delta^l)\theta^l \leq \Delta^R \qquad \text{[Subsidy]}$$

Unfortunately, the previous constraint contains the product of continuous and integer variables, which can only be linearized through introduction of a big number of additional binary variables and constraints.Therefore, we rewrote the constraint in such a way that only inactive bids are linearized (the number of inactive participants was expected to be lower than the number of active fishers). We multiplied the amount sold to sellers σ_i with the discount δ^l and subtracted the amount bought by inactive buyers $\sum_{j \in \mathcal{B}^I} \beta_j^l \delta^l$, which also yields the discounts for the active buyers:

$$\sum_{i \in \mathcal{S}^E} \sum_{l \in \mathcal{L}} q_i^l \delta^l \sigma_i + \sum_{i \in \mathcal{S}} \sum_{l \in \mathcal{L}} q_i^l \delta^l \sigma_i +$$
$$\sum_{l \in \mathcal{L}} (\rho^l - \delta^l)\theta^l - \sum_{j \in \mathcal{B}^I} \beta_j^l \delta^l \leq \Delta^R \qquad \text{[Subsidy]}$$

This constraint is still non-linear, and all 4 terms contain products of two variables. However, the the total number of variables after linearization should be lower than in the first version of the constraint because of smaller number of inactive bids. The product of binary and integer variables in terms 1 and 2 ($\delta^l \sigma_i$) are straightforward to linearize. The terms 3 and 4 are more challenging as they are products of integer and continuous variables. We first convert the integer variables θ^l and β_j^l to a sum of binary variables:

$$\theta^l = \sum_{n=0}^{logN/2} 2^n \theta_n^l \qquad \text{[Government Bids]}$$

$$\beta_j^l = \underline{\beta}_j^l + \sum_{n=0}^{log(\overline{\beta}_j^l - \underline{\beta}_j^l)/2} 2^n \mu_{j,n}^l \qquad \text{[Inactive Bids]}$$

Then we can linearize these terms as products of binary and continuous variables. The number of shares the government can buy was restricted by a parameter N in constraint (Government bids). This parameter N was set to 50% of the supply in a share class.