

Timeshare Exchange Mechanisms

Yu Wang, Aradhna Krishna

Ross School of Business, University of Michigan, 701 Tappan Street, Ann Arbor, Michigan 48109
{yuwz@umich.edu, aradhna@umich.edu}

This paper focuses on the timeshare industry, where members own timeshare “weeks” and can exchange these weeks among themselves without a medium of exchange (such as money). Timeshare exchanges allow for the weeks to be redistributed among members to better match their preferences and thus increase efficiency. As such, the problem falls within the domain of matching problems, which have recently gained much attention in academia. We demonstrate theoretically that the two major timeshare exchange mechanisms used currently (deposit-first mechanism and request-first mechanism) can cause efficiency loss. We propose an alternate exchange mechanism, the top trading cycles chains and spacebank (TTCCS) mechanism, and show that it can increase the efficiency of the timeshare exchange market because TTCCS is Pareto efficient, individually rational, and strategyproof. We test the three exchange mechanisms in laboratory experiments where we simulate exchange markets with networked “timeshare members.” The results of the experiments are robust across four different environments that we construct and strongly support our theory. The research focuses on an industry not studied earlier within academia and extends theoretical work on mechanism design to cases where supply of resources is dynamic, but resources can be stored.

Key words: mechanism design; timeshare; experimental economics; efficiency; simulated market

History: Accepted by Jagmohan S. Raju, marketing; received September 10, 2004. This paper was with the authors 7 months for 2 revisions.

1. Introduction

Many business systems entail resource allocation across agents, where agents have preferences over scarce resources. For example, dormitory rooms are assigned to university students, oversubscribed courses are allocated to students, kidneys are matched between donors and patients, etc. For the systems mentioned above, researchers have proposed theoretically superior mechanisms with improved efficiency (e.g., Abdulkadiroglu and Sonmez 1999). Experimental work has then tested whether the theoretically superior mechanisms are indeed better in practice and how the market of agents behaves, i.e., how closely it adheres to the behavior predicted by the models (e.g., Krishna and Unver 2004).

In this paper we focus on the timeshare industry, where developers sell the right to use a specific property for a certain time interval to individuals who then become “members” of timeshare exchange organizations. These members can then exchange the rights to use these properties among themselves. Thus, a timeshare member who has the right to a vacation property in Chicago for a week in June may be able to exchange it with another member who has the right to a property in Florida for a week in December. For this \$9.4 billion industry (Resort Condominiums International 2003), we demonstrate theoretically that the two exchange systems currently in use (deposit-first and request-first mechanisms) can cause efficiency loss. We propose an alternative exchange mechanism, the top trading cycles chains

and spacebank (TTCCS) mechanism, and show that it can increase the efficiency of the timeshare exchange market because it is Pareto efficient, individually rational, and strategyproof. We then test the three mechanisms in laboratory experiments where we simulate exchange markets with networked “timeshare members.” The results of the experiments strongly support our theory.

To expand a little on the timeshare industry, we note that virtually all timeshares are resort or vacation properties. Currently in this industry, there are over 5,400 resorts in more than 100 countries, and ownership of timeshares increased at an average annual rate of about 12% from 1990 to 2003. Occupancy periods of timeshare resort accommodations are commonly called weeks because they are almost always for one week. As of January 1, 2003, 6.7 million households worldwide owned 10.7 million timeshare weeks (Resort Condominiums International 2003). It is believed that “more than 80% of [timeshare] purchasers buy with the intention of using the exchange option” (Ragatz Associates 1999), whereby a timeshare member can exchange her week for someone else’s at a different resort in a different time of the year, although her *home week* (the week she bought) is fixed at one place in one specific time period. These medium-free exchanges are preferred to payment or bidding because a core appeal of timeshares is the aspect of one-time up-front payment. If people had to bid each time, it would become similar to booking hotels and would lose its attraction.

Most exchanges are made through vacation clubs and exchange companies. The two largest timeshare exchange companies, Resort Condominiums International and Interval International, confirm over 95% of all the timeshare exchanges throughout the world.¹ In this paper, we explore the following questions.

1. Can the existing timeshare exchange mechanisms realize economically efficient allocation outcomes?
2. If the answer to the first question is negative, then is there a more efficient way to allocate the resources (weeks) to members of the exchange system?
3. When implemented among boundedly rational agents (as timeshare members are in the real world), will the exchange mechanisms we study generate allocation outcomes as suggested by the theory?

To formulate the timeshare exchange problem and analyze it theoretically, we build on literature in economics on mechanism design and, more specifically, on matching mechanisms where resources are matched to agents based on their preferences (e.g., Abdulkadiroglu and Sonmez 1999, Sonmez and Unver 2005). Our problem falls into the class of “one-sided matching problems” (where agents have preferences over resources, but resources do not have preferences over agents) as opposed to “two-sided matching problems” such as hospital-intern matching, where interns have preferences over hospitals and hospitals also have preferences over interns. Some of these matching mechanisms have been put into practice. For example, the hospital-intern matching mechanism (Roth and Peranson 1997) was adopted in 1997 by the National Resident Matching Program, and 14 transplant centers in New England plan to implement the economists’ design for kidney exchanges (Roth et al. 2004).

We focus on an industry not studied earlier, namely the timeshare industry, with the hope of improving its operation. Theoretically, we contribute by identifying this industry as one that can benefit from improved mechanism design, formally characterizing the currently used mechanisms and showing their weaknesses, proposing an alternative mechanism, and demonstrating experimentally that it can improve efficiency. The proposed algorithm builds on an earlier one, as is typical in mechanism-design research. Specifically, the modification extends theoretical work on mechanism design to cases where supply of resources is dynamic and resources can

be stored. Laboratory experiments are conducted to test the theory. Studies adopting similar approaches include Ho and Weigelt (1996), Srivastava et al. (2000), Amaldoss et al. (2000), Haruvy et al. (2001), Amaldoss and Jain (2002), Zwick et al. (2003), etc.

While the relevance of mechanism-design work to business systems is obvious, most work on the topic has been done within economics. Thus, another goal of this paper is to encourage the use of mechanism-design work within management science.

2. Existing Timeshare Exchange Mechanisms

At present, the most commonly used timeshare exchange mechanisms are the deposit-first (DF) and request-first (RF) mechanisms. In both, members try to exchange their home week for a more preferred week (its time and/or location is different from that of their home week). They look for a preferred week in a depository of weeks that we call the *spacebank*. Thus, the spacebank is an exchange company’s inventory of weeks. Members who want to make exchanges can “deposit” their home week into the spacebank and “withdraw” a different week from it. Also, unsold weeks are often put into the spacebank by timeshare developers in the hope that members may exchange their home week for these weeks, enjoy them, and buy them in the future.

Deposit-First Mechanism

Under the DF system, in order for a member to see what is available in the spacebank (for exchange with her home week), *she has to deposit her week with the exchange company first*. Thus, a member can withdraw a week under the DF mechanism only if (1) she has already deposited her home week, (2) a week that she wants is available in the spacebank, and (3) no one else requests the week or has a higher priority for getting it (priorities are exogenously given, i.e., priorities are not determined by the exchange mechanisms per se; this will be discussed in greater detail later). It is possible that after a member deposits her week, her preferred weeks do not ever become available in the spacebank. Thus, a member may end up with a worse week than her home week, suggesting that in the DF mechanism members face a risk of being left with a week worse than their own.

Request-First Mechanism

In the RF mechanism, members can deposit their home week *after* they have found a preferred week in the spacebank and withdrawn it successfully (i.e., no one else requests it or has a higher priority for getting it). Upon confirmation that a member has been allocated this preferred week, the right to the use of her home week is immediately assigned to the exchange company.

¹ As of 2002, Resort Condominiums International had about 3,700 affiliated resorts and over 3.5 million member families; Interval International was made up of 1,900 affiliated resorts and nearly 1.5 million member families. Confirmed exchanges in 2001 totaled 2,052,760 and 669,000 for Resort Condominiums International and Interval International, respectively, adding up to about 98% of all exchanges made throughout the world in that year.

Searching Methods

Under both DF and RF mechanisms, members have two options to search for a preferred week. They can proactively search online through a website or phone into a call center to find out what is available in the spacebank. (In the case of DF, they can only do this if they have deposited their home week.) If a member finds this method too exacting, then the second option is to get on a “waiting queue” whereby she tells the exchange company what resorts and times she wants, and then waits to be informed if these become available. By doing so, the member would be taking a “gamble” of exchanging her home week for the right to get on the waiting queue. (She has to make a deposit *without knowing* what week she can get until later, and may end up with a week less preferred than her home week.)²

Potential Problems of the Existing Mechanisms: Inefficiency

Ideally, a one-sided matching mechanism should be *Pareto efficient*, *individually rational*, and *strategyproof* (Abdulkadiroglu and Sonmez 1999, Krishna and Wang forthcoming). In the timeshare exchange setting, *Pareto efficiency* means that the final allocation of weeks is such that it would be impossible to increase one person’s utility from her allocated week without decreasing someone else’s utility from their allocated week. A mechanism is *individually rational* if it assures every member an alternative that is at least as good as her home week. When members have the option not to participate (as in timeshare exchanges), lack of individual rationality means there is risk involved in participation, which results in low participation rate in exchanges (i.e., people opting not to participate). This causes efficiency loss for the system because fewer exchanges are made. Thirdly, a mechanism is *strategyproof* (or *dominant strategy incentive compatible*) if no member can ever benefit by misrepresenting her preferences. A mechanism needs to be strategyproof in order to be efficient, because all the allocations are made based on members’ “revealed preferences,” while their true preferences are private information.

The example below shows why the DF and RF mechanisms tend to result in lower efficiency.

EXAMPLE 1, PART A. Three members, Amy, Bob, and Cindy, own timeshare weeks 1, 2, and 3, respectively. The utility of each week for each member is shown in the table below. A larger number represents higher utility. For example, for Amy, Week 2 is the most preferred (with utility 10), followed by Week 1 (with utility 8), and then Week 3 (with utility 5).

	A (Amy)	B (Bob)	C (Cindy)
w1 (Week 1)	<u>8</u>	10	8
w2 (Week 2)	10	<u>8</u>	9
w3 (Week 3)	5	5	<u>5</u>

A member has two options: to keep her original week (i.e., “opt out”), or to give it up and enter the exchange system (i.e., “opt in”). Thus, *under the DF mechanism*, should all the members enter the exchange system? More importantly—if some of the members do not enter the exchange system, is the outcome going to be Pareto efficient?

First, assume all of the members decide to opt in. Also assume that whenever two, or more than two, members compete for the same week, each has the same probability of getting that week. Thus, Amy has to compete with Cindy for w2 (because for both members w2 is the most-preferred week). On the other hand, no one competes with Bob, so he gets w1 for sure. If Amy is an expected-utility maximizer, her expected utility for entering the exchange system will be: $0.5 * 10 + 0.5 * 5 = 7.5$. (Because Amy has 0.5 probability of getting w2, she also has 0.5 probability of getting w3, which has a utility of 5 for her.) Therefore, for Amy, the optimal strategy is to “opt out” and keep Week 1, which gives a utility of 8 for sure. Up to this point, we already know that a rational Amy will not enter the exchange system.


If we further consider the “unraveling” of the market, then given the rational strategic choice of Amy, Bob will not want to enter the system either (because Bob prefers w1, which will not be deposited to the spacebank by Amy). It is obvious that given the rationality of the members, our original assumption that everyone enters the system is incorrect. Obviously, no exchange happens and the original endowment carries to the end. The final outcome {Amy, w1; Bob, w2; Cindy, w3} is Pareto dominated by: {Amy, w2; Bob, w1; Cindy, w3}.³

Example 1 makes it clear that members with more desirable weeks may hesitate to enter the DF exchange system because they may end up with a week worse than their home week. In other words, the DF mechanism is not individually rational. As a result, lower participation rate leads to fewer exchanges, although some members could benefit from exchanging their weeks (e.g., Amy and Bob could have made exchanges between themselves and both would have benefited). Therefore, the final allocations of the DF mechanism tend to be inefficient.

³ In this example, we assume that all members have accurate beliefs about others’ endowments and preferences. In the real world, members will have some beliefs about the value of their week versus others’ and we will get the same result—members who believe they have a relatively high-value week will not deposit. In the theory developed later, we do not assume accurate beliefs.

² In the RF mechanism, under some plans if a member decides to join the waiting queue, she has to give up her home week only when a preferred week becomes available and is allocated to her.

By contrast, under the RF mechanism, anyone who wishes to obtain a better week will not hesitate to participate. This is because the deposit of the home week is made *after* the confirmation of an exchange. Because a member is guaranteed a timeshare week that is at least as good as her home week, the RF mechanism is individually rational. Theoretically, we would expect a participation rate of 100%. However, the RF mechanism also has its own deficiency. Because a member can hold her week until she finds a better week and confirms the exchange, her home week is not available to other members for exchange during her entire searching period. This can substantially decrease the number of weeks stored in the spacebank. More importantly, mutually beneficial “trading cycles” may remain undetected in such a system. A *trading cycle* is a cycle of member-week pairs where every member prefers another member’s week so that a cycle is formed. Example 1, discussed earlier, shows a simple situation where such a trading cycle exists:



(w2 is Amy’s top choice; w2 belongs to Bob; w1 is Bob’s top choice; w1 belongs to Amy.)

EXAMPLE 1, PART B. While exchange between Amy and Bob is clearly mutually beneficial, it may go undetected *under the RF mechanism* because they both wait for each other to deposit their weeks. As a result, Amy and Bob will not be able to switch their weeks. Cindy will also stay with her home week because no other week is deposited. Thus, the original allocation will carry to the end.

As the example shows, the outcome of the RF mechanism is likely to be inefficient too. Only when all such trading cycles are identified by the system can exchanges be conducted and the efficiency of the system improved. Thus, under both the DF and the RF mechanisms, there is a potential loss of efficiency. Please note that Example 1 captures the essence, but simplifies the operation, of the two exchange mechanisms. Most importantly, for simplicity, a continuity element is not considered yet. We emphasize that the essence of the example still holds even when continuity is explicitly considered.

Given the identified deficiencies in the existing timeshare exchange systems, the question now is: How does one design a Pareto-efficient mechanism that satisfies the greatest number of members with the best possible exchanges? If such a mechanism does exist, it can boost customer satisfaction without requiring extra resources (i.e., without adding more timeshare weeks into the spacebank). We propose a more comprehensive timeshare exchange mechanism, which

1. reduces the risk involved in timeshare exchanges and leads to a higher participation rate compared to the DF mechanism,

2. detects all trading cycles in the system and conducts exchanges accordingly, and

3. results in higher efficiency compared to the RF and DF mechanisms.

More specifically, as we will show, the proposed mechanism is Pareto efficient, individually rational, and strategyproof. In the following section we briefly review extant research, which we build on.

3. Literature Review

The timeshare exchange problem is categorized as a one-sided matching problem. Research on one-sided matching problems began with the study of the housing market where agents (each having one house) trade houses with each other (Shapley and Scarf 1974). The two most relevant problems for our work are the house allocation problem (Abdulkadiroglu and Sonmez 1999, or A&S 1999) and the kidney exchange problem (Roth et al. 2004, or RSU 2004).

A house allocation is an assignment of houses to agents such that each agent (either a current tenant or a new applicant without a house) is assigned a distinct house. A&S 1999 show that real-life mechanisms are inefficient because of their failure to guarantee each existing tenant a house at least as good as the one she currently occupies. They propose a theoretically superior mechanism to solve the house allocation problem, which leads to full participation, sincere revelation of preferences, and Pareto-efficient final outcomes. If in timeshare exchanges *all members got their final allocation simultaneously*, then the timeshare exchange problem can be seen as a special case of the house allocation problem where all the potential applicants are existing tenants. However, a key difference between the timeshare exchange problem and the house allocation problem is that the latter is one-shot, whereas the former entails a continuity element—timeshare members can choose to get on a queue and wait for future opportunities rather than obtaining instant allocations. Also, timeshare members enter and leave the system in a *stochastic* manner.

The continuity element makes the timeshare exchange problem more similar to the kidney exchange problem, which has been investigated by Ross and Woodle (2000), Zenios et al. (2001), and Roth et al. (2004), among many others. In the kidney exchange scenario, patients either get kidneys from living donors or wait in a queue for kidneys from future cadavers, if and when they become available. Because of immunological incompatibility between patients and their own donors, exchanges involving several donor-patient pairs are often made. In RSU 2004, the top trading cycles and chains (TTCC) mechanism is

Table 1 Comparison of Discrete Resource Allocation Problems

	Three discrete resource allocation problems		
	Timeshare exchange	Kidney exchange	House allocation
Nature of the system (continuous or one-shot)	Continuous (1) members enter and leave the system stochastically; (2) members on the queue wait for future opportunities	Continuous (1) patients enter and leave the system stochastically; (2) patients on the queue wait for future opportunities	One-shot allocation: every applicant gets a house when the algorithm ends
Agents holding an original “endowment”	Members with weeks at hand	Donor-patient pairs	Current tenant-house pairs
Agents without any “endowment”	Members without weeks (the weeks have been deposited)	Patients on the queue waiting for kidneys from cadavers	New applicants
Resource other than the endowment held by agents	Spacebank composed of: (1) Weeks deposited by members (2) Weeks deposited by developers	Kidneys from cadavers {The kidney is highly perishable, so there is no “spacebank”}	Vacant houses

proposed to make kidney exchanges more efficient and incentive compatible. The timeshare exchange problem is similar to the kidney exchange problem if we think of the members who still hold their home weeks as the counterpart of donor-patient pairs, and the members who have deposited their home weeks but have not withdrawn desirable weeks as the counterpart of patients on the waiting queue. Another important similarity between these two systems is the continuity element. Because of these important common features of the two problems, we will propose a mechanism that is mainly based on the TTCC mechanism to solve the timeshare exchange problem.

There are also a few differences between the kidney exchange and timeshare exchange systems. First, because cadaver kidneys are highly perishable, in the kidney exchange problem there is no counterpart of a spacebank. The availability of the spacebank necessitates the modification of the TTCC mechanism. Second, in the kidney exchange problem most patients on the waiting queue are without paired donors, whereas in the timeshare exchange problem every member has her home week but chooses to relinquish it to join the waiting queue. In Table 1 we compare the similarities and differences between the timeshare exchange problem, the kidney exchange problem, and the house allocation problem.

4. Solution to the Timeshare Exchange Problem: TTCCS Mechanism

We build on the TTCC mechanism (RSU 2004) to propose the top trading cycles chains and spacebank (TTCCS) mechanism. Similar to the solution to the kidney exchange problem proposed by RSU 2004, our approach is to consider an instance of the dynamic problem and solve that instance (which is static) with a static mechanism. The continuity element is captured by allowing members to get on the waiting

queue, and by running the mechanism repeatedly over time.

Timeshare Exchange Model

We denote a member-week pair by (m_i, w_i) , consisting of a member m_i and her home week w_i . For a system with n member-week pairs at a specific point of time, let W denote the set of weeks from all member-week pairs (i.e., $W = \{w_1, w_2, \dots, w_n\}$). Let (S, W_s) denote the spacebank and the set of weeks in the spacebank, where $W_s = \{w_{s1}, w_{s2}, \dots, w_{sm}\}$. Let q represent the option of relinquishing the week a member holds in order to get on the waiting queue. Each member has preferences over all the potentially available weeks $W \cup W_s$ and the waiting-queue option q . For simplicity, we assume that all the preferences are strict, meaning there are no equally preferred alternatives. Members are expected to reveal their preferences by ordering various options from the most preferred to the least preferred.⁴ Let P_i denote the strict preferences of m_i over $W \cup W_s \cup \{q\}$. For our purpose, *the relevant part of P_i is the ranking up to week w_i or q , whichever ranks higher*. If member m_i ranks w_i , her home week, at the top of her preferences, then she will not participate in the exchange system. If member m_i ranks w_i higher than q , then she will not exchange her original week for a position on the waiting queue. If she ranks q higher than w_i , then if all weeks above q are unavailable, she will relinquish w_i to join the queue.⁵ In an exchange system with a large number of weeks, doing an exhaustive ranking even up to w_i or q , whichever ranks higher, may be quite cumbersome. However, what matters is the ranking only for weeks that members know about and prefer to their own.

⁴ We show later that it is in the best interest of the members to truthfully reveal their preferences, i.e., TTCCS is strategyproof.

⁵ We prove later that the TTCCS mechanism guarantees a week at least as good as one’s own, i.e., it is individually rational.

Formally, at a specific point in time, a timeshare exchange problem consists of

1. a set of member-week pairs $\{(m_1, w_1), \dots, (m_n, w_n)\}$,
2. a set of weeks from the member-week pairs $W = \{w_1, \dots, w_n\}$,
3. a set of weeks from the spacebank $S: W_s = \{w_{s1}, w_{s2}, \dots, w_{sm}\}$, and
4. a strict preference relation P_i over $W \cup W_s \cup \{q\}$ for each member m_i (although only a small part of the preference relation is relevant for our purpose).

The outcome of a timeshare exchange problem is a matching of weeks and waiting-queue option and the members such that (i) each member m_i is assigned either a week from $W \cup W_s$ or the waiting-queue option q , and (ii) no single week is assigned to more than one member, whereas the waiting-queue option q may be assigned to multiple members.

Cycles and Chains

Before solving the timeshare exchange problem with the TTCCS mechanism, we provide formal definitions of a few key concepts: cycles, *S*-chains, and *q*-chains. The core of the mechanism is an algorithm consisting of several rounds. In each round

- each member m_i points towards either a week in $W \cup W_s$ or q , and
- each week w_i points to m_i , the member who holds it.

Given the pointing of the members and the weeks, a *cycle* is defined as an ordered list of member-week pairs $(w'_1, m'_1, w'_2, m'_2, \dots, w'_k, m'_k)$ such that

- week w'_1 points to member m'_1 ,
- member m'_1 points to week w'_2 ,
- ⋮
- ⋮
- week w'_k points to member m'_k , and
- member m'_k points to week w'_1 .

With TTCCS, whenever a cycle is formed, we should carry out exchanges such that the pointer (a week) is assigned to the pointee (a member). Thus, in Figure 1, we should assign w'_2 to m'_1 , w'_3 to m'_2 , and w'_1 to m'_3 . It is worth noting that the exchanges related with cycles do not involve the waiting-queue option, and all members in the cycles can get their week allocations instantly and simultaneously.

We define two types of *chains*. A *q-chain* is an ordered queue of weeks and members $(w'_1, m'_1, w'_2, m'_2, \dots, w'_k, m'_k)$ starting with a week and ending with a member, as illustrated in Figure 2. The second type of chain involves the spacebank *S*. An *S-chain* is an ordered list of member-week pairs $(w'_1, m'_1, w'_2, m'_2, \dots, w'_k, m'_k)$ starting with a week and ending with a member, as illustrated in Figure 3.

We use *chains* as a general term to refer to both the *S*-chains and the *q*-chains, although they are conceptually quite different. With any chain $(w'_1, m'_1,$

Figure 1 Example of a Cycle (Formed by Three Member-Week Pairs)




Figure 2 Example of a q-Chain (Formed by Two Member-Week Pairs and q)





Figure 3 Example of an S-Chain (Formed by Two Member-Week Pairs and a Week in Spacebank S)



$w'_2, m'_2, \dots, w'_k, m'_k)$, we will refer to the pair (w'_k, m'_k) as the *head* and the pair (w'_1, m'_1) as the *tail*. With TTCCS, exchanges associated with chains should be carried out such that the member at the head of the chain either gets an instant allocation from the spacebank (in the case of *S*-chains), or gets a position on the waiting queue (in the case of *q*-chains); the other members in the chain get the instant week allocations that they point to.

Active and Passive Members

TTCCS mechanism involves multiple steps. At a certain step, *active members* refer to those who have not yet been assigned weeks. *Passive members* refer to those who have already been assigned weeks, but are not yet removed from the system. This will become clearer when we provide examples later.

Proposed Solution: TTCCS Mechanism

With the key concepts defined, we now discuss our solution to the timeshare exchange problem, the TTCCS mechanism. At a fixed time, we take as given the set of member-week pairs and the set of weeks in the spacebank. With this fixed set of member-week pairs, all the members in the pairs are prioritized into a single list. Priority can be determined by the value of home weeks so that members with more valuable weeks obtain higher priority; it can be random (i.e., equal), etc. TTCCS can accommodate any priority system that the timeshare company chooses to use. Note that higher priority does not automatically imply that a member has a better shot at getting a desired week. The ability to exchange one's week for a more desired week under TTCCS depends both on priority and on how desirable one's home week is for other members (so that trading cycles can form).

We similarly assume that the operation of the waiting queue is given—this mainly refers to how the allotment is (dynamically) made to the members who choose to join the queue. Again, this can be based on

the value of the home week, be equal, etc. The system for prioritizing members and for operating the waiting queue can be the same or be different.

Then, the TTCCS mechanism determines the exchanges as follows.⁶

1. In each round of the algorithm

- based on her preferences P_i , each remaining active member m_i points to the best available week in $W \cup W_s$ or to the waiting-queue option q , whichever is preferred;
- each remaining passive member continues to point to her assignment;
- each remaining week in W points to its paired member.

2. There is either a cycle, or a chain (S -chain or q -chain), or both.⁷ At this step (Step 2), we make exchanges associated with cycles.

(i) Proceed to Step 3 if there are no cycles. Otherwise, locate one cycle⁸ and carry out the corresponding exchanges. Remove all the members and the weeks in the cycle from the system.

(ii) If there is no active member in the system, go to Step 4. Otherwise, go back to Step 1.

3. Each remaining pair initiates a chain.

(i) Select the chain starting with the highest-priority member among the remaining active members. Carry out exchanges for the members in the selected chain. The assignment is final for these members. The selected chain remains in the system but each member in it is passive thereafter.

(ii) If there is no active member in the system, go to Step 4. Otherwise, go back to Step 1.

4. Remove all the chains (if there are any) and put the weeks at the tails of the chains into the spacebank. The algorithm is finished.

Example 2 below illustrates how the mechanism works.

EXAMPLE 2. Consider a timeshare exchange problem consisting of 12 member-week pairs $(m_1, w_1), \dots, (m_{12}, w_{12})$. There are four extra weeks in the spacebank: w_{13}, w_{14}, w_{15} , and w_{16} . The members are prioritized from m_1 to m_{12} , with m_1 having the highest priority and m_{12} , the lowest. Preferences of the members over the weeks and the waiting-queue option are given below. (Note that for each member this is not, and need not be, an exhaustive ordering of all the weeks.) The dark arrows show what was already allocated based on the cycle identified in Step 2 and the chain selected in Step 3(i).


⁶ The algorithm described here uses a specific chain selection rule. Other rules may also be used as in RSU 2004.

⁷ Please refer to RSU 2004 for proof of this statement.

⁸ Because cycles never intersect, when multiple cycles coexist, removing any one of them does not influence the others. Therefore, the order in removing cycles does not have any effect on the final allocation of the mechanism.


- | | |
|---|---|
| $m_1: w_{11}, w_2, w_1$ | $m_7: w_6, w_4, w_2, w_9, w_7$ |
| $m_2: w_1, w_{16}, w_4, q$ | $m_8: w_7, w_{15}, q$ |
| $m_3: w_{15}, w_9, w_3$ | $m_9: w_8, w_{14}, w_1, w_5, q$ |
| $m_4: w_3, w_5, w_{16}, w_4$ | $m_{10}: w_{14}, w_{11}, q$ |
| $m_5: w_1, w_3, w_{14}, w_4, w_5$ | $m_{11}: w_{15}, w_1, w_3, w_6, w_{11}$ |
| $m_6: w_3, w_8, w_2, w_1, w_5, w_{13}, q$ | $m_{12}: w_1, w_{11}, w_5, w_{12}$ |

Round 1




There is no cycle in this round. Pick the S -chain starting with m_1 : assign w_{11} to m_1 ; assign w_{15} to m_{11} , and keep the chain in the system. m_1 and m_{11} are not active anymore.

Round 2




Now m_3 points to her second choice, w_9 , and a cycle forms. Make exchanges accordingly, then remove (m_3, w_3) , (m_6, w_6) , (m_7, w_7) , (m_8, w_8) , and (m_9, w_9) from the system.

Round 3




Now m_4 points to w_5 . There is no cycle in this round. The S-chain starting with m_2 is picked because now m_2 has the highest priority among all the active members. Assign w_1 to m_2 , and keep the chain in the system.

Round 4




Now m_5 point to w_{14} and m_{12} points to w_5 . Again, there is no cycle. Pick the S-chain starting with m_4 . Assign w_5 to m_4 , assign w_{14} to m_5 , and keep the chain in the system.

Round 5



Now m_{10} points to q and m_{12} points to her home week, w_{12} . A cycle forms. Assign w_{12} to m_{12} and remove the cycle.

Round 6



Now the only active member is m_{10} . Put m_{10} on the waiting queue. Put all the weeks at the tails of the chains (w_2, w_4, w_{10}) into the spacebank. We are done.

The final allocation is:

$$\begin{pmatrix} m_1 & m_2 & m_3 & m_4 & m_5 & m_6 & m_7 & m_8 & m_9 & m_{10} & m_{11} & m_{12} \\ w_{11} & w_1 & w_9 & w_5 & w_{14} & w_3 & w_6 & w_7 & w_8 & q & w_{15} & w_{12} \end{pmatrix}.$$

Theoretical Properties of the TTCCS Mechanism

We now consider whether the TTCCS mechanism meets the three ideal mechanism properties we discussed earlier. First of all, the proposed mechanism is individually rational. The final assignment to each member can be either an instant week allocation or a position on the waiting queue (only if the member prefers the waiting-queue option to her original week). An instant week allocation is at least as good as one’s home week because a week always points to its owner as long as she is still in the system. Thus, it is guaranteed that the owner can form a cycle with her home week whenever she fails to get a better option, so theoretically the participation rate should be 100%. Secondly, the proposed mechanism is Pareto efficient. The idea is that in Round 1, the members who get allocations receive their top choices among all the available options. In Round 2, the members who get allocations receive their top choices among all the remaining options. Proceeding in a similar manner, it is obvious that nobody’s allocation can be improved without hurting anyone who got an allocation in an earlier round. Thus, the mechanism is Pareto efficient. Lastly, RSU 2004 shows that TTCC is strategyproof for certain variants of their algorithm, including TTCC with chain selection rule e . The TTCCS mechanism can be shown to be equivalent to that variant and so is also strategyproof. Because our proposed mechanism is Pareto efficient, individually rational, and strategyproof, theoretically it is superior to the existing mechanisms.

We now test whether the theoretically superior TTCCS mechanism remains superior to DF and RF in practice, when all three mechanisms are implemented among boundedly rational agents. Empirical test of the theory is necessary because it has been shown previously that for some mechanisms there exist persistent “discrepancies between data and theory” (Kagel and Levin 1993).

5. Experimental Studies

To simulate the *timeshare exchange market*, we designed special computer software that allows us to set up the market, monitor individual behavior, carry out exchanges, and collect experimental data.⁹ We simulated four different market environments (four conditions) that differ in the heterogeneity of preferences,

⁹Details about the software are available from the authors on request.

the way priority is determined, and market size (discussed in detail later). Within each condition, there were three sessions, each using one of the three mechanisms (thus, we had 12 experimental sessions in total). We compared results from the three sessions to see if our theory was supported. The four different conditions tested the robustness of our results.

Design

We had 120 participants in total, 36 each for three conditions (3 mechanisms \times 12 participants per mechanism) and 12 for a fourth smaller-market condition (3 mechanisms \times 4 participants per mechanism). The participants were recruited from a paid subject pool in a large university and took part in seven sets of exchanges (in the same condition), with the first set for practice purposes. Each set of exchanges simulated one *market*, where participants started with their home weeks and ended with the final allocation. As in other studies of matching mechanisms (e.g., Kagel and Roth 2000), we ran multiple markets in each session to generate sufficient data points for statistical tests. All participants got a \$3 fixed participation fee, plus the performance-related monetary compensation. A participant's performance-related compensation was determined by *the average number of points she got in the six markets* (number of dollars earned equaled the average points earned). Further details on the experimental procedure are given after we describe the four market environments.

Base Condition (More Heterogeneous Preference and Value-Linked Priority, or MH-VL Condition)

For each of the three sessions (mechanisms) in the base condition, there were 12 members and 15 timeshare weeks (12 home weeks of the members and 3 weeks in the spacebank). In the *DF session*, each market consisted of five consecutive *periods*—this incorporates the dynamic interaction by allowing members to enter the market in Period 1 with pre-assigned home weeks and leave the market at the end of Period 5 with the weeks they obtain through exchanges. Within each period, a member could choose either to make a deposit and then withdraw another week, or to opt out, in which case she would keep her own week and would not be allowed to see what was in the spacebank. If several members tried to withdraw the same week in the same period, each member's chance of getting the week was proportional to the average value (across members) of her home week (please refer to Table 2a for the values of weeks). In specific, the probability of member m_i getting her preferred week was given by $\{\text{Average Value of } w_i / \text{Sum of Average Values of All Competitors' Home Weeks (including } w_i)\}$. Therefore, *the more valuable one's home week was, the higher priority she got*

so that she would have a better chance of getting her preferred week, everything else being equal. We call a priority list determined by this method a *value-linked priority* thereafter—this method is also consistent with the real world, where priorities can be related to property value, and hence should be consistent with average valuation across members.

Members who failed to withdraw a week were eligible to withdraw other weeks in the next period. Every member could choose to make multiple exchanges by depositing and withdrawing repeatedly, just like in the real world. The outcome in the DF market was determined by the value of the week that was held by the member *at the end of Period 5*, when "the exchange market ended." The waiting queue was a special alternative, operationalized as a gamble to reflect the uncertainty involved in waiting. It gave a member a fair draw, i.e., an equal chance, among the 15 weeks. This option was only available in Period 1 under DF and RF and could be ranked in the rank-order list under TTCCS. When a member chose to enter the queue, her final allocation was simply determined by the realization of the gamble, i.e., the member would not go through the five periods step by step, and instead just waited to see what she was allocated.¹⁰ The waiting queue was operated similarly across the three mechanisms so that any difference in our results could be clearly attributed to the mechanism being used and not anything else.

In the *RF session*, everything was the same as the DF session except that the members were allowed to hold their weeks while checking the spacebank. Once a withdrawal was successful, the home week was deposited automatically. Note that participation in RF meant that members had tried to make exchanges by at least checking the spacebank.

In the *TTCCS session*, if a member chose to enter the exchange market, then she was required to submit a "preference list" where she was supposed to rank order all the weeks and the gamble option from "the most preferred" to "the least preferred." (While we asked for this exhaustive ranking in the experiment, remember that all that is needed to implement TTCCS is the preferences up to the home week or the queue, whichever is ranked higher.) Once all

¹⁰ Thus, the realization of this queue did not depend on others' decisions, and the expected value of entering the queue was fixed (i.e., the size of the queue was totally irrelevant). This makes the three mechanisms strategically comparable. Furthermore, in reality the value of the waiting queue may depend on expected future participation, and thus may differ across mechanisms, e.g., because more members are expected to participate under TTCCS, the queue option may be more valuable under TTCCS. In our experiment, the value of the queue was the same across the three mechanisms, which actually biases the results against TTCCS. However, this leads to a stronger test of TTCCS. We thank the reviewers for suggesting this implementation of the queue.

members had made decisions regarding whether they wanted to opt out (and keep their home week) or to enter the market (and submit their preference list), the mechanism determined the final allocations. In this value-linked priority condition, in the priority list to implement TTCCS, the owner of the most popular week (the week that has the highest average value) had the highest priority; the owner of the second popular week had the second-highest priority, etc.

Please note that under DF and RF mechanisms, typically the members need to interact repeatedly in order to make exchanges and get their preferred weeks. Therefore, in the experiment there were multiple “periods” within each market. Under TTCCS, the agents only need to submit their rank-order lists (only once for each market), and the system automatically works out the solution (going through multiple rounds of exchanges if needed, as shown earlier in Example 2).

Values of Weeks. In all three mechanisms (sessions) in the base condition, each of the 12 members started with a home week. In addition, three extra weeks (w13, w14, and w15) were initially in the spacebank and had no prior owner. The *values* of the weeks varied for each member, and the members had different preferences. We designed the parameters very carefully to ensure that preferences were quite heterogeneous in this condition. The average correlation of members’ preferences is 0.04. Table 2a presents all the parameters we used in the base condition.

Numbers in brackets indicate the weeks that were originally owned by the members when they entered the market. For example, the home week of Member 1 (i.e., M1) was w1, which was worth 12 points to her. Each member’s preferences were private information. For example, M1 only saw the first row of parameters, which shows how many *points* each week was worth to her. We calculated the baseline for efficiency comparison: The average Pareto-optimal total earning equals 149.5.¹¹ The initial allocation before any exchanges realizes 60.2% of the optimal result, and we were interested in how much each mechanism could improve on this. Because the waiting queue has a random allocation, the value of the waiting-queue option

for a member should (normatively) be the average points across all weeks for that member. In our experiment this value was equal for all members.

Less Heterogeneous Preference and Value-Linked Priority, or LH-VL Condition

To test the robustness of our results with respect to different preference structures, we designed the LH-VL condition where a different value chart was used (see Table 2b). The average correlation of preferences here is 0.264 (versus 0.04 in the base condition).

The baseline for efficiency comparison is 149.0, given the value chart in Table 2b. The initial allocation realizes 60.4% of the optimal result. Therefore, for this condition everything was kept the same as (or extremely similar to) the base condition, except for a less heterogeneous preference structure.

More Heterogeneous Preference and Random Priority, or MH-R Condition

Although our model takes the priority list as given, it is interesting to test the theory with different prioritizing methods. Therefore, in the MH-R condition we kept everything the same as in the base condition (i.e., same value chart), except that we used a random priority for all the members. That is, under DF and RF, when several members competed for the same week, each was given the same probability of getting that week. Similarly, TTCCS proceeded with a randomly determined priority list.¹²

Small-Sized Market, or MH-VL-S Condition

To test the robustness of our result with respect to market size, we designed the more heterogeneous preference, value-linked priority, and small market-size condition (MH-VL-S), where there were four members and five weeks. The waiting queue gave a member a fair draw (i.e., an equal chance) from the five weeks. A value chart was carefully constructed so that the key features were extremely similar to those of the base condition (see Table 2c). The average Pareto-optimal total earning of the group is 30, and the initial

¹¹ In this complex system, there are multiple Pareto-efficient outcomes. To calculate the average Pareto-optimal total earning as the “baseline” for efficiency comparison, we obtain all the permutations of the 12 members (which gives us 12! different orderings); with each ordering, we let the member at the top pick her best choice among the available options, and so on. We obtain the optimal total earnings for each of these 12!, or 479,001,600, orderings and then calculate the average, which equals 149.5 with a standard deviation of 5.2. We emphasize that we use the words “Pareto efficient” or “Pareto optimal” in their basic economic sense. We are *not* looking for the outcome that gives us the maximum total earning for the whole group.

¹² As discussed earlier, another nicety of TTCCS is that a member’s chance of getting her preferred week depends both on her priority and on the value of her home week because trading cycles are more likely to form for members who have desirable home weeks, i.e., these members are more likely to be in a position where they want other members’ weeks, and other members want their week, so that a trading cycle forms. Priority affects member-week pairs that do not form trading cycles and thus get sorted into *S*-chains and *q*-chains (see Example 2; in the formation of *S*-chains and *q*-chains, priority determines who gets to pick a week first). This is not the case with DF and RF, where a member’s chance of getting a week depends solely on priority (when multiple members want the same week, priority determines who is more likely to get it). So, from fairness concerns, TTCCS is better than DF and RF if priority is random (i.e., in TTCCS, the value of the home week will still matter for trade, whereas for DF and RF it will not).

Table 2a Value Chart for MH-VL and MH-R Conditions

	w1	w2	w3	w4	w5	w6	w7	w8	w9	w10	w11	w12	w13	w14	w15
M1	[12]	15	5	3	1	3	5	1	8	1	9	3	6	8	6
M2	1	[9]	15	12	3	5	1	5	3	3	6	8	1	6	8
M3	15	5	[3]	12	9	6	1	3	1	8	5	6	8	1	3
M4	5	3	8	[12]	15	1	1	6	3	1	6	5	3	9	8
M5	3	1	3	9	[8]	15	12	3	1	5	5	6	1	8	6
M6	6	1	3	15	3	[5]	12	9	1	8	5	6	8	3	1
M7	5	3	1	1	3	1	[12]	15	6	5	3	9	6	8	8
M8	12	8	3	1	8	3	5	[9]	15	6	5	6	1	1	3
M9	12	9	5	6	3	1	15	1	[3]	8	5	6	8	1	3
M10	15	3	8	12	5	9	6	3	1	[6]	1	1	8	5	3
M11	6	5	1	15	3	8	12	3	9	6	[5]	1	8	3	1
M12	12	6	9	3	5	1	15	1	8	3	5	[6]	8	1	3

Table 2b Value Chart for LH-VL Condition

	w1	w2	w3	w4	w5	w6	w7	w8	w9	w10	w11	w12	w13	w14	w15
M1	[12]	5	1	8	3	1	8	15	6	5	1	9	6	3	3
M2	8	[9]	15	6	5	1	12	3	1	6	3	8	1	5	3
M3	6	3	[3]	15	9	5	12	3	1	8	5	6	8	1	1
M4	8	15	5	[12]	3	3	8	1	5	1	9	6	6	1	3
M5	12	1	3	6	[8]	15	9	1	1	5	3	6	5	8	3
M6	12	6	3	8	3	[5]	15	9	1	8	5	1	6	3	1
M7	9	3	5	8	15	1	[12]	6	3	1	6	5	3	8	1
M8	5	8	3	12	8	3	1	[9]	15	6	5	6	1	1	3
M9	15	9	5	12	3	1	8	1	[3]	8	5	6	6	1	3
M10	6	3	8	15	5	9	12	3	1	[6]	1	3	8	5	1
M11	12	5	1	6	3	8	15	3	9	6	[5]	1	8	3	1
M12	15	6	9	12	5	1	3	1	8	3	5	[6]	8	1	3

Table 2c Value Chart for MH-VL-S Condition

	w1	w2	w3	w4	w5
M1	[7]	10	0	4	2
M2	10	[7]	4	0	2
M3	10	0	[2]	4	7
M4	0	10	4	[2]	7

Table 2d Key Features of Experimental Conditions

Condition	Average correlation of preferences	Prioritizing method	Market size	Pareto-optimal total earnings of the group	Initial efficiency (%)
MH-VL (base)	0.040	Value-linked	12 members 15 weeks	149.5	60.2
LH-VL	0.264	Value-linked	12 members 15 weeks	149.0	60.4
MH-R	0.040	Random	12 members 15 weeks	149.5	60.2
MH-VL-S	0.048	Value-linked	4 members 5 weeks	30.0	60.0

efficiency is 60%. The average correlation of preferences is 0.048, indicating a very heterogeneous preference structure. We summarize the four conditions in Table 2d.

Procedure

In each session, 12 (or 4, depending on the condition) participants were seated in front of networked computer terminals. We started with experimental instructions to educate them about the timeshare industry and the exchange mechanism that was implemented

in that specific session.¹³ The participants then made one set of exchanges for practice purposes, followed by six more sets of exchanges. Each set of exchanges simulated one market. At the beginning of each market, every participant was randomly assigned a role among M1, . . . , M12 (or M1, . . . , M4, depending on

¹³ Exact instructions for participants (for the base condition) are available in an online technical appendix to this paper on the *Management Science* website at <http://mansci.pubs.informs.org/ecompanion.html>.

the condition), and learned privately about the value of all the potentially available weeks for her alone. For example, if a participant in the base condition was assigned as M4, then her initial endowment was w_4 and she saw the fourth row of Table 2a. Thus, the preferences of the participants were manipulated by the experimenter. Before they started, the participants were told that they would make exchanges in six independent markets, and that their “role” in each market was assigned randomly—thus, participants believed that in each set of exchanges they were in a different market that was unrelated to earlier ones. Participants then made exchanges according to the requirement of the mechanism that was in use. They were asked to make decisions privately, and no communication of any form was allowed. After each market, participants learned (only) about their own allocation. At the end, the participants were compensated based on their performance, debriefed, thanked, and dismissed.

The interactive custom software needed a network of computers where data were inputted from all participants and individual allocations based on all these data were instantly determined. Further, the software kept track over all markets to determine each participants’ final payment. Each mechanism (and to a lesser extent each condition) needed a different software program.

Predicted Efficiency and Participation

We expect better performance of the TTCCS mechanism compared to DF and RF. In specific, theoretically the overall efficiency of the final allocation is 100% for TTCCS, whereas for RF and DF it is less than 100%. Hence, TTCCS is predicted to be more efficient than RF and DF. Theoretically, all members should be willing to participate in RF and TTCCS, so the participation rate is 100% for RF and TTCCS, whereas it is less than 100% for DF. Thus, RF and TTCCS are predicted to elicit greater participation rate than DF. We emphasize that the theory applies with any priority list that is exogenously given to our model.

Experimental Results

Within each condition we compare the efficiency of the outcomes and the behavior of the participants across sessions. Specifically, the *observed efficiency* of each market is measured by the ratio of the sum of actual earnings of all participants to the Pareto-optimal total earning of the group (which is consistent with extant literature such as Chen and Sonmez 2002). Another important dependent variable is *participation rate*, which is measured by the percentage of participants who decide to enter the system and make exchanges. Under DF and RF, a participant is considered to have entered the market as long as she does not choose to opt out for all the five periods of the market. We report

the observed efficiency and the participation rate of the three mechanisms in Tables 3a and 3b.

In Market 5 of the base condition, for example, we find that TTCCS yields 100% efficiency compared to the Pareto-optimal outcomes, whereas the RF mechanism yields 72.9% efficiency. In other words, the total number of points earned by all 12 members in the RF market was 72.9% of the maximum they could have earned.¹⁴

Efficiency. For each session we pool data from the six sets of exchanges. The Mann-Whitney test¹⁵ shows that in all four conditions, the observed efficiency of TTCCS is significantly higher than that of DF ($p < 0.05$, one-tailed test), and that of RF ($p < 0.05$, one-tailed test).

We find that participants seem to learn fairly quickly (i.e., over one or two markets) what is the best strategy for them to adopt under a certain mechanism—thus, if a participant deposited her home week in the first set of exchanges under DF and did not fare well in her allotment, in further markets she tended not to deposit her home week if it was already fairly satisfactory. This accounts for participation and thus efficiency decreasing over time under DF in the MH-R condition. Also, note that in two of the four conditions TTCCS does not generate the most efficient outcome in the first set of exchanges. However, over one or two markets TTCCS stabilizes as the most efficient mechanism among the three.

Interestingly, although we do not have a decisive theoretical conclusion as to whether DF or RF is less efficient, RF generates worse results than DF in most cases. We speculate that because deposits are made before withdrawals in DF, there tend to be more weeks in the spacebank which facilitate exchanges. In RF there are fewer weeks in the spacebank, so the whole system is more likely to get “stuck,” in the sense that fewer mutually beneficial exchanges are identified and carried out by the system.

Participation. As predicted, the participation rate in the RF and TTCCS mechanisms is considerably higher than in the DF mechanism (Table 3b). Participants quite quickly figured out that their best strategy under RF and TTCCS was to participate in the exchange market, no matter what week they were assigned at the beginning. This explains why under RF and TTCCS the participation rate converges to 100% very quickly, but in DF it never reaches 100%. Because we are interested in the general result, for each session we pool

¹⁴ Note that whenever the observed total payoff is greater than the average Pareto-optimal outcome, we truncate the ratio and report 100% as the observed efficiency.

¹⁵ The Mann-Whitney test is a nonparametric alternative to the classical students *t*-test, used for comparing the central tendency of two independent random samples. (See, e.g., Gibbons 1993, pp. 30–42.)

Table 3a Observed Efficiency of the Markets

Mechanism	Market no.	More heterogeneous, value-linked condition						Less heterogeneous, value-linked condition					
		1	2	3	4	5	6	1	2	3	4	5	6
DF	Total payoff	133.0	134.0	128.0	124.0	130.0	135.0	124.0	128.0	117.7	126.0	132.0	118.0
	Efficiency (%)	89.0	89.6	85.6	82.9	87.0	90.3	83.2	85.9	79.0	84.6	88.6	79.2
RF	Total payoff	112.0	122.7	110.0	108.7	109.0	125.0	105.0	130.4	122.7	132.7	109.0	105.7
	Efficiency (%)	74.9	82.1	73.6	72.7	72.9	83.6	70.5	87.5	82.3	89.1	73.2	70.9
TTCCS	Total payoff	128.4	142.7	148.4	151.7	154.0	154.0	140.7	140.0	153.0	154.7	140.0	156.7
	Efficiency (%)	85.9	95.5	99.3	100.0	100.0	100.0	94.4	94.0	100.0	100.0	94.0	100.0
		More heterogeneous, random condition						Small-sized condition					
DF	Total payoff	135.0	136.0	130.0	133.0	120.0	108.7	28.0	25.0	25.0	28.0	25.0	25.0
	Efficiency (%)	90.3	91.0	87.0	89.0	80.3	72.7	93.3	83.3	83.3	93.3	83.3	83.3
RF	Total payoff	121.0	124.0	107.7	106.7	107.7	109.0	25.6	25.0	25.6	25.0	25.6	25.6
	Efficiency (%)	80.9	82.9	72.0	71.4	72.0	72.9	85.3	83.3	85.3	83.3	85.3	85.3
TTCCS	Total payoff	139.0	138.0	140.0	154.0	154.7	155.0	25.0	31.6	31.0	31.0	31.0	31.6
	Efficiency (%)	93.0	92.3	93.6	100.0	100.0	100.0	83.3	100.0	100.0	100.0	100.0	100.0

Table 3b Participation Rate of the Markets

Mechanism	Market no.	More heterogeneous, value-linked condition (%)						Less heterogeneous, value-linked condition (%)					
		1	2	3	4	5	6	1	2	3	4	5	6
DF	Participation	83.3	83.3	75.0	75.0	83.3	75.0	75.0	83.3	75.0	75.0	83.3	75.0
RF	Participation	83.3	83.3	100.0	100.0	100.0	100.0	83.3	100.0	100.0	100.0	100.0	100.0
TTCCS	Participation	83.3	100.0	100.0	100.0	100.0	100.0	91.7	91.7	100.0	100.0	91.7	100.0
		More heterogeneous, random condition (%)						Small-sized condition (%)					
DF	Participation	83.3	91.7	83.3	83.3	66.7	75.0	75.0	50.0	75.0	75.0	50.0	75.0
RF	Participation	91.7	91.7	91.7	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
TTCCS	Participation	100.0	100.0	100.0	100.0	100.0	100.0	75.0	100.0	100.0	100.0	100.0	100.0

Table 3c Proportions of Participants Who Got Better, Same, or Worse Weeks (Averages Across Six Markets in Each Session)

	More heterogeneous, value-linked condition	Less heterogeneous, value-linked condition	More heterogeneous, random condition	Small-sized condition
DF				
% better-off	68.1	63.9	63.9	54.2
% same	27.8	29.2	26.4	37.5
% worse-off	4.2	6.9	9.7	8.3
RF				
% better-off	54.2	58.3	45.8	50.0
% same	38.9	41.7	48.6	50.0
% worse-off	6.9	0.0	5.6	0.0
TTCCS				
% better-off	82.0	86.1	80.6	91.7
% same	13.9	13.9	18.1	8.3
% worse-off	4.2	0.0	1.4	0.0

data from the six sets of exchanges and then compare the three sessions in the same condition. A t -test of proportions shows that, in each condition, the participation rate of RF is significantly higher than that of DF ($p < 0.01$, two-tailed. For the base condition, $t = 2.78$; for the LH-VL condition, $t = 3.69$; for the MH-R condition, $t = 2.92$; for the MH-VL-S condition, $t = 3.46$); the participation rate of TTCCS is significantly higher than that of DF ($p < 0.01$, two-tailed. For the base condition, $t = 3.50$; for the LH-VL condition, $t = 3.32$; for the MH-R condition, $t = 4.17$; for the MH-VL-S condition, $t = 2.79$); and the difference between TTCCS and RF is not statistically significant ($p > 0.3$, two-tailed), except in the MH-R condition where the difference between TTCCS and RF is marginally significant ($p = 0.08$, two-tailed $t = 1.77$). However, note that by the fourth set of exchanges, participation under RF is 100% in all conditions.

Note that for each session of DF, participation and efficiency show similar patterns (i.e., higher participation rate is accompanied by higher efficiency and vice versa). This complies with theory because the main deficiency of DF is low participation rate—when participation is high, the efficiency of the system should be improved. However, this is not always true for RF sessions, where in the MH-R condition efficiency is low although there is 100% participation. This is not surprising, because when members with better weeks and worse weeks compete for the same popular week, there are two possible results under RF: (i) when a member with a better week gets it, her original home week (which is in the spacebank now) is also popular so others will want it—the exchanges get going and many members come out better off; (ii) when a member with a less popular week gets it, her original (less popular) home week is now in the spacebank but nobody wants it. Therefore the system gets “stuck” here and the overall efficiency is low. Therefore (especially in the random condition), the outcome of RF is largely influenced by this “randomness,” even though participation is 100%.

In Table 3c, we report the proportion of participants who got better, same, or worse (than home) weeks in each session (for brevity, we report averages across the six markets for each session). Again, TTCCS generates better results because a higher percentage of participants got *better weeks*, compared to RF and DF, in all four conditions. A t -test of proportions (pooled across the six sets of exchanges in each session) shows that for TTCCS the percentage of participants who got better weeks is significantly higher than that of DF and RF ($p < 0.05$, two-tailed. TTCCS versus DF: for the LH-VL condition, $t = 3.19$; for the MH-R condition, $t = 2.27$; for the MH-VL-S condition, $t = 3.22$. TTCCS versus RF: for the base condition, $t = 3.75$; for the LH-VL condition, $t = 3.91$; for the

MH-R condition, $t = 4.63$; for the MH-VL-S condition, $t = 3.57$), except in the base condition, where the difference between TTCCS and DF is at a marginal level ($p = 0.053$, $t = 1.95$).

We observe that a few TTCCS participants ended up with worse weeks than their home weeks, a result not consistent with theory. Our belief is that a few participants were not able to comprehend the TTCCS mechanism, especially in the first several sets of exchanges. Thus, two participants misrepresented their preferences and got worse weeks as a result. Note, however, that TTCCS still generated much better efficiency in general. We also cannot rule out the possibility that a participant accidentally made a mistake and so ended up with a worse week.

The outcomes of the simulated exchange markets provide strong support for the theory. With two different preference structures, two different methods of prioritizing, and two different market sizes, we observed that TTCCS results in tremendous improvement in efficiency over DF and RF.

6. General Discussion and Future Research

We bring attention to the timeshare industry by identifying it as a new application for mechanism design, characterizing existing systems within a mechanism-design framework, seeing if an alternate mechanism will perform better and designing such a variant (TTCCS), and then theoretically proving that the current systems are inferior and the newly designed variant is superior in efficiency. We modify the TTCC mechanism into TTCCS to incorporate the unique feature of timeshare exchange systems—the availability of the spacebank. Empirically, our study builds on the existing timeshare exchange methods, retaining their merits and discarding their deficiencies—exchanges associated with cycles may be carried out under the DF mechanism, but the mechanism is not individually rational, resulting in a low participation rate; on the other hand, the RF mechanism is individually rational, but cycles in the system may remain undetected. The TTCCS mechanism reduces the risk involved in timeshare exchanges and leads to a higher participation rate compared to the DF mechanism. It also detects all trading cycles and conducts exchanges accordingly.

Our model assumes fixed preferences at a specific point in time. However, this does not preclude that preferences may change over time. For example, the value of the waiting queue may decline as timeshare members get closer to their vacation period, because there is less time for more desirable weeks to come in. Thus, the ranking of the waiting queue relative to available weeks may decline. However, changing preferences do not affect the implementation of TTCCS because members can be allowed to change their rank-order lists from time to time.

The timeshare exchange problem can be more complicated than what we have presented if we take into consideration some complexities in the real-world setting. A few issues are worth discussing. For example, although in our model each member only owns one week, in practice some households own several weeks at the same time. Under some conditions, this can be easily accommodated into our model—we can consider such a household as several independent members, each having only one home week. As long as the vacation decisions are made separately for these several weeks, such separation makes sense, but obviously, we may not be able to do so if the weeks are consecutive in time at the same resort accommodation. Special cases like this necessitate extensions of our theoretical model. Similarly, there may be value tiers and members may only be allowed exchanges within their own tier. Our mechanism can be adapted for this change, a simple solution being to treat each value tier as a separate exchange system. When members are allowed exchanges within multiple tiers, the problem is more complicated and model extensions will be needed. Future research can also consider extensions of the model with endogenous entry of timeshare members.

Another issue is the feeling of control. In the DF and RF mechanisms members may feel that they are more involved in the exchanges and get a sense of control over exchanges. Under TTCCS, they may tend to blame the exchange company if they do not get a preferred allocation (rather than attributing it to a low-value week that they may own). In a similar vein, under TTCCS, if members forget to submit their rank-order lists before their home weeks come up, they will automatically be assigned their home weeks. However, they may blame the timeshare exchange company for not sending them enough reminders.

A major objective of this research is to show the relevance of mechanism-design work to business applications and to encourage research in the mechanism design area within management science, the area being studied mostly by economists thus far. Further, we focus on an industry not studied earlier within academia, namely the timeshare industry, and show problems with the current timeshare exchange mechanisms and propose an alternative one that we demonstrate to be superior both theoretically and empirically.

An online supplement to this paper is available on the *Management Science* website at <http://mansci.pubs.informs.org/ecompanion.html>.

Acknowledgments

The research is funded by a University of Michigan Doctoral Student Research Grant and a Faculty Research Grant. The authors would like to thank Yan Chen for her help throughout this project. The paper also benefited from comments of participants at the Marketing Science conference 2004,

INSEAD winter camp 2005, and the NUS speaker series 2005. The authors are grateful to Yang Peng for programming for this project.

References

- Abdulkadiroglu, A., T. Sonmez. 1999. House allocation with existing tenants. *J. Econom. Theory* **88**(2) 233–260.
- Amaldoss, W., S. Jain. 2002. David vs. Goliath: An analysis of asymmetric mixed-strategy games and experimental evidence. *Management Sci.* **48**(8) 972–991.
- Amaldoss, W., R. J. Meyer, J. S. Raju, A. Rapoport. 2000. Collaborating to compete. *Marketing Sci.* **19**(2) 105–126.
- Chen, Y., T. Sonmez. 2002. Improving efficiency of on-campus housing: An experimental study. *Amer. Econom. Rev.* **92**(5) 1669–1689.
- Gibbons, J. D. 1993. *Nonparametric Statistics: An Introduction*. Sage Publications, Inc., Newbury Park, CA.
- Haruvy, E., I. Erev, D. Sonsino. 2001. The medium prizes paradox: Evidence from a simulated casino. *J. Risk Uncertainty* **22**(3) 251–261.
- Ho, T., K. Weigelt. 1996. Task complexity, equilibrium selection, and learning: An experimental study. *Management Sci.* **42**(5) 659–679.
- Kagel, J. H., D. Levin. 1993. Independent private value auctions: Bidder behavior in first-, second- and third-price auctions with varying numbers of bidders. *Econom. J.* **103**(419) 868–879.
- Kagel, J. H., A. E. Roth. 2000. The dynamics of reorganization in matching markets: A laboratory experiment motivated by a natural experiment. *Quart. J. Econom.* **115**(1) 201–235.
- Krishna, A., M. U. Unver. 2004. Efficiency enhanced course bidding: Evidence from a field experiment at the University of Michigan Business School. Working paper, University of Michigan, Ann Arbor, MI.
- Krishna, A., Y. Wang. The relationship between top trading cycles mechanism and top trading cycles and chains mechanism. *J. Econom. Theory*. Forthcoming.
- Ragatz Associates, Inc. 1999. The economic benefits of RCI exchange affiliation: 1999 executive edition. Report, Ragatz Associates, Inc., Eugene, OR. <http://www.rciaffiliates.com/industryoverview/industrystudies.asp>.
- Resort Condominiums International. 2003. Resort timesharing worldwide: 2003 edition. Summary report, RCI, Parsippany, NJ. <http://www.rciaffiliates.com/industryoverview/industrystudies.asp>.
- Ross, L. F., E. S. Woodle. 2000. Ethical issues in increasing living kidney donations by expanding kidney paired exchange programs. *Transplantation* **69** 1539–1543.
- Roth, A. E., E. Peranson. 1997. The effects of the change in the NRMP matching algorithm. *J. Amer. Medical Assoc.* **278**(9) 729–732.
- Roth, A. E., T. Sonmez, M. U. Unver. 2004. Kidney exchange. *Quart. J. Econom.* **119**(2) 457–488.
- Shapley, L. S., H. E. Scarf. 1974. On cores and indivisibility. *J. Math. Econom.* **1**(1) 23–37.
- Sonmez, T., M. U. Unver. 2005. House allocation with existing tenants: An equivalence. *Games Econom. Behavior* **52** 153–185.
- Srivastava, J., D. Chakravarti, A. Rapoport. 2000. Price and margin negotiations in marketing channels: An experimental study of sequential bargaining under one-sided uncertainty and opportunity cost of delay. *Marketing Sci.* **19**(2) 163–184.
- Zenios, S., E. S. Woodle, L. F. Ross. 2001. Primum non nocere: Avoiding increased waiting times for individual racial and blood-type subsets of kidney wait list candidates in a living donor/cadaveric donor exchange program. *Transplantation* **66** 406–407.
- Zwick, R., A. Rapoport, A. K. C. Lo, A. V. Muthukrishnan. 2003. Consumer sequential search: Not enough or too much? *Marketing Sci.* **22**(4) 503–519.